universität wien

MASTERARBEIT | MASTER'S THESIS

Titel | Title

Optimal Time-Temperature Curves for Heat-Treatable Aluminum Alloys

verfasst von | submitted by

Bram Bastiaansen

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2025

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 821
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Mathematik
Betreut von | Supervisor:	Ass.-Prof. Yurii Malitskyi PhD

Acknowledgements

This thesis was carried out within the Complex Dynamical Systems research group at the Center for Vision, Automation & Control of the Austrian Institute of Technology (AIT). I would like to take this opportunity to express my sincere gratitude to everyone who supported and contributed to this work. First and foremost, I am deeply grateful to my colleagues and supervisors at AIT, Stephan Strommer and Lukas Grohmann, whose guidance and advice were essential in shaping this thesis. I would also like to thank Prof. Yura Malitskyi, who supervised my work in his official capacity as representative of the University of Vienna. His constructive feedback helped resolve several inconsistencies and strengthened the overall quality of this research. I am further indebted to all those who took the time to review and comment on my work, providing valuable insights and corrections. Finally, I would like to extend my heartfelt gratitude to my parents, whose support made this journey possible in the first place.

Abstract

Aluminum alloys are widely used in industrial settings for their strength and corrosion resistance, and can be made even stronger through heat-treatment. Efficient heat treatment of aluminum alloys requires balancing mechanical performance, energy consumption, and processing time. This thesis investigates the optimal control of time-temperature profiles in an industrial furnace (model N 120/85 HA) for the heat-treatable alloy EN AW-6082 (Al-Mg-Si). A transient, physics-based model of the furnace is developed, capturing convective and radiative heat transfer, transient conduction through multilayered furnace walls, and the thermal interactions between furnace components. The model is implemented in MATLAB using an implicit time integration scheme. Calibration against experimental measurements demonstrates close agreement between predicted and observed temperatures, confirming the model's predictive capability. Optimization over furnace air time-temperature curves is performed to minimize energy use and process duration while achieving target mechanical properties, with surrogate models incorporated to represent tensile strength, yield strength, and elongation. The proposed solution strategy successfully generates feasible temperature profiles across multiple operating scenarios, showing consistent, physically reasonable outcomes. These results demonstrate that the combination of calibrated modeling and trajectory-based optimization provides a reliable and flexible framework for efficient, data-driven heat treatment of aluminum alloys in industrial settings.

Zusammenfassung

Aluminiumlegierungen werden aufgrund ihrer Festigkeit und Korrosionsbeständigkeit häufig in industriellen Anwendungen eingesetzt und können durch Wärmebehandlung weiter verstärkt werden. Eine effiziente Wärmebehandlung von Aluminiumlegierungen erfordert die Abwägung zwischen mechanischer Leistungsfähigkeit, Energieverbrauch und Prozessdauer. In dieser Arbeit wird die optimale Steuerung von Zeit-Temperatur-Profilen in einem Industrieofen (Modell N 120/85 HA) für die aushärtbare Legierung EN AW-6082 (Al-Mg-Si) untersucht. Es wird ein transienter, physikalisch basierter Ofenmodellansatz entwickelt, der konvektiven und radiativen Wärmetransport, transiente Wärmeleitung durch mehrschichtige Ofenwände sowie die thermischen Wechselwirkungen zwischen den Ofenkomponenten berücksichtigt. Das Modell wird in MATLAB unter Verwendung eines impliziten Zeitschrittverfahrens implementiert. Die Kalibrierung anhand experimenteller Messungen zeigt eine enge Übereinstimmung zwischen den vorhergesagten und beobachteten Temperaturen und bestätigt damit die Prognosefähigkeit des Modells. Die Optimierung über zeitabhängige Ofenluft-Temperaturverläufe wird durchgeführt, um den Energieverbrauch und die Prozessdauer zu minimieren und gleichzeitig die geforderten mechanischen Eigenschaften zu erreichen. Hierfür werden Surrogatmodelle zur Abbildung von Zugfestigkeit, Streckgrenze und Dehnung eingesetzt. Die vorgeschlagene Lösungsstrategie erzeugt erfolgreich realisierbare Temperaturprofile für verschiedene Betriebsszenarien und liefert konsistente, physikalisch plausible Ergebnisse. Die Ergebnisse zeigen, dass die Kombination aus kalibriertem Modell und trajektorienbasierter Optimierung einen zuverlässigen und flexiblen Rahmen für eine effiziente, datengestützte Wärmebehandlung von Aluminiumlegierungen in industriellen Anwendungen bietet.

Contents

1 Introduction

1.1 Background and Motivation

Aluminum alloys are widely used in aerospace, automotive, and construction industries due to their high strength-to-weight ratio, corrosion resistance, and ease of fabrication [48]. Within this class of materials, heat-treatable alloys are of particular importance because their mechanical properties can be significantly enhanced through carefully controlled thermal processing [45]. Heat treatment of aluminum alloys requires careful control of complex thermal interactions between the furnace walls, heating elements, circulating air, and the alloy itself. The final performance of these alloys depends strongly on the resulting temperature trajectory.

In industrial practice, heat treatment is both energy-intensive and time-consuming. Standardized and predetermined schedules are typically employed to guarantee product quality under a wide range of conditions, but such approaches are often conservative and do not fully account for real-time furnace dynamics or variability in alloy response [18, 31]. This can lead to excessive power consumption, unnecessarily long processing times, and, in some cases, less-than-optimal material properties [8, 38]. Recent studies highlight the need for adaptive and optimized heat treatment strategies that improve efficiency without sacrificing quality [56, 25, 47].

The challenge, therefore, lies in determining furnace operating strategies that simultaneously reduce energy use and process time while ensuring that the treated alloy meets its mechanical property requirements. This requires capturing the transient thermal interactions between the furnace components and the alloy, and using this knowledge to design improved temperature trajectories.

This thesis addresses that challenge by developing a physics-based model of the N 120/85 HA furnace model, calibrating it against experimental data, and applying mathematical optimization techniques to explore and design time-temperature curves that reduce energy use and processing time without compromising the desired mechanical properties. Rather than providing a single prescriptive solution, the work documents the modeling, calibration, and optimization steps taken toward systematically solving this problem. This thesis was carried out in collaboration with, and under the supervision of, the Austrian Institute of Technology (AIT). Experimental data were provided by Leichtmetallkompetenzzentrum Ranshofen GmbH (LKR), a subsidiary of AIT.

1.2 Objectives and Scope

The primary objective of this thesis is to develop a model-based optimization framework for determining the optimal time-temperature curve of heat-treatable aluminum alloys in an industrial

furnace. This work forms part of a larger project aimed at introducing optimal control methods for achieving high-performance aluminum alloys through predictive modeling and process optimization.

In the context of industrial furnace optimization, the connection between the time-temperature curve and material microstructure is critical. The formation and evolution of precipitates directly dictate mechanical outcomes, meaning that small changes in furnace temperature profiles can significantly affect strength and ductility. By integrating predictive models of mechanical properties with furnace process control, it becomes possible to design thermal schedules that simultaneously minimize energy consumption, reduce processing time, and achieve target material performance (Figure 1.1).

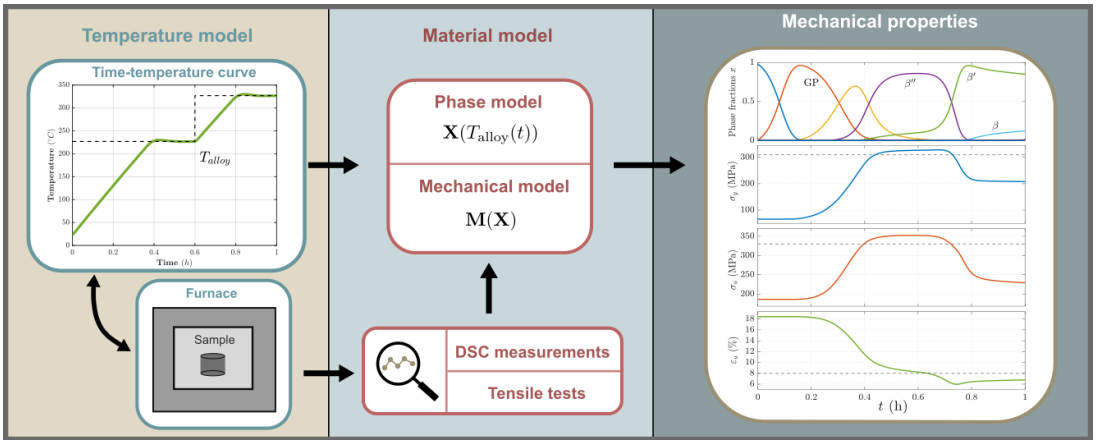

Figure 1.1: Underlying model of the optimization

The specific goals of this thesis are:

1. To develop a comprehensive transient furnace model that explicitly represents the major thermal domains, including the multi-layered wall structure, circulating furnace air, the alloy component, and radiative heating assemblies, with appropriate coupling of conduction, convection, and radiation mechanisms.

2. To implement robust numerical methods, such as finite volume spatial discretization combined with implicit time integration and iterative nonlinear solvers, in order to simulate the coupled thermal response of the furnace system with stability and accuracy over industrially relevant timescales.

3. To calibrate and validate the furnace model through comparison with experimental measurements of temperature evolution and power consumption, thereby ensuring predictive capability under realistic process conditions.

4. To formulate and solve a multi-objective optimization problem that jointly minimizes energy consumption and process duration while maximizing (or minimizing deviations from) target mechanical properties.

The scope of this thesis is limited to simulation-based optimization of a specific furnace model (N 120/85 HA). Given the diversity and complexity of industrial furnaces, the methodology developed here may be informative for other configurations but is not intended as a universally applicable framework. While full-scale industrial deployment lies beyond the scope of this work, the inclusion of experimental measurements aims to demonstrate the practical relevance and adaptability of the model and optimization approach.

The computation of alloy phase transformations and mechanical properties are not addressed in this work. These aspects are captured by a thermokinetic and mechanical property model developed at the Austrian Institute of Technology (AIT) based on atomistic simulations [46]. Within this framework, the evolution of phase fractions is represented by the matrix $\mathbf{X}$, while the resulting mechanical property model is denoted by $\mathbf{M}$. Figure 1.2 illustrates the overall project structure. Within the project, this theses focuses specifcally on the thermal modeling of the furnace system, the simulation of time-temperature curves, and the exploration of temperature control strategies to achieve target alloy properties as predicted by $\mathbf{X}$ and $\mathbf{M}$.

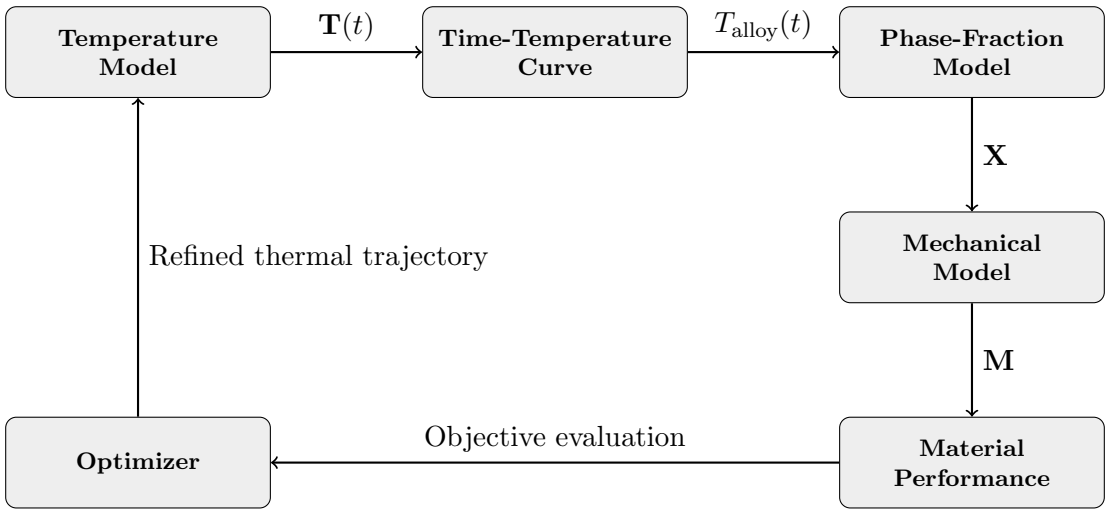

Figure 1.2: Flow diagram of project structure.

1.3 Structure of the Thesis

This thesis is organized into nine chapters, each addressing a distinct aspect of the research:

- **Chapter 1 - Introduction:** Outlines the research motivation, defines the problem statement, establishes the scope of the study, and presents an overview of the thesis structure.
- **Chapter 2 - Heat-Treatable Aluminum Alloys:** Provides background on aluminum alloys, including their classification, heat treatment processes, and the evolution of mechanical properties.
- **Chapter 3 - Furnace Model:** Describes the physical configuration of the furnace system,

identifies the dominant modes of heat transfer, and introduces the key assumptions applied to simplify the thermal dynamical modeling framework.

- **Chapter 4 - Mathematical Formulation:** Develops the governing equations of the system, including energy balances, heat transfer methods, and boundary conditions for each thermal domain.
- **Chapter 5 - Numerical Methods:** Details the finite volume discretization strategy, implicit time integration scheme, and nonlinear solution procedures used to capture the transient thermal response of the furnace and alloy.
- **Chapter 6 - Model Calibration and Validation:** Presents the calibration of model parameters and validation of results against experimental measurements. The chapter assesses model accuracy through comparisons of predicted and measured thermal histories.
- **Chapter 7 - Optimization Framework:** Defines the optimization problem, including the formulation of the objective function, constraints, and solution strategy for identifying optimal furnace time-temperature curves.
- **Chapter 8 - Results and Discussion:** Reports the outcomes of simulation and optimization studies, including analyses of temperature evolution, energy consumption, and mechanical performance. The implications of the findings are discussed in the context of industrial applicability.
- **Chapter 9 - Conclusions and Outlook:** Summarizes the main contributions of the work, reflects on the limitations of the current study, and outlines directions for future research and model development.

Together, these chapters provide a coherent framework that integrates physical modeling, numerical simulation, model calibration, and optimization strategies to improve the efficiency and effectiveness of industrial heat treatment processes for aluminum alloys. All custom MATLAB scripts developed for this work, including the physical furnace model implementation and the optimization routines, are available in a public Github repository (see Appendix 9.4).

2 Heat-Treatable Aluminum Alloys

Heat-treatable aluminum alloys play a crucial role in modern engineering applications, from the building of bridges and buildings to aeroplanes, where specific mechanical performance is required [33, 48, 19]. Unlike pure aluminum, which is generally not suitable for industrial applications due to its low strength and limited mechanical potential, these alloys can have their microstructure—and therefore their mechanical behavior—significantly altered through controlled thermal processing [52]. Understanding this connection between thermal history and material response is essential for both design and manufacturing optimization. This chapter introduces the classification of aluminum alloys, the relevant heat-treatment processes, and the evolution of mechanical properties, with a focus on their relevance to furnace process optimization.

2.1 Classification of Aluminum Alloys

Aluminum alloys are generally classified into two broad categories: *heat-treatable* and *non-heat-treatable*. Non-heat-treatable alloys achieve their strength primarily through mechanical work (cold working), while heat-treatable alloys rely on controlled thermal processing to develop strength through microstructural changes [40].

Heat-treatable alloys, such as those in the 2xxx (Al-Cu), 6xxx (Al-Mg-Si), and 7xxx (Al-Zn-Mg) series, are of particular interest in high-performance engineering because their strength can be significantly enhanced through heat treatment. This precipitation-hardening mechanism is central to their mechanical performance and distinguishes them from non-heat-treatable alloys. In this work, the investigated material is the aluminum alloy EN AW-6082 (Al-Mg-Si). Its chemical composition (in terms of weight percentage) is characterized by the distribution of alloying elements as shown in Table 2.1.

Element	Si	Fe	Cu	Mn	Mg	Cr	Zn	Ti	Al
Content (wt.-%)	0.89	0.46	0.05	0.46	0.64	0.01	0.03	0.02	bal.

Table 2.1: Chemical composition of aluminum alloy EN AW-6082.

2.2 Heat Treatment Process

The heat treatment of aluminum alloys generally consists of three main stages [45, 8]. Each stage changes the internal structure of the metal to improve its strength and durability.

2.2.1 Solution Heat Treatment

In this stage, the alloy is heated to a high temperature (typically 450-550°C) so that the alloying elements—such as magnesium and silicon—fully dissolve into the aluminum base metal. At this temperature, the alloy becomes a single, uniform solid solution, meaning that the atoms of the added elements are evenly distributed within the aluminum's crystal structure (often called the matrix). After heating, the material must be cooled quickly enough to prevent these elements from separating out again. The minimum rate at which this cooling must occur is called the critical cooling rate. If cooling is too slow, unwanted particles or phases can form, reducing the alloy's strength potential.

2.2.2 Quenching

The material is rapidly cooled, typically in water or with forced air. This quick cooling "freezes" the dissolved atoms in place, creating a supersaturated solid solution—a temporary, unstable structure that stores a lot of potential strength. However, this rapid temperature drop can also cause residual stresses and distortion, since the outside of the part cools and contracts faster than the inside. Balancing the quenching speed is therefore critical to achieving good mechanical properties without excessive warping.

2.2.3 Artificial Aging

Following quenching, the alloy is reheated to an intermediate temperature (typically 150-200°C) and held for a defined period. During this time, the trapped alloying atoms slowly move (diffuse) through the metal and cluster together into tiny, evenly distributed particles, called precipitates. These precipitates block the motion of dislocations—small defects that allow metals to deform—thereby making the material harder and stronger. This process is known as precipitation hardening or age hardening [46]. The temperature and duration of this stage determine how the precipitates form and grow. Their size, spacing, and composition directly influence the final mechanical properties of the alloy, which is why precise control of the time–temperature curve is essential during aging.

2.3 Mechanical Properties

The mechanical properties of heat-treatable alloys, such as yield strength, tensile strength, and elongation, depend on the distribution, size, and type of precipitates formed during aging. The evolution of these properties is highly non-linear and strongly coupled to the time-temperature curve [57, 38]. There are other properties—like hardness, fracture toughness, and fatigue—but they are considered outside the scope of the optimization framework.

In this thesis, the analysis is restricted to three commonly characterized properties: the yield strength (σ_y), which is the stress at which permanent deformation begins; the ultimate tensile strength (σ_u), which represents the maximum stress the material can withstand before failure; and the uniform elongation (ε_u), defined as the strain at the onset of necking (i.e. the maximum uniform plastic deformation the material can undergo before localized instability). Figure 2.1 gives the relationship between stress (applied pressure) and strain (amount of deformation),

highlighting the mechanical properties.

Experimental measurement of these properties after heat treatment serves both as a benchmark for validation and as input for predictive models. As described by Singh and Warner [46], atomistic-based multiscale modeling can capture the evolution of precipitate distributions and their effect on macroscopic mechanical properties. Building on this approach, the mechanical model **M** and phase-fraction model **X** can predict these properties for Al-Mg-Si alloys under arbitrary time-temperature curves, providing a direct link between furnace operation and resulting alloy performance.

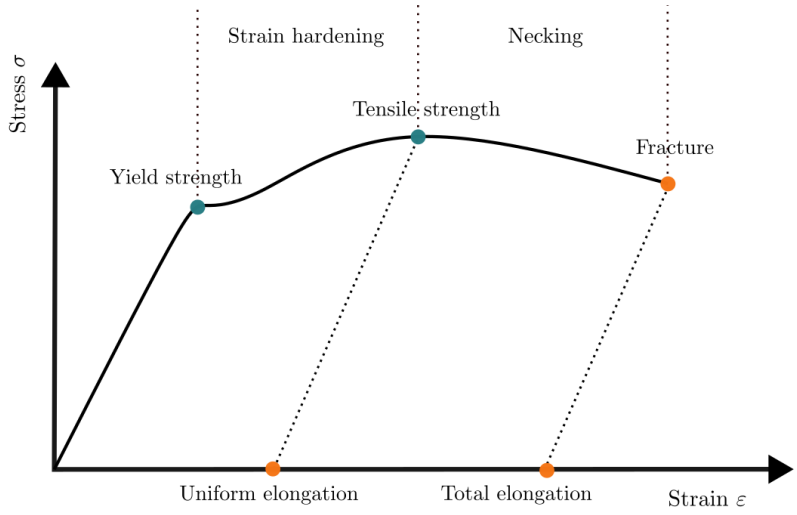

Figure 2.1: Stress-strain curve

The alloy initially exists in a supersaturated solid solution. Upon aging, small clusters of solute atoms form known as Guinier-Preston (GP) zones [37]. During these stages, yield and tensile strength begin to increase. With further heat treatment, these clusters transform into the β'' phase, consisting of predominantly needle-shaped, coherent precipitates with a nearly perfect lattice structure. This metastable phase is strongly correlated with the peak in both yield strength and ultimate tensile strength, and is therefore of particular importance for time-temperature curve optimization. As the alloy continues to evolve, the β'' phase decomposes into the semi-coherent β' phase. These larger precipitates are less effective in hindering dislocation motion, leading to a gradual reduction in strength. Finally, the stable equilibrium β phase is formed, consisting of large, incoherent precipitates. At this stage, strength decreases further, while ductility recovers, reflecting the trade-off between strength and elongation.

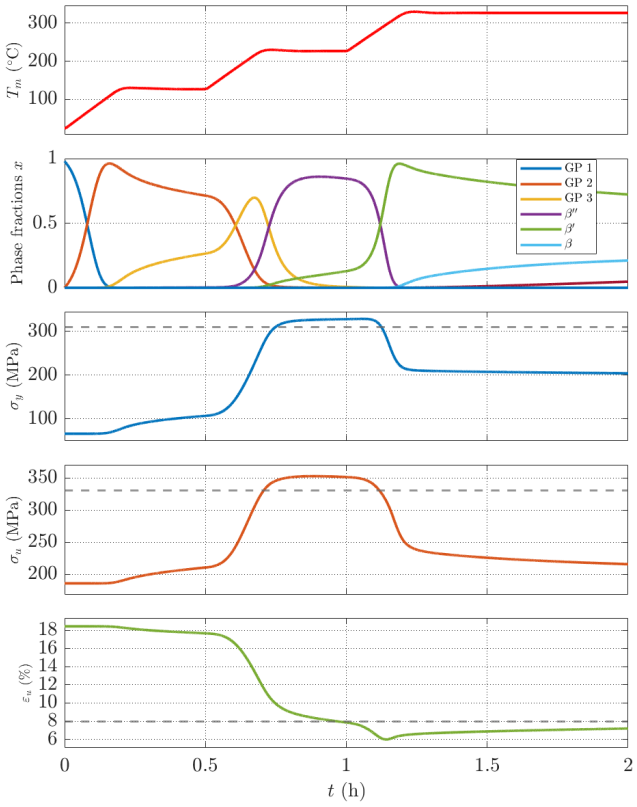

Figure 2.2: Predicted phase fractions and mechanical property evolution for a heat-treatable aluminum alloy.

Figure 2.2 illustrates a representative prediction of phase evolution and its effect on mechanical properties over time. The upper figure corresponds to the imposed time-temperature curve of the alloy. Immediately below, the evolution of phase fractions (denoting the relative percentage of each distinct type of precipitate) is shown, with particular emphasis on the β'' phase (purple). This phase is strongly correlated with the peak in both yield strength σ_y and ultimate tensile strength σ_u. As the β'' phase decomposes, strength decreases correspondingly. Meanwhile, uniform elongation ε_u shows a monotonic decline at elevated temperatures, reflecting the trade-off between strength and ductility. The dashed horizontal lines denote standard minimum performance requirements.

3 Physical and Thermal Model of the Furnace

In many industrial processes, furnaces play a critical role in heating materials for manufacturing and metallurgical applications. Optimizing furnace operation requires a thorough understanding of the underlying heat transfer mechanisms, which directly affect energy consumption, product quality, and overall process efficiency.

This chapter presents a comprehensive heat transfer model for an industrial furnace system (model N 120/85 HA) used to process heat-treatable aluminum alloys. The model captures transient thermal behavior through conduction, convection, and radiation between the furnace air, the multi-layered walls, the internal metal sheets, the radiative heating elements, and the alloy component placed within the furnace chamber. A key feature of the furnace is its heating design: electrical heating elements are housed behind thin metal sheets, which absorb radiative energy and transfer it to the furnace interior via forced convection.

Due to the multi-physics nature of the system and the strong coupling between heat transfer modes, obtaining an analytical solution is intractable. The governing equations are therefore solved numerically, enabling the simulation of temperature evolution for critical components—the alloy, the furnace air, and the walls—over time. Given the complexity of the system, a number of simplifying assumptions are made. These assumptions, along with their motivations, are listed in Section 3.2.

3.1 Physical Model

The furnace is modeled as a closed thermal system comprising five primary domains: a multi-layered insulated wall structure, an enclosed air volume, heating elements in separated chambers, thin metal sheets which separate the main chamber from the heating compartments, and an alloy component positioned within the furnace interior. Heat exchange occurs between all the domains.

A schematic of the furnace system is shown in Figure 3.2. The key physical components are described as follows. Note that $\mathbf{T_w}$ is expressed in boldface to represent the vector of temperatures across all six wall layers, whereas T_f denote the scalar temperature for the lumped furnace air.

- **Wall Structure** ($\mathbf{T_w}$)**:** The furnace walls consist of multiple layers, each with distinct thermal conductivities and thickness, providing insulation and structural support.
- **Furnace Air** (T_f)**:** The internal air volume occupies the space between the walls and the alloy component. A high-efficiency fan located at the back of the chamber circulates the air, promoting forced convective heat transfer and improving temperature uniformity.

- **Heating Elements** (T_h)**:** Cylindrical electrical heating elements are embedded alongside the top, bottom, and side walls of the furnace. They radiate heat to warm up the adjacent surfaces as well as the surrounding air.
- **Alloy Component** ($T_{\mathbf{alloy}}$)**:** The alloy workpiece is centrally located within the furnace chamber and exchanges heat with the surrounding air via forced convection. The heating rate and uniformity depend on the workpiece geometry and surface properties.
- **Thin Metal Sheet** ($T_{\mathbf{ms}}$)**:** This sheet lines the interior-facing boundary of the heating compartments. It absorbs radiative heat from the heating elements and transfers energy to the furnace interior primarily through convection, acting as an intermediary between the heaters and the main chamber.

3.1.1 Furnace Schematics

Figure 3.1 presents the furnace model N 120/85 HA as provided in the manufacturer's manual.

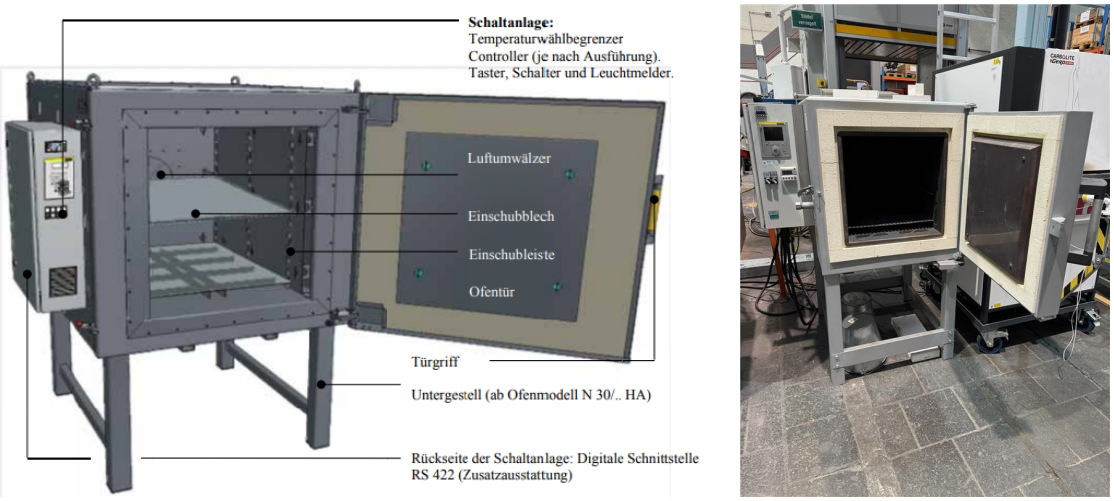

(a) Schematic representation from the manual. (b) Photograph of the actual furnace.

Figure 3.1: Furnace model N 120/85 HA.

Figure 3.2 shows a visual representation of the primary heat transfer mechanisms within the furnace, constructed for clarity in the context of this thesis.

For reference, the symbols used in Figure 3.2 denote the following heat transfer terms:

- $\dot{Q}_{h\to\mathrm{ms}}$: Radiative heat transfer from the heating elements to the thin metal sheet.
- $\dot{Q}_{h\to w}$: Radiative heat transfer from the heating elements to the walls.
- $\dot{Q}_{\mathrm{ms}\to f}$: Convective heat transfer from the metal sheets to the furnace air.
- $\dot{Q}_{f\to\mathrm{alloy}}$: Convective heat transfer from the furnace air to the alloy element.
- $\dot{Q}_{f\to w}$: Convective heat transfer from the furnace air to the walls.

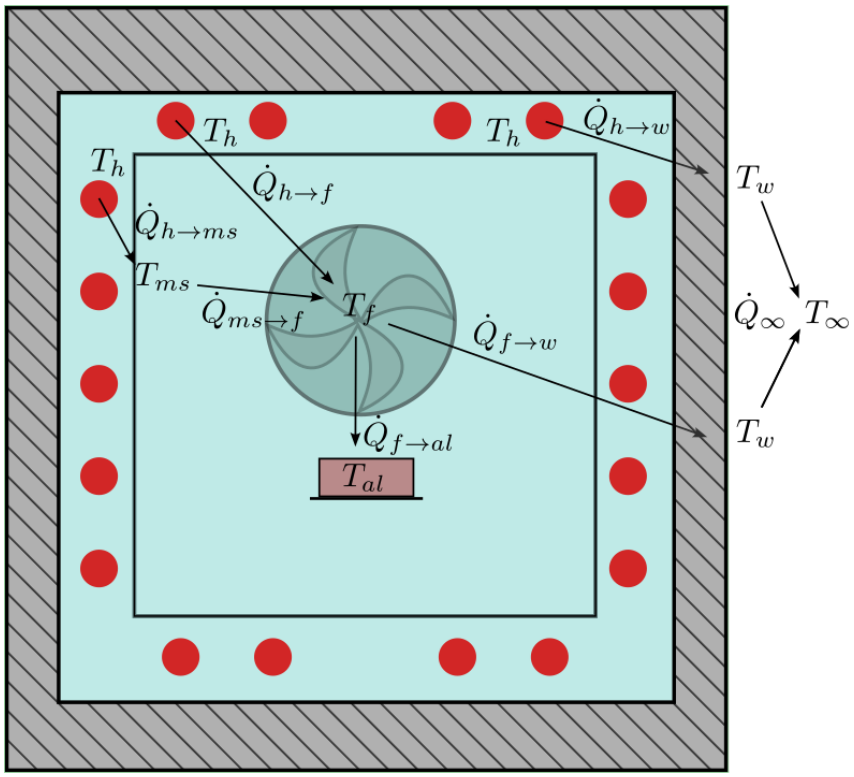

Figure 3.2: Schematic illustration of heat transfer pathways in the furnace.

- $\dot{Q}_{\infty}$: Convective heat transfer from the walls to the external environment.

Finally, Figure 3.3 provides a schematic overview of the multi-layered furnace walls, an example temperature distribution across the layers.

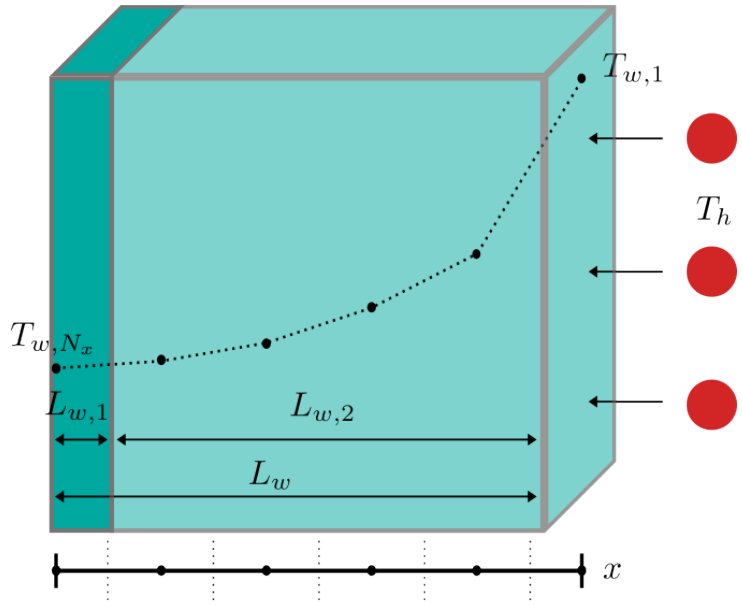

Figure 3.3: Schematic representation of furnace walls.

For notational clarity (as referenced in Figure 3.3),

- L_w denotes the total wall thickness, measured from the exterior to the interior surface.
- $L_{w,1}$ and $L_{w,2}$ refer to the thicknesses of the first and second wall layers, respectively.

3.2 Assumptions

Modeling the transient temperature behavior of an industrial furnace involves inherent complexities due to the unpredictable nature of real-world operation. For instance, furnace doors may be opened intermittently during processing to insert or remove materials. The associated heat loss depends on uncertain variables—such as door aperture, duration, and operator behavior—which are difficult to predict or control mechanically. This is one of many practical scenarios that introduces variability not easily incorporated into a deterministic thermal model.

The goal of this work is not to account for all such operational variances, but rather to develop a tractable model that can capture the dominant thermal phenomena with sufficient accuracy. In particular, we aim to predict the temperature evolution of all furnace domains over time, enabling calibration and comparison to empirical data. With this purpose in mind, the following simplifying assumptions are made:

A.1 Homogeneous wall properties. All layers of the furnace walls have constant, temperature-independent thermal properties (e.g. conductivity, density, heat capacity).

A.2 Uniform wall geometry. Each wall has constant thickness and cross-sectional area; lateral (in-plane) conduction is neglected (one-dimensional model).

A.3 No internal volumetric heat generation. There is no source of heat generation within the solid domains of the furnace walls. Any heat entering or leaving the (wall)-system is specified through boundary conditions.

A.4 Perfectly mixed furnace air. Forced convection by an internal fan ensures the air temperature is spatially uniform at all times.

A.5 Closed mass system. No mass enters or leaves the furnace; air leakage, infiltration, or exhaust is neglected.

A.6 Alloy temperature uniformity. The alloy component has a single, uniform temperature, justified by its relatively small size and high thermal conductivity.

A.7 No internal alloy heat generation. The alloy does not generate heat (e.g. via latent heat of phase change), and only absorbs energy from the furnace environment.

A.8 Metal sheet temperature uniformity. Given the small thickness of the metal sheet (approximately 3 mm), it is reasonable to assume negligible thermal gradients within the sheet itself, treating its temperature as spatially uniform.

A.9 Constant ambient temperature. Ambient (external) temperature is fixed at room temperature throughout each simulation.

A.10 Known initial conditions. Initial temperatures of furnace air, wall nodes, and alloy are prescribed exactly at $t = 0$.

A.11 Radiation path. Heating elements radiate only to adjacent wall layers and a thin metal sheet; direct radiation to air or alloy is neglected, matching the furnace's separated-channel design.

A.12 Negligible cavity reflections. Radiative reflections within the cavity are ignored—walls are treated as perfectly absorbing, even if they are not "black-bodies".

A.13 Uniform sheet emission. The thin internal metal sheet emits radiation uniformly toward the furnace air; spatial variations in emissive flux are neglected.

A.14 Non-participating furnace atmosphere. Furnace gases do not participate in radiation; any absorption or emission by H_2O (water vapor) or CO_2 is lumped into effective emissivities.

These assumptions are selected to balance physical realism with model simplicity and computational feasibility. While some, such as constant ambient temperature or uniform wall geometry, are well justified, others may introduce error. For most assumptions, such as neglecting the lateral conductive heat-flow through the furnace walls, the effects are expected to be minimal. With these assumptions we aim to strike a balance between fidelity and computational efficiency. Their impact on model accuracy is quantified in Chapter 6 by comparing against experimental temperature profile.

3.3 Heat Transfer Mechanisms

Heat transfer is the movement of thermal energy driven by spatial temperature gradients [22]. In the furnace model, three primary modes of heat transfer are considered: conduction, convection, and radiation, as defined below.

Conduction

Conduction is the transfer of thermal energy through a solid or between solids in contact, driven by molecular interactions arising from temperature gradients. In the furnace, conduction is the primary mode of heat transfer through the multilayered walls, where heat propagates from the inner surface (exposed to furnace air and radiative heating elements) to the outer surface (exposed to ambient conditions).

Given the transient nature of the heating process, conduction is modeled as a time-dependent phenomenon. The furnace walls consist of multiple layers, each characterized by distinct thermophysical properties—density (ρ_w), specific heat capacity ($C_{p,w}$), and thermal conductivity (λ_w)—which govern the rate of heat diffusion. At material interfaces, thermal conductivity may change abruptly, requiring careful numerical treatment to capture the heat flow accurately.

The conductive heat flux is described by Fourier's one-dimensional law of heat conduction [22]:

$$q_{\text{cond}} = -\lambda_w \frac{\partial T_w}{\partial x} = -\lambda_w \frac{\Delta T_w}{\Delta x}, \tag{3.1}$$

where q_{cond} is the conductive heat flux (W/m^2) and λ_w is the thermal conductivity of the wall material. The negative sign reflects the natural flow of heat from higher to lower temperatures, analogous to diffusion. Figure 3.4 illustrates conduction across the furnace wall layers over time.

Convection

Convection is the mode of heat transfer that occurs due to the bulk movement of a fluid (gas or liquid) in conjunction with heat diffusion within the fluid. In the furnace system, convection is essential for heat exchange between the furnace air, the alloy, and the wall surfaces.

There are two types of convection:

- **Natural (Free) Convection:** Driven by buoyancy forces arising from temperature-induced density variations in the fluid. No mechanical assistance (e.g., a fan or pump) is involved.
- **Forced Convection:** The fluid motion is driven by an external source such as a fan, blower, or pump, leading to enhanced heat transfer rates.

Generally, convective heat transfer is modeled using Newton's Law of Cooling [22]:

$$q_{\text{conv}} = h\Delta T, \tag{3.2}$$

where:

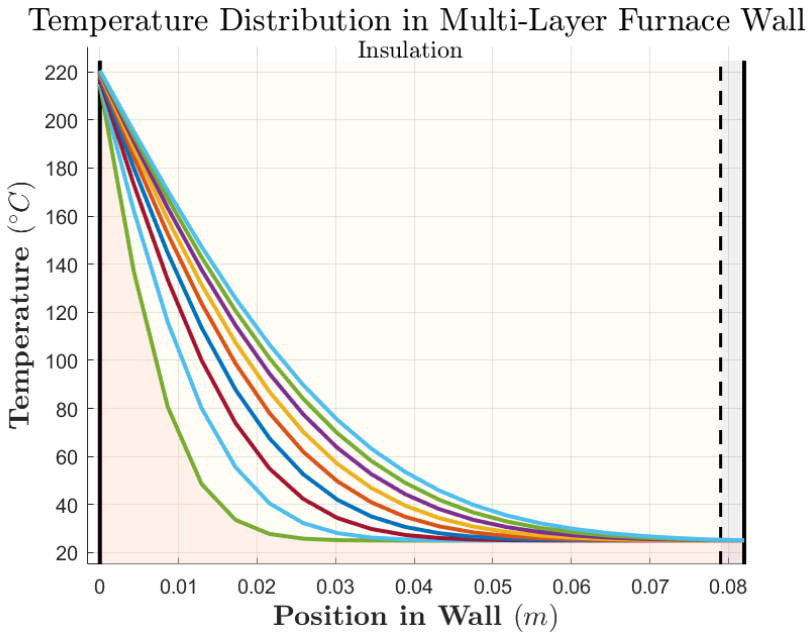

Figure 3.4: Conductive heat transfer within the wall over time.

- q_{conv} is the convective heat flux (W/m^2),
- h is the convective heat transfer coefficient ($\text{W/m}^2\cdot\text{K}$),
- ΔT is the temperature difference between the solid surface and the surrounding fluid.

A positive q_{conv} indicates heat transfer to the surface; a negative value indicates heat leaving it. The coefficient h depends on fluid properties, flow conditions, and surface geometry, and is estimated from empirical correlations, computational-fluid-dynamics (CFD) simulations, or experimental calibration.

In this model, convection occurs in the following regions:

- Between the thin metal sheets and the internal furnace air (primarily forced convection due to fan operation)
- Between the furnace walls and the internal furnace air (natural and/or forced convection depending on fan operation).
- Between the alloy and the furnace air (primarily forced convection due to air circulation).
- Between the outer furnace wall surface and the ambient air (natural convection).

Radiation

Radiation is the emission and absorption of thermal energy in the form of electromagnetic waves, capable of traveling through a vacuum. All objects with a temperature above absolute zero emit

thermal radiation in the form of electromagnetic waves, primarily in the infrared spectrum. The rate of radiative heat transfer depends on the temperature and emissivity of the surfaces involved, as well as the relative geometry between the emitting and receiving bodies.

The fundamental equation for radiative heat transfer is the Stefan-Boltzmann Law [22], which relates the power radiated by a surface to its temperature:

$$Q_{\text{rad}} = \epsilon \, \sigma \, A \, (T^4 - T_\infty^4), \tag{3.3}$$

where:

- Q_{rad} is the radiative heat transfer rate (W),
- ϵ is the emissivity of the surface,
- σ is the Stefan-Boltzmann constant (5.67×10^{-8} W/m^{2}K^4),
- A is the surface area of the emitting body (m^2),
- T is the temperature of the emitting body (K),
- T_∞ is the temperature of the surroundings or ambient medium (K).

For radiative exchange between two finite surfaces, the net heat transfer is given by:

$$\dot{Q}_{1\to2} = \frac{\sigma \epsilon_1 \epsilon_2 A_2 (T_1^4 - T_2^4)}{\left(\frac{1}{\epsilon_1} + \frac{1}{\epsilon_2} - 1\right)}, \tag{3.4}$$

where:

- $\dot{Q}_{1\to2}$ is the net heat transfer rate from surface 1 to surface 2 (W),
- ϵ_1 and ϵ_2 are the emissivities of surfaces 1 and 2, respectively,
- A_1 is the surface area of the emitting body (m^2),
- T_1 and T_2 are the temperatures of surfaces 1 and 2, respectively (K).

This expression assumes two parallel, diffuse gray surfaces. For non-parallel or complex geometries, a *view factor* must be introduced to account for mutual orientation and visibility [4].

View Factors and Radiative Exchange in Enclosures

The view factor (also called the configuration or shape factor) $F_{i\to j}$ quantifies the fraction of radiation leaving surface i that reaches surface j directly [4]. It depends solely on the geometry and orientation of the surfaces involved. These view factors have two key properties:

1. Reciprocity Relation:

$$A_i \cdot F_{i\to j} = A_j \cdot F_{j\to i}. \tag{3.5}$$

2. Summation Rule in an Enclosure:

$$\sum_{j=1}^{N} F_{i\to j} = 1. \tag{3.6}$$

For this model, the specific analytical expression for the view factors are discussed in Section 4.1.3.

3.4 PID Controller

To regulate the furnace temperature T_f toward a desired set-point T_f^{set}, a Proportional-Integral-Derivative (PID) controller is employed. The controller continuously adjusts the power input P_{input} to the heating elements based on the temperature error and its dynamics. PID control is widely used in industrial systems due to its simplicity and effectiveness [1, 41].

Control Law

The PID control signal is computed as a weighted sum of three components:

$$P_{\text{unsat}} = K_p\, e(t) + K_i \int_0^t e(\tau)\, d\tau + K_d\, \frac{de}{dt}, \tag{3.7}$$

where:

- $e(t) = T_f^{\text{set}} - T_f(t)$ is the instantaneous temperature error,
- K_p, K_i, and K_d are the proportional, integral, and derivative gains,
- The integral term eliminates steady-state offset by accounting for accumulated past errors,
- The derivative term anticipates future error trends, improving responsiveness and reducing overshoot.

Power Saturation and Anti-Windup

In practical implementations, the power input is constrained through clamping:

$$P_{\text{input}} = \min\left(\max(0, P_{\text{unsat}}), P_{\max}\right) \quad \Longrightarrow \quad 0 \le P_{\text{input}} \le P_{\max}. \tag{3.8}$$

When the unsaturated controller output exceeds these limits, the integral term may continue to accumulate error, causing excessive overshoot and degraded transient performance. To prevent this, an anti-windup mechanism is implemented: the integral term is updated only when the output is not saturated, or when the error is driving the power back into the valid range. In this model, clamping (conditional integration) is used to stop integration of the error when the output is saturated and the error would drive further saturation.

Practical Considerations

Tuning the PID gains K_p, K_i, and K_d is essential to balance heating rates, overshoot, and stability. The controller must handle actuator saturation gracefully, especially during large transients such as heat-up or cool-down phases. With these considerations, the PID controller provides robust and responsive regulation of furnace temperature, while remaining simple and computationally efficient for time-stepping simulations.

4 Mathematical Formulation

To model the thermal behavior of the furnace system, this paper constructs a set of governing equations based on the principle of energy conservation. These equations describe the time-dependent temperature evolution of the major components: the furnace air, the alloy, the furnace walls, the thin metal sheets, and the heating elements. By coupling these components through appropriate boundary conditions and heat transfer mechanisms—conduction, convection, and radiation—the dynamic thermal response of the furnace during operation can be simulated.

The mathematical formulation specifies the energy balance equations for each component. Depending on the spatial and temporal characteristics, these balances are expressed as either partial differential equations (PDEs) or ordinary differential equations (ODEs). The resulting systems are solved numerically using discretization schemes, with carefully defined boundary and initial conditions to ensure both physical realism and numerical consistency, as discussed in Chapter 5.

In this chapter, the following are introduced:

- The governing equations for modes of heat transfer.
- The energy balance equations for the furnace components.
- The boundary and initial conditions defining the overall system.

Together, these formulations enable prediction of temperature evolution throughout the furnace and provide the foundation for optimizing performance with respect to energy efficiency and the attainment of desired mechanical properties.

4.1 Governing Equations

The three modes of heat transfer—conduction, convection, and radiation—were introduced in Section 3.3. In this section, the relevant equations are adapted to describe the governing and energy balance equations.

4.1.1 Conduction

Conduction governs heat transfer within the multilayered furnace walls. In this model, one-dimensional transient heat conduction is assumed in the spatial direction normal to the furnace wall surfaces. This assumption is justified by the slab-like geometry and insulation design of the furnace. The general form of the one-dimensional transient heat conduction equation (diffusion equation) in a solid medium is:

$$\rho\, C_p \frac{\partial T(x,t)}{\partial t} = \frac{\partial}{\partial x}\left(\lambda \frac{\partial T(x,t)}{\partial x} + \dot{q}(x,t)\right), \tag{4.1}$$

a parabolic (time-dependent) partial differential equation, where:

- $T(x,t)$ is the temperature $[K]$ at position x and time t,
- ρ is the material density $[kg/m^3]$,
- C_p is the specific heat capacity $[J/(kg \cdot K)]$,
- λ is the thermal conductivity coefficient $[W/(m \cdot K)]$,
- $\dot{q}(x,t)$ is the heat generating source term $[W/m^3]$.

The furnace consists of six walls—top, bottom, front, back, and two sides—each with slight differences in geometry and composition. Since wall composition affects conduction, the diffusion equation must be applied to each wall individually. Only the left and right walls are identical in thickness, allowing them to be modeled together. The temperature distribution within each wall is determined by solving the heat diffusion equation. For a generic wall, the governing equation reduces to:

$$\rho_w(x)\, C_{p,w}(x)\, \frac{\partial T_w(x,t)}{\partial t} = \frac{\partial}{\partial x}\left(\lambda_w(x)\, \frac{\partial T_w(x,t)}{\partial x}\right), \tag{4.2}$$

Here:

- $T_w(x,t)$ is the temperature $[K]$ of a given wall at position x and time t,
- $\rho_w(x)$ is the density of the wall material as a function of position $[kg/m^3]$,
- $C_{p,w}(x)$ is the specific heat capacity as a function of position $[J/(kg \cdot K)]$,
- $\lambda_w(x)$ is the thermal conductivity coefficient as a function of position $[W/(m \cdot K)]$.

Under Assumption **A.3** (see Section 3.2), no internal volumetric heat generation occurs within the solid wall domains. Thus, the source term is set to zero (i.e., $\dot{q}(x,t) = 0$). Instead, energy input from radiative heating elements and convective exchange with the furnace air is incorporated through boundary conditions at the inner wall surface, while heat loss to the ambient is applied at the outer wall surface. This approach is justified by the effective thermal insulation of the walls, which localizes heat exchange at their boundaries.

Wall Material Properties

The furnace walls are composed of multiple layers, each characterized by distinct thermophysical properties: density (ρ_w), specific heat capacity ($C_{w,p}$), and thermal conductivity (λ_w). These material-dependent properties govern the rate and distribution of heat conduction.

Thermal Conductivity Across Layers

Thermal conductivity λ_w is discontinuous at material interfaces. To preserve energy conservation and numerical stability within the finite volume discretization, the effective thermal conductivity at an interface is computed using the harmonic mean. Following the recommendation of

Moukalled, Mangani, and Darwish [32], this yields accurate interfacial values when properties are discontinuous:

$$\lambda_{\text{eff}} = \frac{2\lambda_1\lambda_2}{\lambda_1 + \lambda_2}, \tag{4.3}$$

where λ_1 and λ_2 are the thermal conductivities of adjacent layers.

Density and Specific Heat Capacity

The density (ρ_w) and specific heat capacity ($C_{p,w}$) are assumed uniform within each layer but vary between layers. They are defined piecewise as:

$$\rho_w(x) = \begin{cases} \rho_{w,1} & 0 \le x \le L_{w,1}, \\ \rho_{w,2} & L_{w,1} < x \le L_w, \end{cases} \tag{4.4}$$

$$C_{p,w}(x) = \begin{cases} C_{p,w,1} & 0 \le x \le L_{w,1}, \\ C_{p,w,2} & L_{w,1} < x \le L_w, \end{cases} \tag{4.5}$$

where $L_{w,1}$ denotes the thickness of the first (insulated) material layer, and L_w the total thickness of the furnace wall.

4.1.2 Convection

Convective heat transfer governs the exchange of thermal energy between the well-mixed furnace air and surrounding components, including the alloy, the thin metal sheet, the inner wall surfaces, and the heating elements. Owing to the fan located at the back wall, the air is assumed to be spatially uniform in temperature $T_f(t)$ at each time step, while convective heat fluxes at different interfaces are modeled by Newton's law of cooling [59]. The following subsections specify this relation for each component.

Air-Alloy Convection

The alloy component is heated primarily by convection from the circulating air:

$$\dot{Q}_{f\to\text{alloy}}(t) = h_{f\to\text{alloy}}\, A_{\text{alloy}}\, (T_f - T_{\text{alloy}}), \tag{4.6}$$

where A_{alloy} is the effective exposed surface area of the alloy element and $h_{f\to\text{alloy}}$ is the corresponding convective heat transfer coefficient. T_f denotes the temperature of the furnace air, and T_{alloy} the (external) temperature of the alloy element.

Air-Metal Sheet Convection

The thin metal sheet lining the heating chambers absorbs radiation from the heating elements and transfers this energy to the furnace air by convection:

$$\dot{Q}_{\text{ms}\to f} = h_{\text{ms}\to f}\, A_{\text{ms}}\, (T_{\text{ms}} - T_f), \tag{4.7}$$

where A_{ms} is the effective surface area of the sheet and T_{ms} its temperature.

Air-Wall Convection

The furnace walls also exchange heat convectively with the furnace air, typically acting as a sink once the air temperature exceeds that of the surrounding structure [3]:

$$\dot{Q}_{f\to w} = \sum_i \dot{Q}_{f\to w_i} = h_{f\to w_i}\, A_{w_i}\, (T_f - T_{w_i,1}), \tag{4.8}$$

where $T_{w_i,1}$ is the innermost node temperature of wall i and A_{w_i} its surface area.

Air-Heater Convection

In addition to radiative transfer, the heating elements contribute convective heat to the furnace air:

$$\dot{Q}_{h\to f} = h_{h\to f}\, A_h\, (T_h - T_f), \tag{4.9}$$

where T_h is the heater temperature, A_h its effective surface area, and $h_{h\to f}$ the corresponding convective coefficient.

Fan-Induced Enhancement of Convective Heat Transfer

In the present model, convective coefficients h are treated as constants. However, they can be refined to capture the effect of fan-induced air velocity by employing Nusselt number correlations. The dimensionless Nusselt number is defined as:

$$Nu = \frac{hL}{\lambda}, \tag{4.10}$$

where L is a characteristic length and λ the thermal conductivity of the furnace air. It depends on the Reynolds and Prandtl numbers:

$$Re = \frac{\rho u L}{\mu}, \qquad Pr = \frac{c_p \mu}{k}, \tag{4.11}$$

with u the fan-induced velocity, ρ the air density, μ its dynamic viscosity, and c_p its specific heat capacity. Standard correlations for forced convection over a flat plate are [22]:

$$Nu = \begin{cases} 0.332\, Re^{1/2}\, Pr^{1/3}, & Re < 5\times 10^5 \quad \text{(laminar)}, \\ 0.037\, Re^{4/5}\, Pr^{1/3}, & Re \geq 5\times 10^5 \quad \text{(turbulent)}. \end{cases} \tag{4.12}$$

From these, the convective coefficient follows as $h = Nu\lambda/L$. Incorporating this dependency would enable the model to account explicitly for fan speed as a controllable parameter influencing heat transfer rates.

4.1.3 Radiation

Radiative heat transfer plays a central role in the furnace model, particularly in transferring thermal energy from the heating elements to the internal surfaces: the furnace walls and the thin metal sheets separating the heating chambers from the main interior. The heating elements emit thermal radiation, which is partially absorbed by these surfaces and subsequently re-emitted or conducted into surrounding domains via convection.

Heater-Wall Radiation

The net radiative flux from the heating elements to the interior wall is modeled using the Stefan-Boltzmann law (Equation (3.3)) for diffuse-gray surfaces:

$$\dot{Q}_{h\to w} = \sum_i \sigma\, \epsilon_{w_i}\, F_{h\to w_i}\, A_{w_i}\, (T_h^4 - T_{w_i,1}^4), \tag{4.13}$$

where:

- σ is the Stefan-Boltzmann constant,
- ϵ_{w_i} is the emissivity of interior wall i,
- $F_{h\to w_i}$ is the view factor from the heating elements to interior wall i,
- A_{w_i} is the surface area of interior wall i,
- T_h is the temperature of the heating elements,
- $T_{w_i,1}$ is the temperature of the interior-surface wall node.

Heater-Metal Sheet Radiation

Similarly, the radiative flux between the heating elements and the thin metal sheets is expressed as:

$$\dot{Q}_{h\to \mathrm{ms}} = \sigma\, \epsilon_{\mathrm{ms}}\, F_{h\to\, \mathrm{ms}}\, A_{ms}\, (T_h^4 - T_{\mathrm{ms}}^4), \tag{4.14}$$

where:

- ϵ_{ms} is the emissivity of the thin metal sheets,
- $F_{h\to \mathrm{ms}}$ is the view factor from the heating elements to the thin metal sheets,
- A_{ms} is the surface area of the thin metal sheet exposed to the heating elements,
- T_{ms} is the temperature of the thin metal sheets.

Under the assumption that both the interior walls and the thin metal sheet behave as infinite planes parallel to the heating elements, it is reasonable to set $A_{\mathrm{ms}} = A_{w_i} = A_w$ and $F_{h\to \mathrm{ms}} = F_{h\to w_i} = F_{h\to w}$, simplifying the calculations.

View Factors

The view factor quantifies the fraction of radiation leaving one surface that reaches another. For the furnace, the cylindrical heating elements are equidistantly positioned near the walls and thin metal sheets. Under the infinite-plane assumption, the analytical view factor from a cylinder to a plane is given by [4]:

$$F_{i\to j} = 1 - \left[1 - \left(\frac{D}{S}\right)^2 + \left(\frac{D}{S}\right) \tan^{-1}\left[\left(\frac{S^2 - D^2}{D^2}\right)^{\frac{1}{2}}\right]\right], \tag{4.15}$$

where:

- D is the diameter of the cylindrical heating elements,
- S is center-to-center spacing between adjacent heating elements,
- i denotes the plane surface, and j denotes the cylindrical element.

In practice, the values of D and S were estimated based on the furnace schematics of the N 120/85 HA model. This approximation maintains the equidistant assumption of the analytical formula while remaining consistent with the physical layout of the furnace. Using these estimated values, the view factors $F_{h\to w}$ and $F_{h\to \mathrm{ms}}$ could be calculated analytically, providing a reasonable approximation of radiative exchange without requiring complex numerical ray-tracing or finite-element radiation modeling.

4.2 Energy Balance Equations

This section presents the final energy balance equations for each component of the furnace system, capturing the dynamic exchange of thermal energy via conduction, convection, and radiation.

4.2.1 Furnace Wall Energy Balance

The energy balance for the furnace walls follows the one-dimensional transient heat conduction equation outlined in Section 4.1.1:

$$\rho_w C_{p,w} \frac{\partial T_w}{\partial t} = \frac{\partial}{\partial x}\left(\lambda_w \frac{\partial T_w}{\partial x}\right), \tag{4.16}$$

where all quantities are as defined previously.

4.2.2 Furnace Air Energy Balance

The temperature of the furnace air evolves due to convective exchange with the walls, the alloy, and the thin metal sheet, as well as heat received from the heating elements. The governing energy balance is:

$$m_f C_{p,f} \frac{\mathrm{d}T_f}{\mathrm{d}t} = \dot{Q}_{h\to f} + 2\cdot \dot{Q}_{\mathrm{ms}\to f} - \sum_{i=1}^{6} \dot{Q}_{f\to w_i} - \dot{Q}_{f\to \mathrm{alloy}}, \tag{4.17}$$

where:

- $m_f = \rho_f \cdot V_f$ is the mass of the furnace air,
- $C_{p,f}$ is the specific heat capacity of the furnace air,
- $\dot{Q}_{h\to f}$ is the convective heat flux from heating elements to the air,
- $\dot{Q}_{\mathrm{ms}\to f}$ is convective heat flux from the thin metal sheets (doubled for both surfaces),
- $\dot{Q}_{f\to w_i}$ is the convective heat flux to wall i,
- $\dot{Q}_{f\to \mathrm{alloy}}$ is the convective heat flux to the alloy.

4.2.3 Alloy Energy Balance

The alloy is heated primarily by convection from the surrounding furnace air:

$$m_{\text{alloy}} C_{p,\text{alloy}} \frac{\mathrm{d}T_{\text{alloy}}}{\mathrm{d}t} = \dot{Q}_{f \to \text{alloy}}, \tag{4.18}$$

where m_{alloy} and $C_{p,\text{alloy}}$ denote the mass and specific heat capacity of the alloy, respectively.

Temperature-Dependent Material Properties

Some thermophysical properties vary with temperature, influencing the accuracy of high-temperature simulations. In this model, we account for the temperature dependence of the furnace air density and the alloy's specific heat capacity. Other properties, such as wall thermal conductivity or air heat capacity, are assumed constant due to their relatively minor influence on energy transport and to avoid unnecessary computational complexity.

Furnace Air Density

The density of furnace air is computed using the ideal gas law [30]:

$$\rho_f(T_f) = \frac{PM}{RT_f}, \tag{4.19}$$

where:

- P is the pressure inside the furnace $[Pa]$,
- M is the molar mass of the air mixture $[kg/mol]$,
- R is the universal gas constant $[8.314\ J/(kg \cdot K)]$,
- T_f is the air temperature $[K]$.

This expression is derived from the ideal gas law in its molar form:

$$PV = nRT.$$

Dividing both sides by volume V and using the definition of density $\rho = \frac{m}{V}$ and $n = \frac{m}{M}$, where m denotes mass and n denotes the number of moles, we obtain:

$$P = \frac{nRT}{V} = \frac{mRT}{MV} \Rightarrow \rho = \frac{PM}{RT}.$$

This inverse relationship with temperature reflects the physical behavior of gas mixtures expanding as they are heated. While the ideal gas law is itself an approximation, many gases are well-approximated by ideal gas behavior.

Alloy Specific Heat Capacity

The specific heat capacity of the alloy is temperature dependent:

$$C_{p,\text{alloy}}(T_{\text{alloy}}) = C_{p,\text{mech}}(T_{\text{alloy}}) + C_{p,\text{temp}}(T_{\text{alloy}}), \tag{4.20}$$

where $C_{p,\text{temp}}$ represents thermal contributions due to lattice vibrations. For practical simulations, $C_{p,\text{temp}}$ is often modeled using experimentally validated polynomial fits [17, 55]:

$$C_{p,\text{alloy}}(T_{\text{alloy}}) = a + bT_{\text{alloy}} + cT_{\text{alloy}}^2, \tag{4.21}$$

where the coefficients a, b, and c are material-specific and obtained from thermophysical databases (e.g., NIST [34]). For simplicity, potential mechanical contributions to C_p under stress due to deformation or other microstructural changes are neglected.

4.2.4 Thin Metal Sheet Energy Balance

The thin metal sheets absorb radiation from the heating elements and exchange heat convectively with the furnace air:

$$m_{\text{ms}} C_{p,\text{ms}} \frac{\mathrm{d}T_{\text{ms}}}{\mathrm{d}t} = \dot{Q}_{h\to\text{ms}} - 2 \cdot \dot{Q}_{\text{ms}\to f}, \tag{4.22}$$

where m_{ms} and $C_{p,\text{ms}}$ denote the mass and specific heat capacity of the sheet, respectively.

4.2.5 Heating Elements Energy Balance

The heating elements (treated as a lumped mass) are powered by an external source $\dot{Q}_{\text{power}}$ and release energy through radiation and convection:

$$m_h C_{p,h} \frac{\mathrm{d}T_h}{\mathrm{d}t} = \dot{Q}_{\text{power}} - \dot{Q}_{h\to\text{ms}} - \sum_{i=1}^{6} \dot{Q}_{h\to w_i} - \dot{Q}_{h\to f}. \tag{4.23}$$

Here, m_h and $C_{p,h}$ denote the mass and specific heat capacity of the heating elements, and $\dot{Q}_{\text{power}}$ represents the heat generated by the electrical power input.

4.3 Boundary and Initial Conditions

To solve the transient heat transfer problem in the furnace system, appropriate boundary and initial conditions are specified for each modeled subsystem: the furnace air, the alloy, the furnace walls, the thin metal sheet, and the heating elements. These conditions ensure a well-posed problem and reflect the physical assumptions and operational constraints of the furnace.

Initial Conditions

At the start of the heating cycle, the entire system is assumed to be isothermal at an ambient initial temperature T_0. The initial conditions are:

- Furnace air: $T_f(t=0) = T_f^0$,

- Alloy : $T_{\text{alloy}}(t=0) = T^0_{\text{alloy}}$,
- Walls: $T_{w_i}(x, t=0) = T^0_{w_i}(x)$(distribution across layers),
- Thin metal sheet: $T_{\text{ms}}(t=0) = T^0_{\text{ms}}$,
- Heating elements: $T_h(t=0) = T^0_h$.

Boundary Conditions for Furnace Walls

The furnace walls are modeled as one-dimensional multilayer slabs. They are heated from the interior furnace chamber and exchange heat with the ambient environment through convection. The boundary conditions are defined as follows:

- **Interior surface (adjacent to heating chamber)**:

$$-\left(\lambda_{w_i}\frac{\partial T_{w_i}}{\partial x}\right)\Bigg|_{x=0} = h_{f\to w_i}\left(T_f - T_{w_i,1}\right) + \sigma\,\epsilon_{w_i}\,F_{h\to w_i}\left(T_h^4 - T_{w_i,1}^4\right) \tag{4.24}$$

$$= \dot{q}_{f\to w_i} + \dot{q}_{h\to w_i}. \tag{4.25}$$

  Here $T_{w_i,1}$ is the temperature of the inner node of wall i, $\dot{q}_{f\to w_i}$ is the convective heat flux from the furnace air, and $\dot{q}_{h\to w_i}$ is the radiative heat flux from the heating elements to wall i.

- **Exterior surface (exposed to ambient)**:

$$-\left(\lambda_{w_i}\frac{\partial T_{w_i}}{\partial x}\right)\Bigg|_{x=L} = h_{w_i\to\infty}(T_{w_i,L} - T_\infty) \tag{4.26}$$

$$= \dot{q}_{w_i\to\infty}. \tag{4.27}$$

  Here $T_{w_i,L}$ is the temperature of the outer node of wall i, T_∞ is the ambient temperature, $h_{w_i\to\infty}$ is the convective heat transfer coefficient to the surroundings, and $\dot{q}_{w_i\to\infty}$ is the corresponding heat flux. The ambient temperature is assumed constant at room conditions.

5 Numerical Methods

This chapter outlines the numerical methods used to simulate the transient thermal behavior of the industrial furnace system described in Chapters 3 and 4. Given the complexity of the physical processes involved—including conduction through multi-layered walls, convection between fluid and solid domains, and radiative heat exchange within enclosures—a combination of discretization techniques, time-integration schemes, and coupled solution strategies are employed to approximate solutions.

The primary objective of the numerical model is to accurately compute the temporal evolution of temperatures across the various furnace components. These temperature histories are subsequently used to evaluate the mechanical properties of the heated alloy, forming the basis for system optimization. In Chapter 6, the numerical solution of the model is compared to physical measurement data.

5.1 Spatial Discretization: Finite-Volume-Method

To numerically solve the heat transfer equations, the furnace wall is discretized into N_x control volumes using the finite-volume-method (FVM). The governing heat diffusion equation (Equation (4.2)) is then integrated over each control volume individually to obtain a discretized form.

5.1.1 Grid Generalization

The wall is divided into N_x cells, each containing a node located at its center. The cells are arranged such that the first and last nodes coincide with the physical boundaries of the wall. Specifically, node $i = 1$ corresponds to the lateral position $x = 0$, at the interior surface of the wall, and node $i = N_x$ corresponds to $x = L$, at the exterior surface.

With this arrangement, all internal cells have a width of Δx and a cross-sectional area of A_w. The first and last cells, however, have a width of $\frac{1}{2}\Delta x$, which is reflected in the corresponding control volumes:

$$\begin{aligned} \mathbf{V}_w &= (V_{w,1}, V_{w,2}, \ldots, V_{w,N_x}), \\ V_{w,1} &= V_{w,N_x} = \frac{1}{2}\Delta x \cdot A_w, \\ V_{w,i} &= \Delta x \cdot A_w \qquad \text{for } i = 2, \ldots, N_x - 1. \end{aligned}$$

For simplicity, a uniform and constant grid discretization is assumed here, although non-uniform grids can be used in general. Temperatures are stored at the cell centers (nodes), while fluxes are evaluated at the interfaces. Figure 5.1 provides a schematic overview of the grid layout.

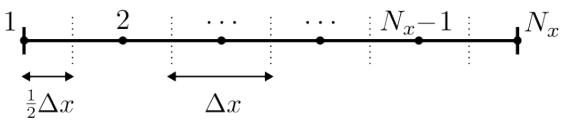

Figure 5.1: Schematic representation of the grid used for discretization.

5.1.2 Discretization

Applying the FVM, Equation (4.2) is integrated over each control volume $V_{w,i}$ (assumed uniform) to yield:

$$\int_{V_{w,i}} \rho_w C_{p,w} \frac{\partial T_w}{\partial t} \mathrm{d}V = \int_{V_{w,i}} \frac{\partial}{\partial x} \left(\lambda_w \frac{\partial T_w}{\partial x} \right) \mathrm{d}V, \tag{5.1}$$

Focusing on the right-hand side of Equation (5.1), this term represents the net heat flux into the control volume, i.e., the heat flux entering minus the heat flux leaving. The conductive heat flux at an interface is given by Fourier's law [22]:

$$\dot{Q} = -\lambda A \frac{\partial T}{\partial x} = -\lambda A \frac{\Delta T}{\Delta x}, \tag{5.2}$$

which is discretized as:

$$\dot{Q}_{i-\frac{1}{2}} = -\left(\lambda_w A_w\right)_{i-\frac{1}{2}} \frac{T_{w,i} - T_{w,i-1}}{\Delta x_{i,i-1}}, \tag{5.3a}$$

$$\dot{Q}_{i+\frac{1}{2}} = -\left(\lambda_w A_w\right)_{i+\frac{1}{2}} \frac{T_{w,i+1} - T_{w,i}}{\Delta x_{i+1,i}}. \tag{5.3b}$$

Here, $\dot{Q}_{i-\frac{1}{2}}$ represents the heat flux into the control volume, while $\dot{Q}_{i+\frac{1}{2}}$ corresponds to the flux leaving the control volume in one dimension.

Since the control volumes are assumed uniform, the cross-sectional area A_w and the spacing Δx remain constant across the wall. The thermal conductivity at the interfaces is calculated using the harmonic mean (Equation (4.3)):

$$\lambda_{w,i-\frac{1}{2}} = \frac{\lambda_{w,i}\lambda_{w,i-1}(\partial x_i + \partial x_{i-1})}{\lambda_{w,i-1}\partial x_i + \lambda_{w,i}\partial x_{i-1}} = \frac{2\lambda_{w,i}\lambda_{w,i-1}}{\lambda_{w,i} + \lambda_{w,i-1}}, \tag{5.4a}$$

$$\lambda_{w,i+\frac{1}{2}} = \frac{\lambda_{w,i}\lambda_{w,i+1}(\partial x_i + \partial x_{i+1})}{\lambda_{w,i}\partial x_{i+1} + \lambda_{w,i+1}\partial x_i} = \frac{2\lambda_{w,i}\lambda_{w,i+1}}{\lambda_{w,i} + \lambda_{w,i+1}}. \tag{5.4b}$$

Substituting this discretization into Equation (5.1) yields the following derivation for the interior nodes ($i = 2, \ldots N_x - 1$):

$$\begin{aligned}\left(\rho_w C_{p,w} V_w \frac{\partial T_w}{\partial t}\right)_i &= \dot{Q}_{i-\frac{1}{2}} - \dot{Q}_{i+\frac{1}{2}} \\ &= \left(\lambda_w A_w\right)_{i+\frac{1}{2}} \frac{T_{w,i+1} - T_{w,i}}{\Delta x_{i+1,i}} - \left(\lambda_w A_w\right)_{i-\frac{1}{2}} \frac{T_{w,i} - T_{w,i-1}}{\Delta x_{i,i-1}}.\end{aligned} \tag{5.5}$$

At the inner wall ($i = 1$), the convective heat flux from the fluid, $\dot{Q}_{f\to w}$, replaces the absent conductive flux $\dot{Q}_{i-\frac{1}{2}}$. Similarly, at the outer wall ($i = N_x$), the heat flux leaving the wall into the environment, $\dot{Q}_{w\to\infty} = A_w\,\dot{q}_{w\to\infty}$, replaces $\dot{Q}_{i+\frac{1}{2}}$. The resulting discretized equations for the boundary nodes are:

$$\left(\rho_w C_{p,w} V_w \frac{\partial T_w}{\partial t}\right)_i = \begin{cases} (\lambda_w A_w)_{i+\frac{1}{2}} \frac{T_{w,i+1}-T_{w,i}}{\Delta x_{i+1,i}} + h_{f\to w}(T_f - T_w) & \text{for } i = 1, \\ -h_{w\to\infty}(T_w - T_\infty) - (\lambda_w A_w)_{i-\frac{1}{2}} \frac{T_{w,i}-T_{w,i-1}}{\Delta x_{i,i-1}} & \text{for } i = N_x. \end{cases} \tag{5.6}$$

5.2 Implicit Euler Method

Instead of employing an explicit scheme for the time discretization of the wall temperature evolution, which is subject to stability restrictions, an implicit scheme is considered. The implicit Euler method provides unconditional numerical stability and does not impose limitations on the size of the time steps [53]. This approach is broadly applicable and allows for more efficient simulations, particularly for models that aim to capture temperature evolution over extended periods. In the following sections, the mathematical formulation of this method is outlined for each component of the model.

5.2.1 Wall Implicit Temperature Update

The heat conduction within the furnace wall is governed by the transient heat diffusion equation:

$$\rho_w C_{p,w} \frac{\partial T_w}{\partial t} = \frac{\partial}{\partial x}\left(\lambda_w \frac{\partial T_w}{\partial x}\right).$$

Integrating over the control volumes $V_{w,i} = A_w \Delta x$ and applying the divergence theorem yields:

$$\int_{V_{w,i}} \rho_w C_{p,w} \frac{\partial T_w}{\partial t} \mathrm{d}V = \int_{V_{w,i}} \frac{\partial}{\partial x}\left(\lambda_w \frac{\partial T_w}{\partial x}\right) \mathrm{d}V.$$

Using the implicit Euler method, the temporal derivative is approximated as:

$$(\rho_w C_{p,w} V_w)_i \frac{T_{w,i}^{(k+1)} - T_{w,i}^{(k)}}{\Delta t} = \int_{V_{w,i}} \frac{\partial}{\partial x}\left(\lambda_w \frac{\partial T_w^{(k+1)}}{\partial x}\right) \mathrm{d}V.$$

Here, the temperature at the unknown $(k+1)$-th time step, $T_w^{(k+1)}$, appears on the right-hand side. This is the defining feature of the implicit method, whereas in the explicit Euler scheme, the previous time-step value $T_w^{(k)}$ would be used instead.

Applying the finite-volume method (FVM) to approximate the conduction terms gives the following discrete formulation:

$$(\rho_w C_{p,w} V_w)_i \frac{T_{w,i}^{(k+1)} - T_{w,i}^{(k)}}{\Delta t} = (\lambda_w A_w)_{i+\frac{1}{2}} \frac{T_{w,i+1}^{(k+1)} - T_{w,i}^{(k+1)}}{\Delta x_{i+1,i}} \tag{5.7}$$

$$- (\lambda_w A_w)_{i-\frac{1}{2}} \frac{T_{w,i}^{(k+1)} - T_{w,i-1}^{(k+1)}}{\Delta x_{i,i-1}}. \tag{5.8}$$

Rewriting the system for interior nodes ($i = 2, \ldots, N_x - 1$) leads to:

$$\alpha_i T_{w,i}^{(k+1)} + \beta_i T_{w,i+1}^{(k+1)} + \gamma_i T_{w,i-1}^{(k+1)} = b_i, \tag{5.9}$$

with coefficients defined as:

$$\begin{aligned}
\alpha_i &= \frac{(\rho_w C_{p,w} V_w)_i}{\Delta t} + \frac{(\lambda_w A_w)_{i+\frac{1}{2}}}{\Delta x_{i+1,i}} + \frac{(\lambda_w A_w)_{i-\frac{1}{2}}}{\Delta x_{i,i-1}}, \\
\beta_i &= -\frac{(\lambda_w A_w)_{i+\frac{1}{2}}}{\Delta x_{i+1,i}}, \\
\gamma_i &= -\frac{(\lambda_w A_w)_{i-\frac{1}{2}}}{\Delta x_{i,i-1}}, \\
b_i &= \frac{(\rho_w C_{p,w} V_w)_i}{\Delta t} T_{w,i}^{(k)}.
\end{aligned}$$

This system can be compactly written in matrix form as:

$$A\, \mathbf{T_w}^{(k+1)} = \mathbf{b}, \tag{5.10}$$

where $\mathbf{T_w}^{(k+1)}$ is the N_x-dimensional vector of wall temperatures at the $(k+1)$-th time step, and $\mathbf{b}$ contains the known temperature distribution at the previous time step.

The matrix A is tridiagonal:

$$A = \begin{bmatrix}
\alpha_1 & \beta_1 & 0 & 0 & \cdots & 0 \\
\gamma_2 & \alpha_2 & \beta_2 & 0 & \cdots & 0 \\
0 & \gamma_3 & \alpha_3 & \beta_3 & \cdots & 0 \\
\vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\
0 & 0 & \cdots & \gamma_{N_x-1} & \alpha_{N_x-1} & \beta_{N_x-1} \\
0 & 0 & \cdots & 0 & \gamma_{N_x} & \alpha_{N_x}
\end{bmatrix}. \tag{5.11}$$

Before proceeding to the solution of this system, the coefficients for the boundary nodes must be defined according to the imposed convective boundary conditions.

Interior Boundary Condition

Consider the left boundary at the furnace side ($i = 1$). Heat is transferred to the wall via convection from the furnace. From Equation (4.24), the heat flux at this boundary is:

$$-\left(\lambda_w \frac{\partial T_w}{\partial x}\right)\bigg|_{x=0} = h_{f\to w}(T_f - T_w) + \sigma \epsilon_w F_{h\to w}(T_h^4 - T_w^4) = \dot{q}_{f\to w} + \dot{q}_{h\to w}.$$

For brevity, we define an effective radiation coefficient:

$$h_{\text{rad}} = \sigma\, \epsilon_w\, F_{h\to w}. \tag{5.12}$$

Applying the finite volume method at this boundary gives:

$$\begin{aligned}(\rho_w C_{p,w} V_w)_1 \frac{T_{w,1}^{(k+1)} - T_{w,1}^{(k)}}{\Delta t} &= (\lambda_w A_w)_{1+\frac{1}{2}} \frac{T_{w,2}^{(k+1)} - T_{w,1}^{(k+1)}}{\Delta x_{2,1}} \\ &\quad + h_{f\to w} A_w \left(T_f^{(k+1)} - T_{w,1}^{(k+1)}\right) \\ &\quad + h_{\text{rad}} A_w \left((T_h^{(k+1)})^4 - (T_{w,1}^{(k+1)})^4\right). \end{aligned} \tag{5.13}$$

The term

$$\dot{Q}_{h\to w}^{(k+1)} = h_{\text{rad}} A_w \left(\left(T_h^{(k+1)}\right)^4 - \left(T_{w,1}^{(k+1)}\right)^4\right), \tag{5.14}$$

introduces nonlinearity. To maintain an implicit but linear system, this term is linearized using a first-order Taylor expansion about $T_{w,1}^{(k)}$:

$$\left(T_h^{(k+1)}\right)^4 - \left(T_{w,1}^{(k+1)}\right)^4 \approx \left(T_h^{(k+1)}\right)^4 - \left(\left(T_{w,1}^{(k)}\right)^4 + 4\left(T_{w,1}^{(k)}\right)^3 \left(T_{w,1}^{(k+1)} - T_{w,1}^{(k)}\right)\right). \tag{5.15}$$

Thus, the radiative heat flux can be approximated as:

$$\dot{Q}_{h\to w}^{(k+1)} \approx h_{\text{rad}} A_w \left(\left(T_h^{(k+1)}\right)^4 - \left(T_{w,1}^{(k)}\right)^4\right) - 4\, h_{\text{rad}} A_w \left(T_{w,1}^{(k)}\right)^3 \left(T_{w,1}^{(k+1)} - T_{w,1}^{(k)}\right). \tag{5.16}$$

Substituting this into Equation (5.13) yields the linearized boundary condition:

$$\begin{aligned}(\rho_w C_{p,w} V_w)_1 \frac{T_{w,1}^{(k+1)} - T_{w,1}^{(k)}}{\Delta t} &= (\lambda_w A_w)_{1+\frac{1}{2}} \frac{T_{w,2}^{(k+1)} - T_{w,1}^{(k+1)}}{\Delta x_{2,1}} \\ &\quad + h_{f\to w} A_w \left(T_f^{(k+1)} - T_{w,1}^{(k+1)}\right) \\ &\quad + h_{\text{rad}} A_w \left(\left(T_h^{(k+1)}\right)^4 - \left(T_{w,1}^{(k)}\right)^4\right) \qquad (5.17) \\ &\quad - 4\, h_{\text{rad}} A_w \left(T_{w,1}^{(k)}\right)^3 \left(T_{w,1}^{(k+1)} - T_{w,1}^{(k)}\right). \qquad (5.18) \end{aligned}$$

Rewriting in standard linear form:

$$\alpha_1 T_{w,1}^{(k+1)} + \beta_1 T_{w,2}^{(k+1)} = b_1, \tag{5.19}$$

with:

$$\begin{aligned}\alpha_1 &= \frac{(\rho_w C_{p,w} V_w)_1}{\Delta t} + \frac{(\lambda_w A_w)_{1+\frac{1}{2}}}{\Delta x_{2,1}} + h_{f\to w} A_w + 4\, h_{\text{rad}} A_w \left(T_{w,1}^{(k)}\right)^3, \\ \beta_1 &= -\frac{(\lambda_w A_w)_{1+\frac{1}{2}}}{\Delta x_{2,1}}, \\ b_1 &= h_{f\to w} A_w T_f^{(k+1)} + \frac{(\rho_w C_{p,w} V_w)_1}{\Delta t} T_{w,1}^{(k)} + h_{\text{rad}} A_w \left(\left(T_h^{(k+1)}\right)^4 + 3\left(T_{w,1}^{(k)}\right)^4\right). \end{aligned}$$

Exterior Boundary Condition

At the outer boundary ($i = N_x$), heat is lost via convection to the surrounding environment:

$$\dot{q}_{w\to\infty} = h_{w\to\infty}(T_{w,N_x} - T_\infty).$$

Applying the finite-volume method:

$$(\rho_w C_{p,w} V_w)_{N_x} \frac{T_{w,N_x}^{(k+1)} - T_{w,N_x}^{(k)}}{\Delta t} = -h_{w\to\infty} A_w (T_{w,N_x}^{(k+1)} - T_\infty) \tag{5.20}$$

$$- (\lambda_w A_w)_{N_x - \frac{1}{2}} \frac{T_{w,N_x}^{(k+1)} - T_{w,N_x-1}^{(k+1)}}{\Delta x_{N_x,N_x-1}}. \tag{5.21}$$

Rewriting in standard linear form:

$$\gamma_{N_x} T_{w,N_x-1}^{(k+1)} + \alpha_{N_x} T_{w,N_x}^{(k+1)} = b_{N_x}, \tag{5.22}$$

with:

$$\begin{aligned}
\alpha_{N_x} &= \frac{(\rho_w C_{p,w} V_w)_{N_x}}{\Delta t} + \frac{(\lambda_w A_w)_{N_x-\frac{1}{2}}}{\Delta x_{N_x,N_x-1}} + h_{w\to\infty} A_w, \\
\gamma_{N_x} &= -\frac{(\lambda_w A_w)_{N_x-\frac{1}{2}}}{\Delta x_{N_x,N_x-1}}, \\
b_{N_x} &= h_{w\to\infty} A_w T_\infty + \frac{(\rho_w C_{p,w} V_w)_{N_x}}{\Delta t} T_{w,N_x}^{(k)}.
\end{aligned}$$

Full Linear System

To summarize, the heat diffusion equation for the wall can now be expressed as a linear system:

$$A\mathbf{T_w}^{(k+1)} = \mathbf{b}, \tag{5.23}$$

where the vector $\mathbf{b}$ is given by:

$$\mathbf{b} = \begin{bmatrix}
h_{f\to w} A_w T_f^{(k+1)} + \frac{(\rho_w C_{p,w} V_w)_1}{\Delta t} T_{w,1}^{(k)} + h_{\text{rad}} A_w \left(\left(T_h^{(k+1)} \right)^4 + 3 \left(T_{w,1}^{(k)} \right)^4 \right) \\
\frac{(\rho_w C_{p,w} V_w)_2}{\Delta t} T_{w,2}^{(k)} \\
\vdots \\
\frac{(\rho_w C_{p,w} V_w)_{N_x-1}}{\Delta t} T_{w,N_x-1}^{(k)} \\
h_{w\to\infty} A_w T_\infty + \frac{(\rho_w C_{p,w} V_w)_{N_x}}{\Delta t} T_{w,N_x}^{(k)}
\end{bmatrix}. \tag{5.24}$$

Here, $\mathbf{T_w}^{(k+1)}$ denotes the temperature vector for the wall nodes. For a furnace with six walls, it is convenient to solve for each wall separately. We denote the temperature vector for wall i as $\mathbf{T_{w_i}}^{(k+1)}$ for $i = 1, \ldots, 6$, with the following index convention:

- $i = 1$: left wall,

- $i = 2$: right wall,
- $i = 3$: top wall,
- $i = 4$: bottom wall,
- $i = 5$: front wall,
- $i = 6$: back wall.

5.2.2 Solution Method: Thomas Algorithm

Since the coefficient matrix A in Equation (5.23) is tridiagonal, the system can be solved efficiently using the *Thomas algorithm* (Tridiagonal Matrix Algorithm, TDMA) [20]. This method performs a forward elimination followed by backward substitution, reducing the computational complexity from $\mathcal{O}(N_x^3)$ for standard Gaussian elimination to $\mathcal{O}(N_x)$.

The algorithm is numerically stable for diagonally dominant systems, a condition that is satisfied here since

$$|\alpha_i| > |\gamma_i| + |\beta_i|, \quad i = 1, \ldots, N_x, \tag{5.25}$$

as discussed in [2, 12].

At each time step, the tridiagonal system is solved using the Thomas algorithm to obtain the wall temperature distribution $\mathbf{T_w}^{(k+1)}$. The method was implemented directly in MATLAB to solve the discretized wall conduction equation at every iteration.

5.2.3 Furnace System Implicit Temperature Updates

The implicit temperature updates for the walls are handled efficiently using the Thomas algorithm. However, the remaining furnace components——the enclosed furnace air, the thin metal sheet, the alloy element, and the heating elements——do not form a linear system, and therefore require a different approach.

Furnace, Metal Sheet, Alloy and Heater Implicit Updates

Applying implicit Euler, the energy balances for these components are:

$$\begin{aligned}
m_f C_{p,f} \frac{T_f^{(k+1)} - T_f^{(k)}}{\Delta t} &= \dot{Q}_{h\to f}^{(k+1)} + 2 \cdot \dot{Q}_{\text{ms}\to f}^{(k+1)} - \Sigma_{i=1}^{6} \dot{Q}_{f\to w_i}^{(k+1)} - \dot{Q}_{f\to\text{alloy}}^{(k+1)}, \\
m_{\text{alloy}} C_{p,\text{alloy}} \frac{T_{\text{alloy}}^{(k+1)} - T_{\text{alloy}}^{(k)}}{\Delta t} &= \dot{Q}_{f\to\text{alloy}}^{(k+1)}, \\
m_{\text{ms}} C_{p,\text{ms}} \frac{T_{\text{ms}}^{(k+1)} - T_{\text{ms}}^{(k)}}{\Delta t} &= \dot{Q}_{h\to\text{ms}}^{(k+1)} - 2 \cdot \dot{Q}_{\text{ms}\to f}^{(k+1)}, \\
m_h C_{p,h} \frac{T_h^{(k+1)} - T_h^{(k)}}{\Delta t} &= \dot{Q}_{\text{power}}^{(k)} - \dot{Q}_{h\to\text{ms}}^{(k+1)} - \Sigma_{i=1}^{6} \dot{Q}_{h\to w_i}^{(k+1)} - \dot{Q}_{h\to f}^{(k+1)}.
\end{aligned}$$

These equations are nonlinear due to the radiative terms $\dot{Q}_{h\to\mathrm{ms}}$ and $\dot{Q}_{h\to w_i}$, which involve fourth-order temperature dependence. Note that, unlike the wall update, we do not linearize these fourth-order terms here. Doing so would require two additional Taylor expansions around $T_h^{(k)}$ and $T_{\mathrm{ms}}^{(k)}$, which may negatively impact the potential accuracy of the thermal simulation model. By separating the wall update from the other components, we keep the Thomas algorithm for the walls and handle only the 4×4 nonlinear system for the remaining components, which is more tractable.

5.2.4 Solution Method: Newton-Raphson Iteration

The Newton-Raphson method is an iterative technique for solving nonlinear systems [49]. Given a function $\mathbf{F}(\mathbf{T})$ and the root condition $\mathbf{F}(\mathbf{T}) = 0$, the method updates an initial guess $\mathbf{T}^{(k+1)}$ by solving the linearized system:

$$\mathbf{J}(\mathbf{T}^{(k)})\,\Delta\mathbf{T}^{(k)} = -\mathbf{F}(\mathbf{T}^{(k)}), \tag{5.26}$$

where $\mathbf{J}(\mathbf{T}) = \partial\mathbf{F}/\partial\mathbf{T}$ is the Jacobian matrix. The solution increment $\Delta\mathbf{T}^{(k)}$ is then used to update the estimate:

$$\mathbf{T}^{(k+1)} = \mathbf{T}^{(k)} + \Delta\mathbf{T}^{(k)}. \tag{5.27}$$

This procedure is equivalent to linearizing $\mathbf{F}$ around $\mathbf{T}^{(k)}$ via a first-order Taylor expansion:

$$\mathbf{F}(\mathbf{T}^{(\mathbf{k+1})}) \approx \mathbf{F}(\mathbf{T}^{(k)}) + \mathbf{J}(\mathbf{T}^{(k)})(\mathbf{T}^{(\mathbf{k+1})} - \mathbf{T}^{(k)}). \tag{5.28}$$

Setting the linear approximation to zero ($\mathbf{F}(\mathbf{T}^{(k+1)}) = 0$) gives the Newton-Raphson update formula. Because $\mathbf{F}(\mathbf{T})$ is locally smooth and well-approximated by its first-order Taylor expansion near the root, the Newton-Raphson iteration converges quadratically when the initial guess is sufficiently close to the true solution. This rapid convergence makes it particularly effective for solving the nonlinear algebraic systems that arise in implicit time-stepping schemes.

Residual System for Furnace Components

For the implicit update of T_f, T_{alloy}, T_{ms}, and T_h, we define the residual vector:

$$\mathbf{F}(\mathbf{T}^{(k+1)}) = \mathbf{F}(T_f^{(k+1)}, T_{\mathrm{alloy}}^{(k+1)}, T_{\mathrm{ms}}^{(k+1)}, T_h^{(k+1)}) = \begin{bmatrix} F_1 \\ F_2 \\ F_3 \\ F4 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \tag{5.29}$$

with entries:

$$F_1 = T_f^{(k+1)} - T_f^{(k)} - \frac{\Delta t}{m_f C_{p,f}} \left[\dot{Q}_{h\to f}^{(k+1)} + 2 \cdot \dot{Q}_{\mathrm{ms}\to f}^{(k+1)} - \Sigma_{i=1}^{6}\, \dot{Q}_{f\to w_i}^{(k+1)} - \dot{Q}_{f\to\mathrm{alloy}}^{(k+1)} \right], \tag{5.30}$$

$$F_2 = T_{\mathrm{alloy}}^{(k+1)} - T_{\mathrm{alloy}}^{(k)} - \frac{\Delta t}{m_{\mathrm{alloy}} C_{p,\mathrm{alloy}}} \left[\dot{Q}_{f\to\mathrm{alloy}}^{(k+1)} \right], \tag{5.31}$$

$$F_3 = T_{\mathrm{ms}}^{(k+1)} - T_{\mathrm{ms}}^{(k)} - \frac{\Delta t}{m_{\mathrm{ms}} C_{p,\mathrm{ms}}} \left[\dot{Q}_{h\to\mathrm{ms}}^{(k+1)} - 2 \cdot \dot{Q}_{\mathrm{ms}\to f}^{(k+1)} \right], \tag{5.32}$$

$$F_4 = T_h^{(k+1)} - T_h^{(k)} - \frac{\Delta t}{m_h C_{p,h}} \left[\dot{Q}_{\mathrm{power}}^{(k+1)} - \dot{Q}_{h\to\mathrm{ms}}^{(k+1)} - \Sigma_{i=1}^{6}\, \dot{Q}_{h\to w_i}^{(k+1)} - \dot{Q}_{h\to f}^{(k+1)} \right]. \tag{5.33}$$

At each iteration, the temperature increments $\Delta\mathbf{T}^{(k)}$ are obtained by solving the linear system described in Equation (5.26), and the temperatures are updated via Equation (5.27).

Jacobian Matrix and Iteration

The Jacobian $\mathbf{J}$ contains partial derivatives of the residuals with respect to each unknown temperature:

$$\mathbf{J} = \begin{bmatrix} J_{11} & J_{12} & J_{13} & J_{14} \\ J_{21} & J_{22} & J_{23} & J_{24} \\ J_{31} & J_{32} & J_{33} & J_{34} \\ J_{41} & J_{42} & J_{43} & J_{44} \end{bmatrix} = \begin{bmatrix} \frac{\partial F_1}{\partial T_f^{(k+1)}} & \frac{\partial F_1}{\partial T_{\text{alloy}}^{(k+1)}} & \frac{\partial F_1}{\partial T_{\text{ms}}^{(k+1)}} & \frac{\partial F_1}{\partial T_h^{(k+1)}} \\ \frac{\partial F_2}{\partial T_f^{(k+1)}} & \frac{\partial F_2}{\partial T_{\text{alloy}}^{(k+1)}} & \frac{\partial F_2}{\partial T_{\text{ms}}^{(k+1)}} & \frac{\partial F_2}{\partial T_h^{(k+1)}} \\ \frac{\partial F_3}{\partial T_f^{(k+1)}} & \frac{\partial F_3}{\partial T_{\text{alloy}}^{(k+1)}} & \frac{\partial F_3}{\partial T_{\text{ms}}^{(k+1)}} & \frac{\partial F_3}{\partial T_h^{(k+1)}} \end{bmatrix}. \tag{5.34}$$

Analytical expressions for the Jacobian entries are given by:

$$J_{11} = 1 + \frac{\Delta t}{m_f C_{p,f}} \left(h_{h\to f} A_h + 2\, h_{\text{ms}\to f} A_{\text{ms}} + 6\, h_{f\to w} A_w + h_{f\to\text{alloy}} A_{\text{alloy}} \right),$$
$$J_{12} = -\frac{\Delta t}{m_f C_{p,f}} h_{f\to\text{alloy}} A_{\text{alloy}},$$
$$J_{13} = -\frac{\Delta t}{m_f C_{p,f}} 2\, h_{\text{ms}\to f} A_{\text{ms}},$$
$$J_{14} = -\frac{\Delta t}{m_f C_{p,f}} h_{h\to f} A_h,$$

$$J_{21} = -\frac{\Delta t}{m_{\text{alloy}} C_{p,\text{alloy}}} h_{f\to\text{alloy}} A_{\text{alloy}},$$
$$J_{22} = 1 + \frac{\Delta t}{m_{\text{alloy}} C_{p,\text{alloy}}} h_{f\to\text{alloy}} A_{\text{alloy}},$$
$$J_{23} = 0,$$
$$J_{24} = 0,$$

$$J_{31} = -\frac{\Delta t}{m_{\text{ms}} C_{p,\text{ms}}} 2\, h_{\text{ms}\to f} A_{\text{ms}},$$
$$J_{32} = 0,$$
$$J_{33} = 1 + \frac{\Delta t}{m_{\text{ms}} C_{p,\text{ms}}} \left(4\sigma\epsilon_{\text{ms}} F_{h\to\text{ms}} A_{\text{ms}} \left(T_{\text{ms}}^{(k+1)} \right)^3 + 2\, h_{\text{ms}\to f} A_{\text{ms}} \right),$$
$$J_{34} = -\frac{\Delta t}{m_{\text{ms}} C_{p,\text{ms}}} \left(4\sigma\epsilon_{\text{ms}} F_{h\to\text{ms}} A_{\text{ms}} \left(T_h^{(k+1)} \right)^3 \right),$$

$$J_{41} = -\frac{\Delta t}{m_h C_{p,h}} h_{h\to f} A_h,$$
$$J_{42} = 0,$$
$$J_{43} = -\frac{\Delta t}{m_h C_{p,h}} \left(4\sigma\epsilon_{\text{ms}} F_{h\to\text{ms}} A_{\text{ms}} \left(T_{\text{ms}}^{(k+1)} \right)^3 \right),$$

$$J_{44} = 1 + \frac{\Delta t}{m_h C_{p,h}} \left(4\sigma \left(\epsilon_{\text{ms}} F_{h\to\text{ms}} A_{\text{ms}} + \sum_{i=1}^{6} \epsilon_{w_i} F_{h\to w_i} A_{w_i} \right) \left(T_h^{(k+1)} \right)^3 + h_{h\to f} A_h \right).$$

This system is solved iteratively until the residual vector $\mathbf{F}$ falls below a prescribed convergence tolerance. Because the temperature changes between time steps are typically small, using the previous time-step's temperatures as the initial guess usually yields rapid convergence. A further improvement can be achieved by applying simple linear extrapolation from earlier time steps. For example, the estimate

$$T_f^{(k+1)} = 2\,T_f^{(k)} - T_f^{(k-1)},$$

often accelerates convergence, particularly when the temperature evolution is smooth.

Coupling the Newton-Raphson Method with the Implicit Wall Temperature Solver

A central feature of the thermal model is the coupling between the nonlinear Newton-Raphson iteration—used to compute the temperatures of the furnace air, alloy, heater elements, and thin metal sheet—and the implicit solution of the wall conduction equation. Although the residual system described above depends explicitly on the wall surface temperature, this temperature is itself determined through an implicit scheme (e.g., the Thomas algorithm). The overall solution strategy therefore remains fully implicit by employing a tightly coupled iterative procedure.

Specifically, the furnace air energy balance includes the convective heat loss term $\dot{Q}f \to w$, which depends on the current wall surface temperature $\mathbf{T_{w,1}}^{(k+1)}$. Because the wall temperature distribution is updated via an implicit finite-volume scheme, it must be consistent with the evolving furnace air temperature during the Newton-Raphson process. To enforce this consistency, the following procedure is applied within each Newton-Raphson iteration:

1. A trial value for the furnace air temperature $T_f^{(k+1)}$ is assumed.
2. Using this value, the linear system for the wall conduction problem is solved (e.g., with the Thomas algorithm), producing a temporary estimate $T_{w,1}^{(k+1)^*}$.
3. This updated wall temperature is then used in evaluating the nonlinear residual $\mathbf{F}$.
4. After the Newton-Raphson update, convergence of both the bulk state temperatures and the wall distribution is assessed. If convergence is not yet reached, the refined $T_f^{(k+1)}$ serves as input to a new wall temperature solve.

In this way, the implicit nature of the wall conduction update is preserved and dynamically incorporated into the nonlinear Newton-Raphson solver. This ensures that at each time step, the interdependence between the bulk state domain (furnace air, alloy, heater elements and metal sheet) and the boundary wall surface is resolved.

Convergence Criteria

In nonlinear systems solved via iterative methods such as Newton-Raphson, an appropriate convergence criterion is essential to ensure both accuracy and computational efficiency. Generally, convergence can be assessed using either:

- The **norm of the residual vector** $\|\mathbf{F}\|$, which indicates how close the current solution is to satisfying the governing equations.
- The **norm of the solution update** $\|\Delta\mathbf{T}\|$, which reflects how much the solution is changing between successive iterations.

In our implementation, we adopt a *solution-based convergence criterion*, where the iterations are terminated when the updates to both the furnace system state temperatures and the wall temperatures fall below a specified tolerance. Specifically, we define the residual norm $\|\Delta T\|$ and declare convergence when:

$$\|\Delta T\| < \varepsilon, \quad \text{and} \quad \|\Delta T_w\| < \varepsilon,$$

where ε is a prescribed level of tolerance.

This criterion ensures that not only the nonlinear residual system has stabilized, but also that the implicit wall conduction update is consistent with the latest furnace state temperature guess.

Although more rigorous or adaptive strategies could be implemented, such as relative tolerances,

$$\frac{\|\Delta T\|}{\|T\|} < \varepsilon,$$

to account for scaling effects, or monitoring $|\mathbf{F}(T)|$ directly to enforce equation satisfaction to machine precision, the present solution-based criterion provides a practical compromise. It offers a transparent and computationally efficient stopping rule that is sufficient for the current stage of model development.

5.3 Model Verification

Before conducting experimental validation and parameter calibration, it is essential to verify that the implemented model behaves in a physically consistent and numerically correct manner. This verification phase ensures that the numerical methods are properly implemented, the coupling strategies function as expected, and the model responds plausibly to various inputs and edge cases. Below, we describe several verification strategies employed in this work.

5.3.1 Limiting Case Behavior

The model is tested under extreme but controlled conditions where the expected outcome is known a priori:

- **Adiabatic Furnace:** By setting all heat losses (e.g., convective wall losses) to zero and removing the alloy, the system should retain all input energy. The temperature should rise monotonically according to temperature profile prescribed by the power input controller.
- **Zero Power Input:** If no power is supplied, all components should gradually equilibrate to the ambient temperature due to heat losses.

- **Thermal Steady-State:** With constant power and long simulation time, the system should reach a steady temperature distribution where net heat input equals total heat loss.

These checks help ensure the system's transient behavior and steady-state tendencies are physically meaningful.

Starting with the **adiabatic** case, Figure 5.2 shows that the furnace retains all heat when external heat transfer is disabled and the alloy element is removed for simplicity. The power input decreases consistently as the wall temperatures gradually converge toward the internal furnace temperature, which in this test case is prescribed as a constant 500 °C. Eventually, after sufficient time, the power demand approaches zero and the furnace temperature stabilizes, as expected.

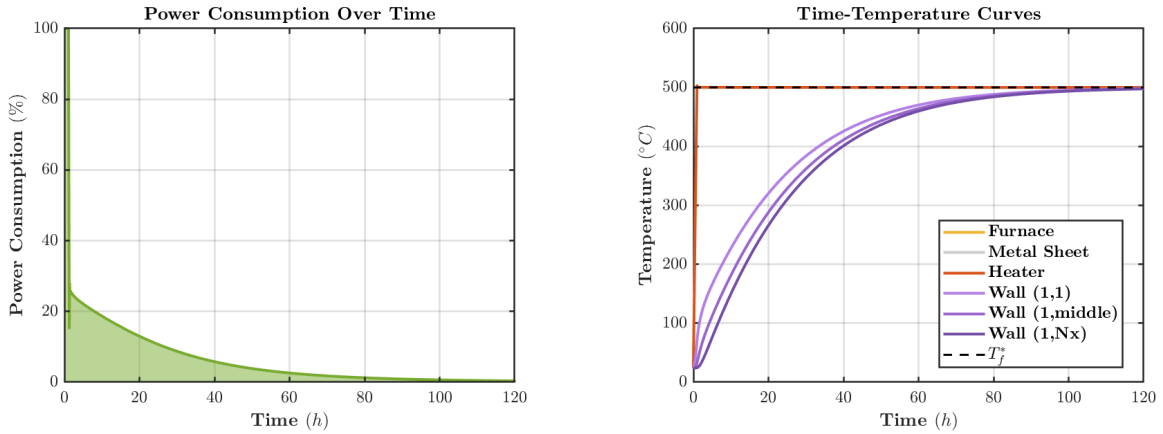

(a) Power curve

(b) Abiatic temperature curves

Figure 5.2: Abiatic Furnace

For the **zero power input** case, the furnace was initialized at 100 °C. Since no power is supplied, the furnace loses heat to the environment and gradually approaches thermal equilibrium with the ambient (room) temperature. Figure 5.3 confirms this expected cooling behavior.

Finally, Figure 5.4 demonstrates the **thermal steady-state** condition. With constant full power input, the internal gas and wall node temperatures asymptotically approach constant values, indicating that the system has reached equilibrium where heat input balances heat loss.

5.3.2 Wall Solver Validation

To verify the accuracy of the wall conduction solver, which employs a finite volume discretization and the Thomas algorithm for solving the resulting tridiagonal system, the implementation was tested against two canonical steady-state conduction problems. These problems were selected because they admit analytical solutions, enabling direct error quantification and solver validation. Agreement with the analytical results provides confidence in both the spatial discretization and the matrix formulation.

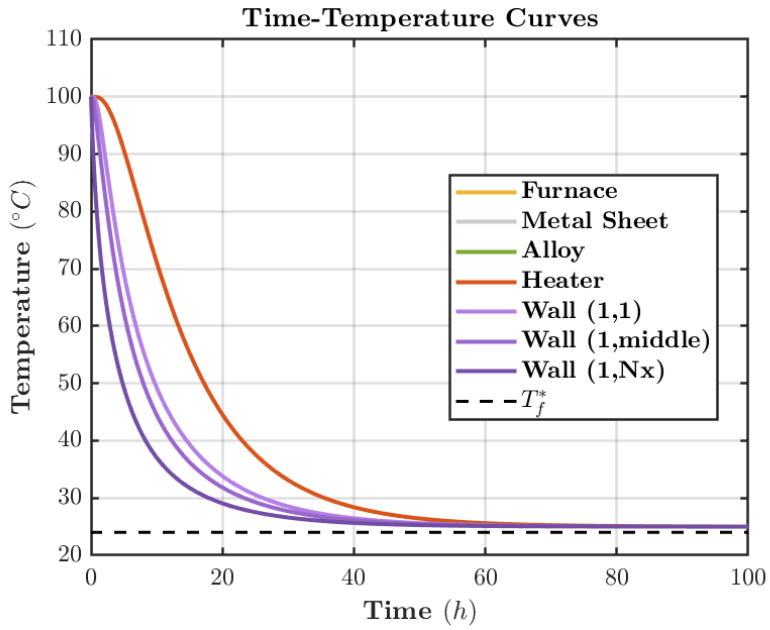

Figure 5.3: Zero Power Input

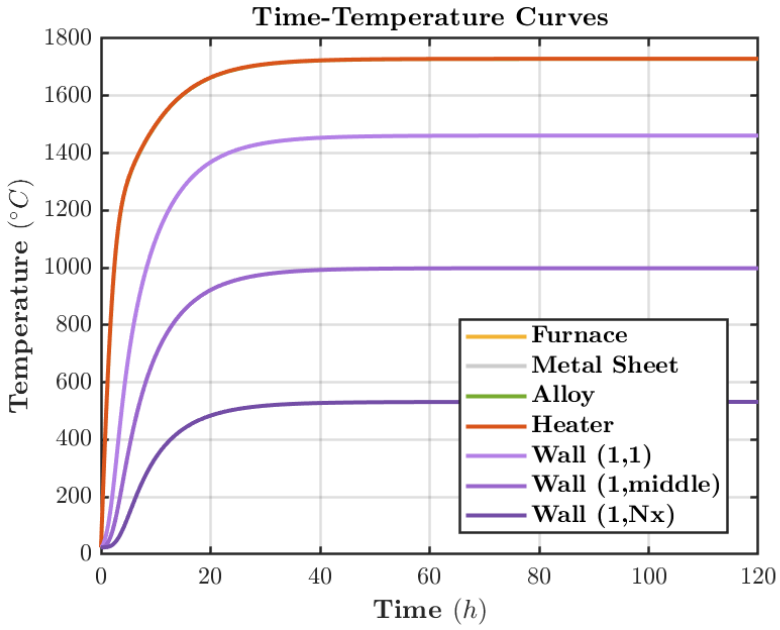

Figure 5.4: Thermal Steady-State

Case 1: Constant Temperature Boundaries

In the first verification case, constant Dirichlet boundary conditions were imposed:

$$T(0) = T_{\text{left}}, \qquad T(L) = T_{\text{right}}, \tag{5.35}$$

where L is the wall thickness. With a uniform thermal conductivity k, the analytical solution for the steady-state temperature profile is linear:

$$T(x) = T_{\text{left}} + \frac{T_{\text{right}} - T_{\text{left}}}{L}\, x. \tag{5.36}$$

This test verifies both the spatial discretization and the correct implementation of the Thomas algorithm. The numerical solutions were compared against the analytical solution across multiple mesh resolutions ($N_x = 5, 10, 20, 50$), as shown in Figure 5.5. The numerical results match the analytical profiles closely, even on finer grids, confirming the solver's correctness for fixed-temperature boundary conditions.

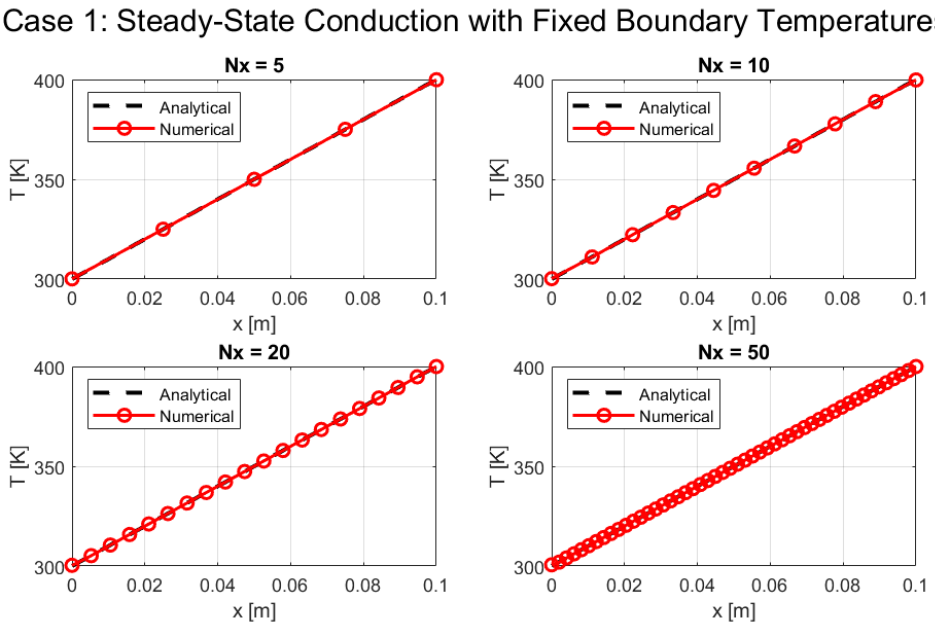

Figure 5.5: Constant Dirichlet boundary conditions

Case 2: Prescribed Heat Flux Boundary

In the second verification case, a constant heat flux q'' is imposed on the left boundary, while the right boundary is held at a fixed temperature:

$$-k \left.\frac{\partial T}{\partial x}\right|_{x=0} = q'', \qquad T(L) = T_{\text{right}}. \tag{5.37}$$

The analytical steady-state solution is:

$$T(x) = T_{\text{right}} + \frac{q''}{k}(L - x). \tag{5.38}$$

This case evaluates the solver's ability to handle mixed boundary conditions, where a Neumann flux condition must be accurately incorporated. As shown in Figure 5.6, the numerical results again agree closely with the analytical solution, demonstrating correct implementation of both heat flux and temperature boundary conditions.

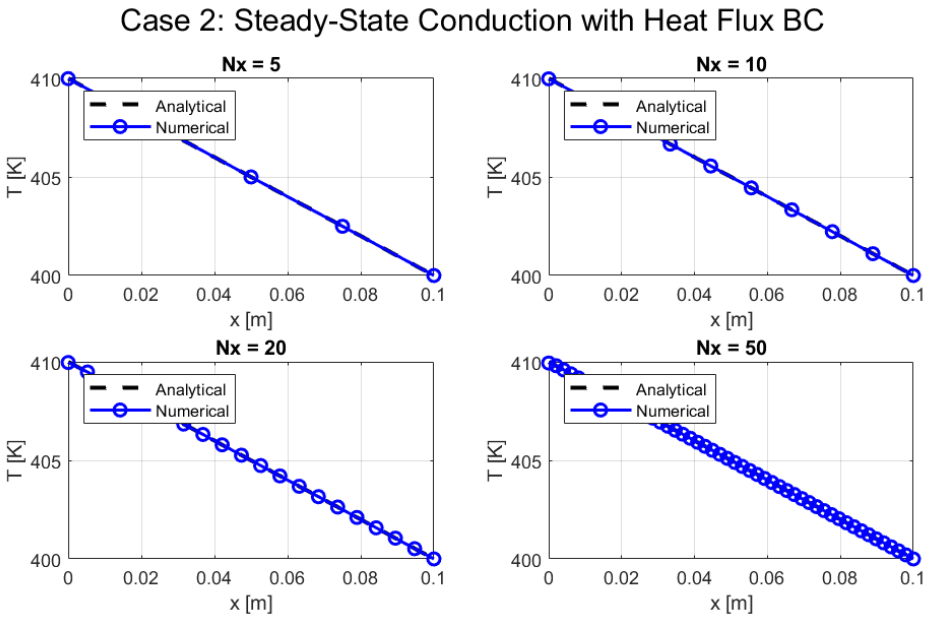

Figure 5.6: Constant heat flux at boundary

Interpretation of Results

The results from both cases confirm that the wall conduction solver:

- Correctly implements the finite volume discretization for steady-state conduction.
- Accurately handles both Dirichlet (fixed temperature) and Neumann (prescribed flux) boundary conditions.
- Produces solutions that converge towards the analytical profile as the spatial resolution increases.

These tests validate the correctness of the spatial discretization and the implementation of the Thomas algorithm within the wall conduction model. With these verifications, the solver can be confidently applied within the full furnace model, where transient effects, convection, and radiation are coupled to the wall conduction problem.

5.3.3 Newton-Raphson Solver Validation

To verify the correctness of the custom Newton-Raphson (NR) solver, its output was compared against MATLAB's built-in nonlinear solver `fsolve`, available in the Optimization Toolbox. The same nonlinear residual system from Equation (5.29), namely

$$\mathbf{F}(\mathbf{T}) = 0,$$

was passed to `fsolve` using identical initial guesses, tolerances, and solver options wherever possible.

The comparison focused on the following aspects:

- **Convergence Agreement:** For a given initial guess, both solvers should converge to the same solution, up to numerical precision.
- **Residual Norms:** The norm of the residual vector at the converged solution should be nearly identical, confirming that both solvers satisfy the same convergence criteria.
- **Execution Time and Efficiency:** While not strictly a verification step, comparing computational performance helps evaluate the trade-offs between using a tailored Newton-Raphson method and a general-purpose solver.

We highlight a representative case: heating the furnace to 400 K over the course of one hour. The resulting temperature curve obtained from both solvers are indistinguishable, which is reflected in Table 5.1, where the residual norms of the converged solutions are nearly identical. Moreover, the execution time of the custom Newton-Raphson solver is substantially lower than that of MATLAB's `fsolve`, supporting the conclusion that the tailored NR solver is more efficient for the specific structure of this problem.

Table 5.1: Comparison of Newton-Raphson and `fsolve` methods for a transient furnace heating simulation.

Quantity	**L_2 Norm Difference**	**Interpretation**
$\|\|T_f^{\mathrm{NR}} - T_f^{\mathrm{FS}}\|\|_2$	2.0×10^{-4}	Negligible
Execution Time		
Newton-Raphson	0.593 s	
`fsolve`	3.939 s	

The representative example confirms that the custom solver produces solutions identical to `fsolve` within machine precision, as was the case for all tested scenarios. This provides strong evidence that the Jacobian matrix and residual formulations are implemented correctly, and that the iterative scheme functions as intended. For practical purposes, the custom solver is adopted in the full furnace simulations due to its tighter integration with the implicit wall temperature solver and its superior computational efficiency for the problem at hand.

5.3.4 Stability Considerations

Stability is a critical aspect of any time-dependent simulation, particularly when solving nonlinear systems involving strong thermal coupling between different components. In our model, stability is enhanced by the use of an implicit time integration scheme (backward-Euler), which is unconditionally stable for linear problems and generally more stable than explicit methods for stiff systems [53].

However, despite the theoretical stability of the time discretization, numerical instabilities may still arise from the nonlinear solution procedure or from inappropriate time step choices. To monitor and assess stability, we perform the following checks:

- **Time Step Sensitivity:** Simulations are run with progressively smaller time steps to ensure that the solution does not significantly change.
- **Wall Node Sensitivity:** Simulations are run with progressively larger numbers of interior wall nodes to ensure that the solution does not significantly change.

Figure 5.7 illustrates the effect of progressively smaller time steps on temperature and power curves. The results highlight that the time step must still be chosen carefully: overly large values can produce non-smooth curves. By contrast, the negligible differences between the results obtained with the smallest time steps (e.g., $\Delta t = 0.2$ s and $\Delta t = 0.1$ s) confirm that the solution stabilizes as Δt decreases. For our purposes, we restrict time steps to 10 seconds or less.

The small relative errors between consecutive simulations with decreasing Δt are further emphasized in Table 5.2. As the table shows, progressively smaller time steps lead to solutions that converge towards stability, with diminishing differences between successive runs.

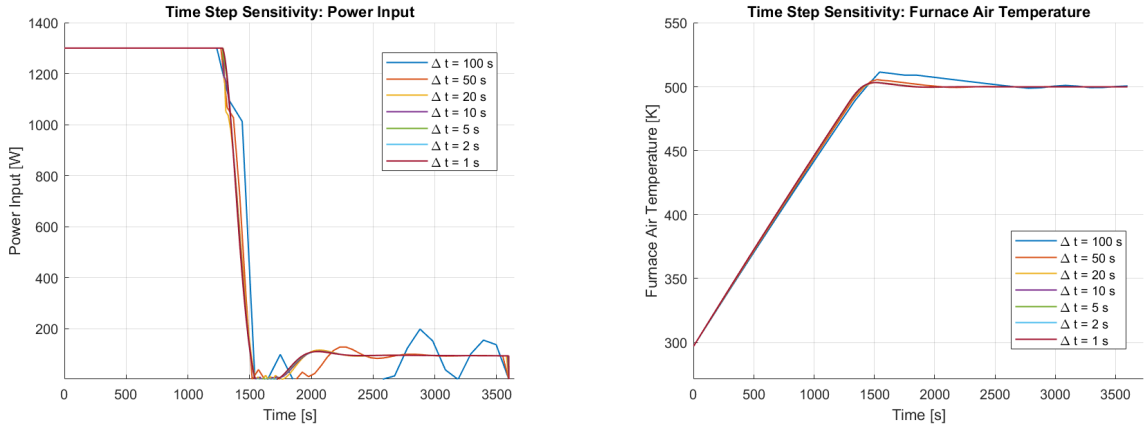

(a) Power curves

(b) Temperature curves

Figure 5.7: Time-step sensitivity analysis: simulations with progressively smaller Δt.

Figure 5.8 presents the wall-node sensitivity analysis. Here, a simple simulation of heating the furnace to 500 K for one hour is used for demonstration.

Table 5.2: Time-step sensitivity analysis for the furnace model. The relative error is computed as the normalized L_2 norm difference between two consecutive simulations.

Δt_1 [s]	Δt_2 [s]	**Relative Error**
100	50	6.83×10^{-3}
50	20	2.11×10^{-3}
20	10	5.33×10^{-4}
10	5	2.69×10^{-4}
5	2	1.61×10^{-4}
2	1	5.37×10^{-5}
1	0.5	2.675×10^{-5}
0.5	0.2	1.571×10^{-5}
0.2	0.1	5.305×10^{-6}

The leftmost plot shows an approximately linear relationship between mesh size and simulation time: as the mesh is refined (i.e., more grid points), the execution time increases proportionally.

The middle plot shows the final temperature values at the end of the simulation for different mesh sizes. The lack of significant variation confirms that the temperature predictions remain stable as the mesh is refined.

Finally, the rightmost plot highlights the relative root-mean-square error between furnace temperature curves and the reference solution obtained with the finest mesh. The results demonstrate that the solution changes negligibly even when the number of wall nodes is increased to very high values. Because simulation time scales linearly with the number of wall nodes, we adopt 15 nodes as a practical choice. This resolution strikes a balance between computational cost and accuracy: the relative difference in root-mean-square error is only about 0.01 compared to the high-resolution case with 405 nodes.

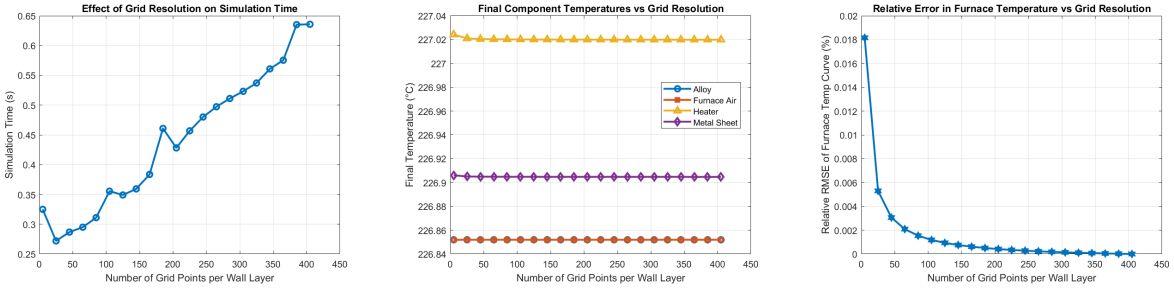

(a) Grid resolution vs. simulation time

(b) Grid resolution vs. final temperature values

(c) Grid resolution vs. relative RMS error

Figure 5.8: Wall-node sensitivity analysis.

6 Parameter Calibration Using Experimental Data

Accurate modeling of industrial furnaces requires more than adherence to fundamental physical laws; it also depends critically on the precise identification of model parameters. While thermal properties, heat transfer coefficients, and other quantities can be estimated from theoretical derivations or literature values, real-world furnace behavior is often affected by unmodeled dynamics, manufacturing tolerances, and environmental variability. Consequently, parameter calibration becomes essential for ensuring that simulation outputs align with experimental observations.

In this chapter, we present a systematic approach to calibrating the transient thermal furnace model developed earlier. A broad set of parameters is considered, including thermal properties, convective and radiative heat transfer coefficients, and free parameters such as heater mass, the number of heaters, PID controller weights, and effective thermal conductivities. These parameters are tuned using measured furnace data, consisting of temperature and power profiles recorded expressly for this purpose.

We formulate calibration as a data-driven optimization problem, where uncertain model parameters are treated as decision variables. Given measured furnace data—including temperature and power profiles—the objective is to minimize a loss function that quantifies the discrepancy between model predictions and experimental results. Conceptually, this parallels supervised learning in machine learning, where the model "learns" parameters from data to improve predictive performance. To ensure robustness and avoid poor local minima, the optimization is performed in two stages: an initial particle swarm optimization explores the parameter space globally, followed by a constrained local refinement using *fmincon*. This hybrid approach effectively balances global exploration with local accuracy.

The calibration process consists of three main components:

1. Parameter selection — identifying the most sensitive and uncertain parameters to optimize.
2. Loss function formulation — defining a metric that captures deviations in both temperature and power profiles over time.
3. Numerical optimization — applying the two-stage optimization strategy to minimize the loss while respecting physical and operational constraints.

By incorporating experimental data into the modeling workflow, we test the predictive capability of the furnace model. This calibrated model forms the foundation for subsequent tasks such as control strategy design and process optimization.

6.1 Selection of Parameters for Calibration

Accurate model calibration requires selecting a subset of parameters that are both **uncertain** and **sensitive** to the model outputs of interest. Optimizing parameters that are either well-known or have negligible influence on the system response would unnecessarily increase computational cost and risk overfitting.

6.1.1 Criteria for Parameter Selection

We distinguish between two key properties:

- **Uncertainty:** Parameters whose values cannot be accurately obtained from literature or manufacturer specifications due to factors such as geometric complexity, surface effects, material variability, or incomplete documentation.
- **Sensitivity:** Parameters whose variation causes significant changes in the model's key outputs, such as the alloy's transient temperature profile or the power required to track a set-point trajectory. Sensitivity can be quantified using local derivative-based metrics (e.g., finite differences) or global techniques (e.g., variance-based sensitivity analysis).

Ideally, only parameters satisfying both criteria are selected for calibration. This ensures that optimization efforts focus on influential unknowns, improving predictive performance without introducing unnecessary degrees of freedom.

6.1.2 Calibrated Parameters

Based on these criteria, the following parameters are treated as decision variables in the optimization problem:

$$\boldsymbol{\theta} = [h_{f\to w},\, h_{f\to\text{alloy}},\, h_{\text{ms}\to f},\, h_{h\to f},\, h_{\text{out}},\, \varepsilon_{\text{ms}},\, \varepsilon_w,\, m_h,\, n_h,\, \lambda_{w,1},\, \lambda_{w,2},\, K_p,\, K_d,\, K_i]\,.$$

Grouped by physical meaning, these include:

1. Convective heat transfer coefficients: $\{h_{f\to w},\, h_{f\to\text{alloy}},\, h_{\text{ms}\to f},\, h_{h\to f},\, h_{\text{out}}\}$. These govern the exchange of heat between furnace air, walls, alloy, metal sheets and heaters.
2. Radiative emissivities: $\{\varepsilon_{\text{ms}},\, \varepsilon_w\}$. These affect radiative heat transfer from the heater elements to the rest of the furnace.
3. Heater characteristics: Lumped heater mass m_h and effective number of heating element n_h. Since the heaters are enclosed and undocumented, their physical properties cannot be directly measured and must be inferred from data.
4. PID controller weights: $\{K_p,\, K_d,\, K_i\}$. These determine how aggressively the furnace controller responds to set-point deviations.
5. Wall thermal conductivities: $\{\lambda_{w,1},\, \lambda_{w,2}\}$. The thermal response of the furnace walls strongly depends on the conductivity of the insulation and steel layers.

It should be noted that we will not attempt to calibrate all of these parameters in one go. Multiple iterations over various subsets of this group, strategically chosen, will reduce computational strains and produce a solution that iteratively becomes a closer match to the measurement data. In the case that calibration suggests further improvement is possible, other parameters may be adjusted within a reasonable range.

Fixed Parameters and Approximations

Parameters that are either well-characterized or appear in the model only in lumped combinations are excluded from calibration. Examples include:

- Heater surface area A_h and specific heat capacity $C_{p,h}$, which always appear in the combined terms $m_h\, C_{p,h}$ and $A_h\, h_{h\to f}$. Fixing one avoids redundancy.
- Alloy geometry and mass, which were measured directly from test samples.
- Known constants such as the Stefan-Boltzmann constant, density of steel, or air properties at ambient conditions.

This approach reduces identifiability issues and focuses optimization on parameters with the greatest influence on predictive accuracy.

Alloy Element Parameter Estimates and Measurements

In this section, we establish reasonable initial guesses for the thermophysical properties of the alloy element used in the furnace model. The alloy under consideration is EN AW-6082, an aluminum alloy primarily composed of Al-Mg-Si. It is commonly used in structural applications and is known for its good corrosion resistance and moderate strength.

In the furnace model, the alloy specimen is represented as a standardized tensile test sample with a "dog-bone" geometry (Figure 6.1). These samples are placed flat on a mesh, exposing most of their surface area to the circulating furnace air. The key geometric parameters, such as length, width, and thickness, were measured manually using a ruler. From these measurements, we compute the specimen's volume V, surface area A, and density ρ using the relationship:

$$m_{\text{alloy}} = \rho\, V. \tag{6.1}$$

The derived density closely matches reference values for aluminum alloys [16], lending confidence to the physical measurements. Table 6.1 summarizes the geometric and physical properties used in the model.

Temperature-Dependent Specific Heat Capacity

Accurate calibration requires temperature-dependent estimates of the alloy's specific heat capacity $C_p(T)$ over the furnace's operating range: using only room-temperature values provide insufficient accuracy for the modeling transient heating. Reported data for aluminum alloys show a consistent increase in C_p with temperature, typically ranging from about $900\,\mathrm{J\,kg^{-1}\,K^{-1}}$ near $300\,K$ to roughly $1100\,\mathrm{J\,kg^{-1}\,K^{-1}}$ at $700\,K$. For instance, Kaschnitz et al. [23] measured

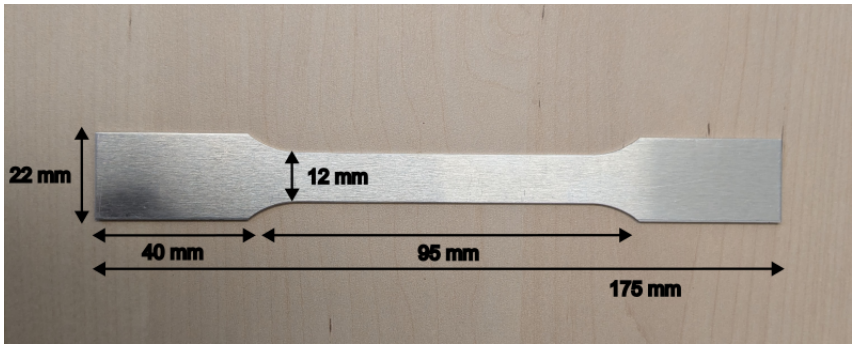

Figure 6.1: Dog-bone tensile test specimen used in furnace measurements

Parameter	Symbol	Value & Units
Geometrical Properties		
Height	h	1.50×10^{-4} m
Width (main body)	w_{mb}	2.20×10^{-3} m
Width (bridge)	w_{b}	1.20×10^{-3} m
Length (main body)	L_{mb}	4.00×10^{-3} m
Length (bridge)	L_{b}	9.50×10^{-3} m
Volume (main body)	V_{mb}	$h \cdot w_{\mathrm{mb}} \cdot L_{\mathrm{mb}}$ [m^3]
Volume (bridge)	V_{b}	$h \cdot w_{\mathrm{b}} \cdot L_{\mathrm{b}}$ [m^3]
Total volume	V	$2 \cdot V_{\mathrm{mb}} + V_{\mathrm{b}}$ [m^3]
Surface area	A	$2 \cdot w_{\mathrm{mb}} \cdot L_{\mathrm{mb}} + w_{\mathrm{b}} \cdot L_{\mathrm{b}}$ [m^2]
Physical Properties		
Density	ρ	2700 kg m^{-3}
Mass	m	$\rho \cdot V$ [kg]

Table 6.1: Geometric, thermal, and physical properties of the aluminum alloy component.

$C_p(T)$ for the Al-7Si-0.3Mg alloy using pulse-heating methods and found results following this trend. Similar magnitudes are suggested for welded heat-treatable Al-Mg-Si alloys in the study by Wiechmann et al. [55], who used differential scanning calorimetry (DSC). On this basis, the values listed in Table 6.2 were adopted as representative estimates for the present simulations.

Temperature [K]	C_p ($\mathrm{J\,kg^{-1}\,K}$)
300	900
400	950
500	1000
600	1050
700	1100

Table 6.2: Alloy Temperature-Dependent Specific Heat Values

To represent $C_p(T)$ continuously, a linear fit was applied to the tabulated data:

$$C_p(T) = a + bT, \tag{6.2}$$

where the coefficients after solving are:

$$C_p(T) = 750 + 0.5\,T \qquad [\mathrm{J\,kg^{-1}\,K^{-1}}]. \tag{6.3}$$

At higher temperatures this linear relationship may not hold, but this linear model approximately captures the observed temperature dependence of C_p across the measured range. Potential mechanical contributions to C_p under stress due to microstructural changes are neglected for simplicity.

Heat Transfer Coefficients Estimate

Realistic estimates for the convective heat transfer coefficients are derived from engineering correlations and tabulated values for air, as found in [50]. These estimates serve as initial guesses and inform the parameter bounds during calibration:

- Furnace air to alloy surface:

$$h_{f\to\mathrm{alloy}} \approx 15\text{-}25\,\mathrm{W\,m^{-2}K^{-1}}.$$

- Metal sheet surface to furnace air:

$$h_{\mathrm{ms}\to f} \approx 15\text{-}25\,\mathrm{W\,m^{-2}K^{-1}}.$$

- Furnace air to interior wall surface:

$$h_{f\to w} \approx 10\text{-}20\,\mathrm{W\,m^{-2}K^{-1}}.$$

- Heater surface to furnace air:

$$h_{h\to f} \approx 25\text{-}100\,\mathrm{W\,m^{-2}K^{-1}}.$$

- Wall exterior to ambient air:

$$h_{\text{out}} \approx 5\text{-}15\,\text{W}\,\text{m}^{-2}\text{K}^{-1}.$$

For calibration, these coefficients are constrained to physically plausible bounds, typically

$$0 \leq h \leq 50\,\text{W}\,\text{m}^{-2}\text{K}^{-1},$$

the exception being $h_{h\to f}$, which is allowed up to $100\,\text{W}\,\text{m}^{-2}\text{K}^{-1}$ to account for forced convection induced by the fan.

Radiative Emissivity Estimates

Accurate modeling of radiative heat transfer requires realistic values for the emissivity of surfaces within the furnace. Because emissivities are treated as free calibration parameters, only literature-based initial estimates are required. We estimate emissivity values for the thin metal sheet separating the heater elements from the main furnace chamber, as well as for the furnace walls, based on material and surface condition assumptions.

Metal Sheet Emissivity (ε_{ms}): The thin metal sheet is assumed to be made of 3 mm steel, likely oxidized due to repeated thermal cycling. Emissivity for oxidized steel surfaces at elevated temperatures (400-800°C) is typically in the range

$$\varepsilon_{\text{ms}} \approx 0.70\text{-}0.85.$$

For modeling purposes, a nominal value of $\varepsilon_{\text{ms}} = 0.80$ is adopted, consistent with literature values [22, 28].

Wall Emissivity (ε_w): The furnace walls are modeled as high-temperature rock wool insulation panels, which are non-metallic with a rough surface texture. These refractory materials typically exhibit high thermal emissivity, in the range

$$\varepsilon_{\text{wall}} \approx 0.85\text{-}0.95.$$

For calibration, we adopt $\varepsilon_{\text{wall}} = 0.90$, consistent with standard references [22, 28, 15].

Furnace Parameters

The furnace volume is known a priori from physical measurements and schematics. Additionally, the furnace air density is computed as described in Section 4.2.3. The only remaining parameter to approximate is the specific heat capacity, which depends on the combustion gas composition (primarily N_2, O_2, CO_2, and H_2O). At elevated temperatures (500-1000°C), water vapor and carbon dioxide increase the effective heat capacity relative to dry air. Using representative molar fractions (concentration of components in furnace air) and high-temperature data [34, 9, 51], an average estimate gives:

$$C_{p,f} \approx 1150\ \text{J}\,\text{kg}^{-1}\,\text{K}^{-1}.$$

This value is used as a constant approximation throughout the simulation.

Wall Parameters

The furnace walls are modeled as two layers: high-temperature rock wool insulation and a thin structural steel sheet. Layer 1 (rock wool) dominates thermal behavior, while the steel layer is sufficiently thin such that its influence on the overall heat transfer is expected to be minor. According to standard datasheets and industrial manufacturers like ROCKWOOL Technical Insulation [42], the parameters can be approximated as follows ([15, 7]):

- Thermal conductivity, layer 1 (rock wool): $\lambda_{w,1} \approx 0.05\ \mathrm{W\,m^{-1}K^{-1}}$.
- Density, layer 1: $\rho_{w,1} \approx 200\ \mathrm{kg\,m^{-3}}$.
- Specific heat capacity, layer 1: $C_{p,w,1} \approx 900\ \mathrm{J\,kg^{-1}K^{-1}}$.
- Thermal conductivity, layer 2 (steel): $\lambda_{w,2} \approx 50\ \mathrm{W\,m^{-1}K^{-1}}$; the layer is thin, limiting its effect on the model.
- Density, layer 2: $\rho_{w,2} \approx 7800\ \mathrm{kg\,m^{-3}}$.
- Specific heat capacity, layer 2: $C_{p,w,2} \approx 500\ \mathrm{J\,kg^{-1}K^{-1}}$.

Metal Sheet Parameters

The thin metal sheet mediates radiative-convective heat transfer between the heater elements and the furnace chamber. The sheet mass was estimated from measurements with a thickness of $t = 3\,\mathrm{mm}$, a total surface area of $A_{\mathrm{ms}} \approx 1.4112\,\mathrm{m^2}$, and a steel density of $\rho \approx 7850\,\mathrm{kg\,m^{-3}}$. The specific heat capacity is approximated from typical steel values. The sheet mass was computed from the known geometry and density of steel:

$$m_{\mathrm{ms}} = \rho_{\mathrm{steel}}\, A_{\mathrm{ms}}\, t = 7850\,\mathrm{kg\,m^{-3}} \times 1.4112\,\mathrm{m^2} \times 3 \times 10^{-3}\,\mathrm{m} \approx 33.2\,\mathrm{kg}.$$

Its physical parameters are summarized in Table 6.3.

Table 6.3: Parameters for the thin metal sheet separating the heating elements from the main furnace chamber.

Parameter	**Symbol**	**Value**	**Units**
Mass	m_{ms}	33.2	kg
Specific heat capacity	$C_{p,\mathrm{ms}}$	500	$\mathrm{J{\cdot}kg^{-1}{\cdot}K^{-1}}$
Emissivity	$\varepsilon_{\mathrm{ms}}$	0.8	–
Area	A_{ms}	$4 \times 0.49 \times 0.72$	$\mathrm{m^2}$

Heater Parameters

The heating elements' exact physical properties are unspecified. Their dynamic response is captured using an effective lumped thermal capacitance $m_{\mathrm{h}}\, C_{p,\mathrm{h}}$. For modeling purposes:

- Specific heat capacity: $C_{p,\mathrm{h}} = 100\,\mathrm{J\,kg^{-1}\,K^{-1}}$, reflecting a reduced effective thermal mass.

- $A_\mathrm{h} = 1\,\mathrm{m}^2$, since its dependency on heat transfer is captured by $h_{h\to f}$ in the combined term $A_\mathrm{h}\, h_{h\to f}$.
- Effective mass m_h is treated as a free calibration parameter.

This approach allows the model to reproduce the furnace thermal response despite missing physical data.

PID Controller Parameters

The furnace controller is modeled using a PID feedback loop with the following initial parameters:

- Maximum power: $P_\mathrm{max} = 13{,}000\,\mathrm{W}$ (from furnace manual).
- Proportional gain: $K_p = 100$. Selected to provide a strong initial response to temperature deviations while avoiding excessive overshoot.
- Integral gain: $K_i = 1$. Included to eliminate steady-state error in the temperature control, with a modest gain to prevent integral windup.
- Derivative gain: $K_d = 10$. Helps dampen oscillations in the system response, providing stability in the presence of rapid temperature changes.
- Number of heating elements: $n_\mathrm{heaters} = 18$, treated as a free calibration parameter due to uncertainty.

Initial PID gains were manually tuned to achieve a stable thermal response. In the final calibration stage, the gains (K_p, K_i, K_d) are optimized while all other parameters are fixed, enabling the controller to match measured furnace dynamics [1].

Calibratable Parameters and Bounds

Table 6.4 summarizes all the parameters in the furnace model that are subject to calibration. These include geometric, convective, radiative, material, and control parameters. Initial guesses and physically reasonable bounds are provided for each parameter, consistent with the discussion in the preceding subsections.

6.2 Calibration Loss Function

To calibrate the furnace model parameters, a loss function is required to quantify the discrepancy between simulated and experimentally measured temperature profiles. In the machine learning analogy, this corresponds to regression objective such as mean squared error (MSE) that penalizes deviations between the model output and experimental data.

Two distinct experimental datasets are available: one corresponding to a step-wise heating schedule with intermediate holds (*step* dataset), and another representing a continuous monotonic heating ramp (*ramp* dataset). To ensure that the calibrated parameters are consistent across

Table 6.4: Calibratable parameters in the furnace model, including initial guesses, bounds, and units.

Parameter	Initial Guess	Bounds	Units	Description
Convective and Radiative Parameters				
$h_{f\to\text{alloy}}$	20	$[0, 50]$	$\mathrm{W\,m^{-2}\,K^{-1}}$	Conv. coeff. furnace → alloy
$h_{\text{ms}\to f}$	20	$[0, 50]$	$\mathrm{W\,m^{-2}\,K^{-1}}$	Conv. coeff. sheet → furnace
$h_{f\to w}$	10	$[0, 50]$	$\mathrm{W\,m^{-2}\,K^{-1}}$	Conv. coeff. furnace → wall
$h_{h\to f}$	25	$[0, 100]$	$\mathrm{W\,m^{-2}\,K^{-1}}$	Conv. coeff. heater → furnace
h_{out}	10	$[0, 50]$	$\mathrm{W\,m^{-2}\,K^{-1}}$	Heat loss to ambient
ε_{ms}	0.8	$[0.01, 1]$	–	Metal sheet emissivity
ε_w	0.9	$[0.01, 1]$	–	Furnace wall emissivity
Material Properties				
$\lambda_{w,1}$	0.05	$[0.01, 1]$	$\mathrm{W\,m^{-1}\,K^{-1}}$	Wall layer 1 thermal conductivity
$\lambda_{w,2}$	50	$[1, 100]$	$\mathrm{W\,m^{-1}\,K^{-1}}$	Wall layer 2 thermal conductivity
m_h	10	$[1, 100]$	kg	Heater effective mass
Controller Parameters				
K_p	100	$[0, 1000]$	–	PID proportional gain
K_i	1	$[0, 10]$	–	PID integral gain
K_d	10	$[0, 100]$	–	PID derivative gain
n_{heaters}	18	$[8, 20]$	–	Number of heating elements

both operating regimes, the calibration is formulated as a simultaneous optimization problem. The total loss function is therefore expressed as

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{s\in\{f,\ \text{alloy}\}} \sum_{j\in\{\text{step},\ \text{ramp}\}} \sum_{i=1}^{N_j} w_s\, w_j \left(T_{s,\text{model},j}(t_i;\boldsymbol{\theta}) - T_{s,\text{measured},j}(t_i)\right)^2, \tag{6.4}$$

where $\boldsymbol{\theta}$ denotes the vector of calibration parameters, N_j is the number of time samples in dataset j, and w_j are weighting coefficients that balance the relative influence of the two datasets.

In this formulation, $T_{s,\text{model},j}(t_i;\boldsymbol{\theta})$ represents the simulated temperature at time t_i for dataset j, and $T_{\text{measured},j}(t_i)$ is the corresponding measured temperature, where $s \in \{f,\ \text{alloy}\}$ denotes the index representing the furnace air temperature T_f or the alloy material temperature T_{alloy}.

The optimization problem consists of finding the parameter vector $\boldsymbol{\theta}^*$ that minimizes the total loss:

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}}\ \mathcal{L}(\boldsymbol{\theta}). \tag{6.5}$$

This procedure is carried out using the numerical optimization strategy discussed in Section 6.4. It should be noted, however, that perfect reproduction of the experimental data is neither expected nor desirable, as measurement noise, unmodelled disturbances, and physical simplifications can cause unavoidable deviations. The goal of the calibration is instead to achieve a balanced compromise that captures the essential system dynamics while maintaining physically interpretable parameter values.

6.3 Measured Data and Signal Processing

In this section, the measured temperature and power profiles from the furnace experiments are analyzed and pre-processed for subsequent model calibration. Two experimental datasets are available, each containing the time evolution of the supplied power and corresponding furnace temperature measurements.

Raw Data

Figures 6.2 and 6.3 show the measured furnace temperature and power profiles for the two datasets.

1. “Step” dataset: A three-step heating profile with a hold at intermediate temperature, followed by a natural cooling cycle.

2. “Ramp” dataset: A monotonic heat ramp followed by a brief hold at $300\,K$.

While the overall heating trends are clear, the power measurements exhibit significant high-frequency fluctuations and apparent noise artifacts. These can likely be attributed to measurement uncertainty and sensor resolution limitations, rapid actuation of the furnace’s internal control system, and possible discretization effects from the data acquisition hardware.

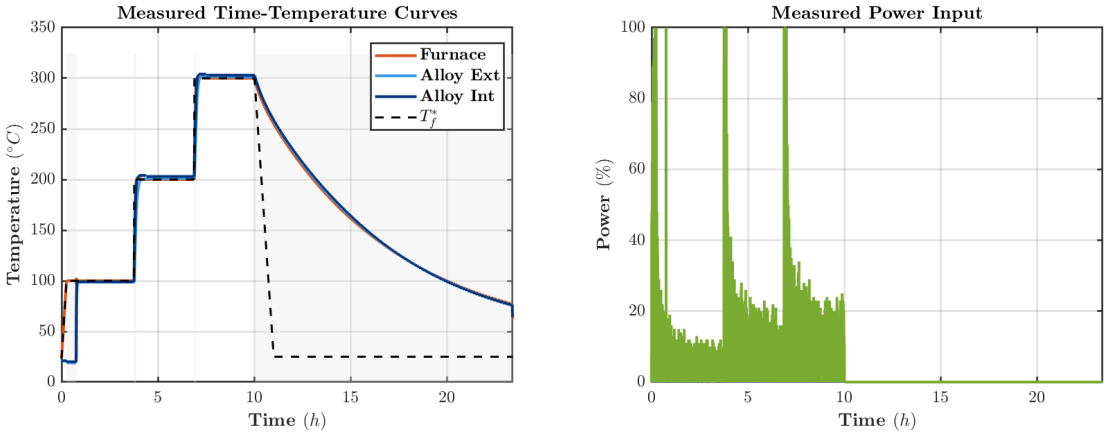

Figure 6.2: Raw measured temperature and power profiles for “step” dataset

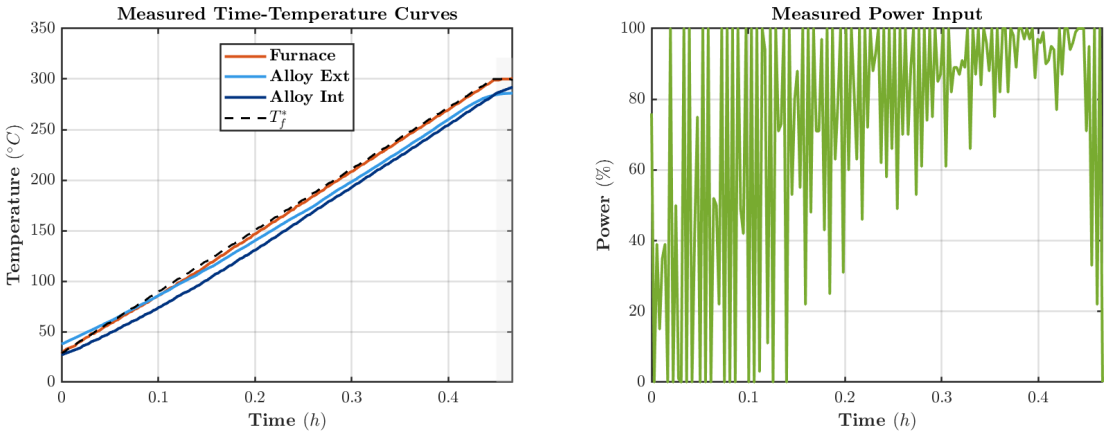

Figure 6.3: Raw measured temperature and power profiles for “ramp” dataset

Observations on the datasets

"Step" dataset

The step dataset consists of three heating stages with an intermediate temperature hold, followed by a natural cooling cycle during which the power is turned off. The alloy sample is introduced into the furnace with a short delay after heating begins, and two independent sensors recorded the alloy temperature: one external and one internal. A small dip in the measured alloy temperature occurs just before insertion into the furnace; this is assumed to be sensor-related and is otherwise irrelevant. For calibration purposes, the external sensor data is used, because in our model the alloy is represented as a lumped thermal mass without internal gradients. Additionally, both the internal alloy and furnace air temperature measurements from the data are rounded to integers, reducing accuracy. The external alloy measurements therefore represent a higher degree of accuracy.

"Ramp" dataset

The ramp dataset features a continuous temperature increase followed by a short hold, and includes a subsequent cooling phase in which several "slits" are first opened to increase airflow into the furnace, followed by the furnace door itself, facilitating faster heat loss. These additional interactions with the slits and furnace door are not considered in the present model and are therefore omitted. A closer inspection of the dataset reveals several inconsistencies: the recorded time steps are irregular, and furnace temperatures are rounded to integer values, reducing accuracy; in the initial phase, the measured alloy temperature appears higher than the furnace air temperature, which is physically unlikely; and the power measurements exhibit extreme high-frequency oscillations and sudden near-zero spikes.

6.3.1 Power Signal Preprocessing and Smoothing

The power-related flags listed above suggest that the power may have been controlled using an on-off thyristor-based PID controller combined with a sensor that records instantaneous rather than averaged power [1]. As a result, the signal appears unsuitable for direct use and may require preprocessing. To assess whether these artifacts could affect model calibration, a preprocessing and smoothing procedure was developed as described below.

Preprocessing Procedure

The measured power signals $P(t)$ were preprocessed using the following steps:

1. **Normalization:** The raw power signal $P(t)$ was normalized as a percentage of the known maximum power capacity $P_{\max}$:

$$P_{\%}(t) = \frac{P(t)}{P_{\max}} \cdot 100.$$

2. **Detection of shutdown region:** To preserve the actual shutdown behavior of the system, a shutdown index t_{shut} was identified as the first instance of where at least N_z consecutive readings satisfied:

$$P_{\%}(t) < \delta_{\text{noise}},$$

where δ_{noise} is a small threshold (e.g., 10%) and N_z is the minimum number of consecutive near-zero values (e.g., 10 samples).

3. **Interpolation of spurious low values:** For $t < t_{\text{shut}}$, low-power values were flagged for interpolation if surrounded by confidently high-power values:

$$P_{\%}(t) < \delta_{\text{noise}}, \quad \text{and} \quad P_{\%}(t \pm \Delta t) > \delta_{\text{trust}},$$

where δ_{trust} is a threshold indicating high-confidence values (e.g., 90%). These suspicious points were replaced via shape-preserving piecewise cubic interpolation over time.

4. **Smoothing via moving average:** The resulting interpolated signal $\tilde{P}(t)$ was then smoothed using a moving average filter with window size w:

$$\tilde{P}(t) = \frac{1}{w} \sum_{k=-w/2}^{w/2} P(t + k\Delta t).$$

5. **Preservation of high-power segments:** To ensure that genuine full-power segments were not attenuated, any values with $P_{\%}(t) > \delta_{\text{trust}}$ were copied directly from the original signal to the final processed output:

$$P_{\text{processed}}(t) = \begin{cases} P(t), & \text{if } P_{\%}(t) > \delta_{\text{trust}}, \\ \tilde{P}(t), & \text{otherwise.} \end{cases}$$

This approach produced a smooth and physically consistent power profile that suppressed measurement artifacts while retaining authentic high-power regions (Figure 6.4). The resulting curves, using a window size of 15, are shown in Figure 6.4.

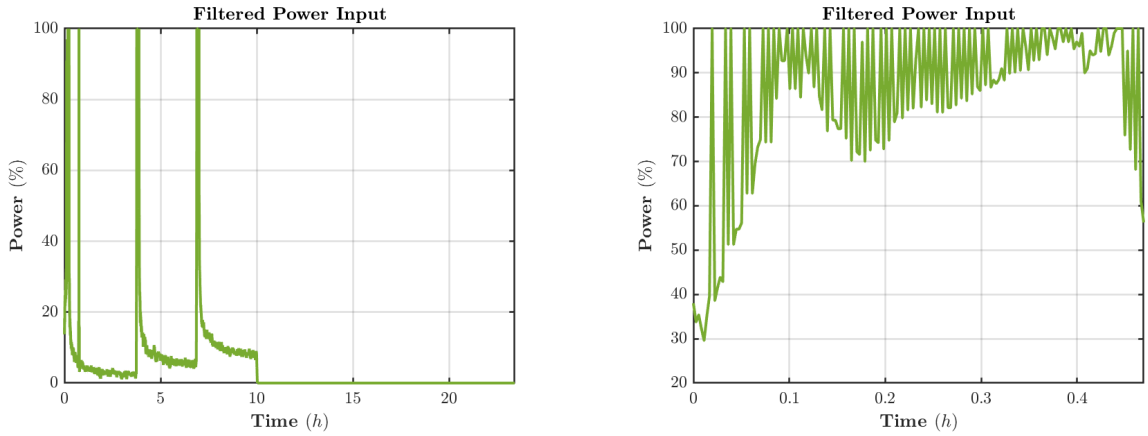

(a) "Step" power input

(b) "Ramp" power input

Figure 6.4: Smoothed and filtered power profiles

Evaluation and Final Choice

Both the raw and smoothed power profiles were tested during model calibration. Interestingly, calibration using the raw (unfiltered) data resulted in more realistic parameter estimates and overall model performance. This suggests that the apparent noise in the raw power data reflects genuine control system behavior rather than purely measurement error. Consequently, the raw power signals were retained for the final model calibration, rather than the smoothed curves. It is important to note that the true instantaneous heater power may nonetheless deviate from the raw measurements due nmeasured system dynamics, while rounding errors, sensor inaccuracies, and irregular sampling intervals introduce residual uncertainties.

6.4 Calibration Methods

Calibrating the furnace model involves minimizing the loss function $\mathcal{L}(\boldsymbol{\theta})$ defined in Section 6.2, where $\boldsymbol{\theta}$ is the vector of uncertain and sensitive model parameters. This is a nonlinear least-squares problem that may be high-dimensional, non-convex, and potentially noisy due to simulation artifacts or measurement error. As such, careful selection of an optimization strategy is prudent.

Problem Statement

We seek to solve the following optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\theta}} \quad & \mathcal{L}(\boldsymbol{\theta}), \\ \text{subject to} \quad & \boldsymbol{\theta}_{\min} \leq \boldsymbol{\theta} \leq \boldsymbol{\theta}_{\max}, \end{aligned}$$

where $\boldsymbol{\theta}_{\min}$ and $\boldsymbol{\theta}_{\max}$ are lower and upper bounds based on physical plausibility or prior knowledge as outlined in Section 6.1. The constraints enforce reasonable parameter values, helping to prevent overfitting and unphysical behavior.

6.4.1 Gradient-Free vs. Gradient-Based Approaches

Two major classes of numerical methods are available for this task:

Gradient-free methods. These include algorithms such as: Nelder-Mead simplex, Particle Swarm Optimization (PSO), Genetic Algorithms (GA), Covariance Matrix Adaptation Evolution Strategy (CMA-ES), and Bayesian Optimization. These methods are well-suited for noisy or discontinuous objective functions and do not require derivative information. They are particularly useful when the simulation is expensive or when the loss surface has many local minima. However, they tend to be slower and may require many function evaluations.

Gradient-based methods. If the model is sufficiently smooth, methods such as conjugate gradient descent can be applied: either with finite-difference gradients or analytical/automatic differentiation. These methods offer faster convergence, especially near local optima, but require the loss function to be continuously differentiable (i.e., $\mathcal{C}^1$-smooth) so that its gradient is well-defined and behaves predictably.

6.4.2 Chosen Optimization Strategy

The objective function is nonconvex, and potentially discontinuous due to iterative solvers and conditional statements. Although automatic differentiation (AD) would, in principle, enable efficient gradient computation, several practical challenges in MATLAB—such as limited AD support for implicit solvers, linear system calls, and custom time-stepping routines—made it infeasible within this project's scope. Given these characteristics, a purely gradient-based approach would be prone to premature convergence to local minima or even numerical instability in regions where derivatives are ill-conditioned or poorly defined [35]. To address these issues, a two-stage hybrid strategy was adopted:

1. **Global exploration:** A gradient-free metaheuristic, Particle Swarm Optimization (PSO), is first employed to perform a broad global search. PSO is well-suited for complex, nonconvex optimization landscapes with interacting parameters [14, 44], and has been successfully applied to a wide range of engineering optimization and thermal system identification problems [36].

2. **Local refinement:** The best-performing solution from the PSO stage is used to initialize a local, gradient-based optimization using `fmincon`. This second stage exploits local curvature information to refine the solution and is run iteratively on different subsets of the parameter space to explore promising solutions.

6.4.3 Particle Swarm Optimization

Particle Swarm Optimization (PSO) is a population-based, derivative-free global optimization metaheuristic inspired by social foraging behaviors in bird flocks and fish schools [14, 44]. Each candidate solution, or *particle*, moves through the search space under the influence of:

- Its own best-found position $\mathbf{p}_i$ (the "cognitive" component), and
- The swarm's global best $\mathbf{g}$ (the "social" component).

At iteration t, particle i has position $\mathbf{x}_i^{(t)}$ and velocity $\mathbf{v}_i^{(t)}$, updated as:

$$\mathbf{v}_i^{(t+1)} = w\,\mathbf{v}_i^{(t)} + c_1\,r_1\big(\mathbf{p}_i - \mathbf{x}_i^{(t)}\big) + c_2\,r_2\big(\mathbf{g} - \mathbf{x}_i^{(t)}\big), \quad \mathbf{x}_i^{(t+1)} = \mathbf{x}_i^{(t)} + \mathbf{v}_i^{(t+1)},$$

where w is the inertia weight, c_1 and c_2 are acceleration coefficients, and $r_1, r_2 \sim U(0,1)$. The inertia term w balances exploration (larger w) against exploitation (smaller w), while the cognitive and social terms guide particles toward promising regions [36]. At each iteration, PSO evaluates the loss over the swarm, updates personal and global bests, and advances the particles. Convergence is declared when either a maximum number of iterations is reached, the best-so-far fitness stalls for a specified number of steps, or an objective tolerance is met [27]. For our furnace-model calibration, PSO's ability to explore nonconvex, noisy loss landscapes without requiring gradients makes it a particularly attractive option.

6.5 Calibration Results

This section presents the results of the furnace model calibration procedure described in Section 6.4. The objective is to determine a single parameter set that reproduces the measured thermal response of the physical furnace under multiple distinct heating protocols. Two experimental datasets are available for this purpose: a multi-step heating schedule with intermediate holds, and a monotonic heating ramp. Note that calibrating the model against both datasets requires balancing trade-offs, as the two measured profiles differ significantly in duration and dynamics. In particular, the "step" dataset spans nearly 24 hours, while the "ramp" dataset covers only a 30-minute heating experiment. The "step" dataset thus contains far more data points than the "ramp" dataset, which would bias a naive least-squares fit toward the long-duration profile. By introducing weights in the composite loss function (Equation (6.4)), we explicitly balance the contribution of each dataset, effectively telling the optimizer to listen to both datasets proportionally. After manual tuning of these weights, the final calibrated solution captures the main thermal behaviors of both scenarios, as shown in Figure 6.5.

Replacing the measured power inputs with a tuned PID controller, as in Figure 6.6, further improves the model's agreement with the temperature profiles in both datasets, while keeping all parameters fixed. Figure 6.7 illustrates this improvement, which implies that the raw power measurements may not reflect the actual energy input as closely as would be desireable.

Figure 6.8 showcases the full thermal dynamics of the calibrated model. Since the only temperature measurements provided were for the furnace air and the alloy element, there is no way to validate the other temperature curves like those of the heater or the walls.

6.5.1 Discussion

Table 6.5 summarizes the parameters obtained from the calibration efforts. Several trends emerge that are worth commenting on. The wall emissivity ε_w is much lower than initially expected, possibility due to shielding effects that reduce direct radiative heat transfer from the heaters to the walls. The convective heat transfer coefficient $h_{f\to w}$ is also significantly lower than predicted, reflecting the high insulation efficiency of the furnace walls. Consequently, the walls physically heat up slower, which aligns with the low calibrated coefficients. Conversely, the metal sheet emmissivity converged to unity, implying that the calibration treats it as a blackbody surface. Whether this is physically realistic or an overcompensation effect from the optimizer cannot be verified with the limited data available.

Another plausible explanation for the very low wall-related coefficients is that the internal air circulation may have been highly localized near the alloy and the heating compartments, resulting in minimal air movement adjacent to the wall surfaces. This would effectively create stagnant air layers that act as an additional insulating boundary, significantly reducing convective and radiative coupling to the wall. The original assumption that forced convection through the fan facilitated a spatially uniform temperature distribution within the internal air may not reflect the physical reality, but does help keep the model simple. Thus, the low calibrated values may represent effective, lumped parameters capturing these unmodeled spatial effects rather than true material wall properties.

Furthermore, the calibrated lumped mass m_h of the heating elements converged to a relatively

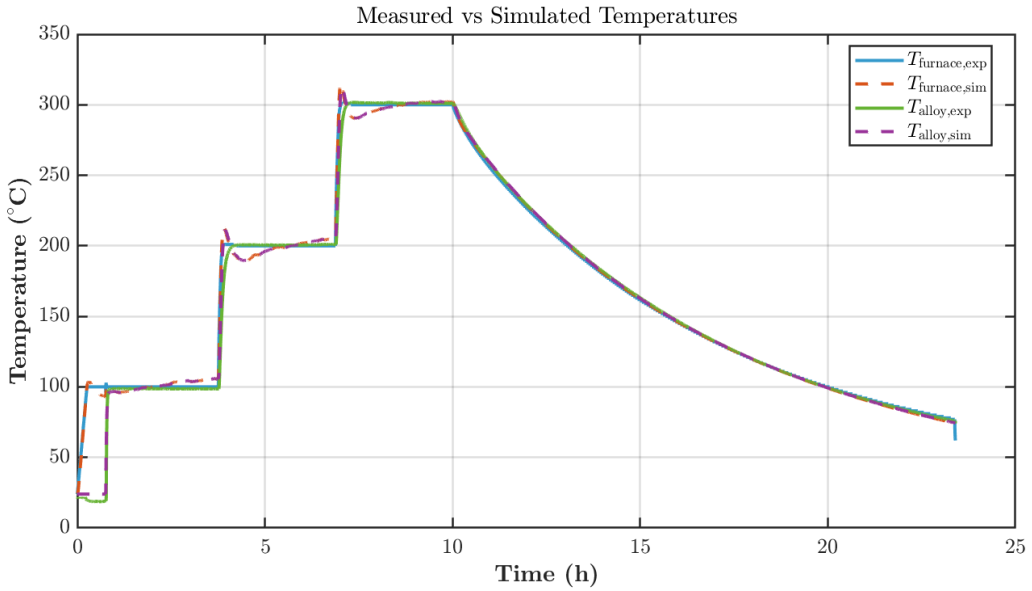

(a) “Step” dataset

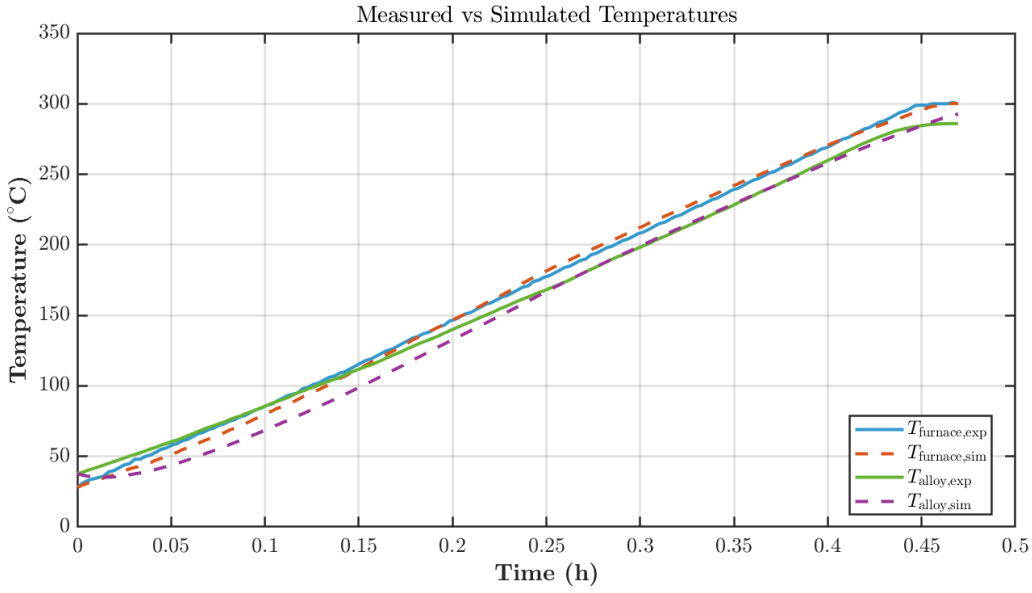

(b) “Ramp” dataset

Figure 6.5: Simulated vs. measured temperature profiles

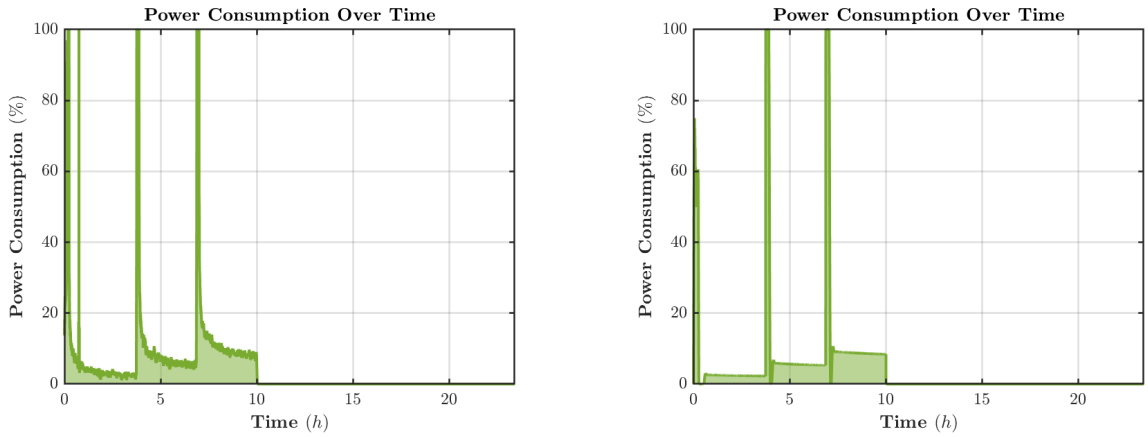

(a) Original "Step" power profile

(b) Tuned PID "Step" power profile

Figure 6.6: Measured vs. tuned PID power profile for the "Step" dataset.

large value. This outcome most likely compensates for the decision to fix the combined radiative surface area of the heaters ($A_h = 1\,\mathrm{m}^2$), a simplification that, in hindsight, may be physically inconsistent with the actual geometry of the furnace. The optimization of m_h therefore effectively absorbed the influence of this assumption, yielding a realistic overall thermal response despite the fixed-area approximation. A more detailed geometric characterization of the heating elements would remove this compensation effect in future model iterations.

Lastly, the number of heating elements, though initially assumed to be 18, was calibrated to 10. Given the strong influence this parameter has on shifting the overall temperature profile, treating it as a free parameter proved a prudent choice. The fact that the optimizer converged to an even value is encouraging given the symmetry of the furnace design. However, in the absence of direct temperature measurements of the heating elements, walls, or metal sheets, we cannot fully assess how faithfully the thermal model reproduces the real temperature distribution within the furnace.

Overall, the calibrated model reproduces measured furnace behavior with high fidelity across both thermal datasets. Significant improvements were gained by subsituting the raw power measurments with a tuned PID controller, highlighting the importance of accurate power measurement and the limitations of sweeping assumptions. Despite uncertainties in the physical power measurements, these results confirm that the model captures realistic furnace behavior and is sufficiently robust for subsequent process optimization and control design.

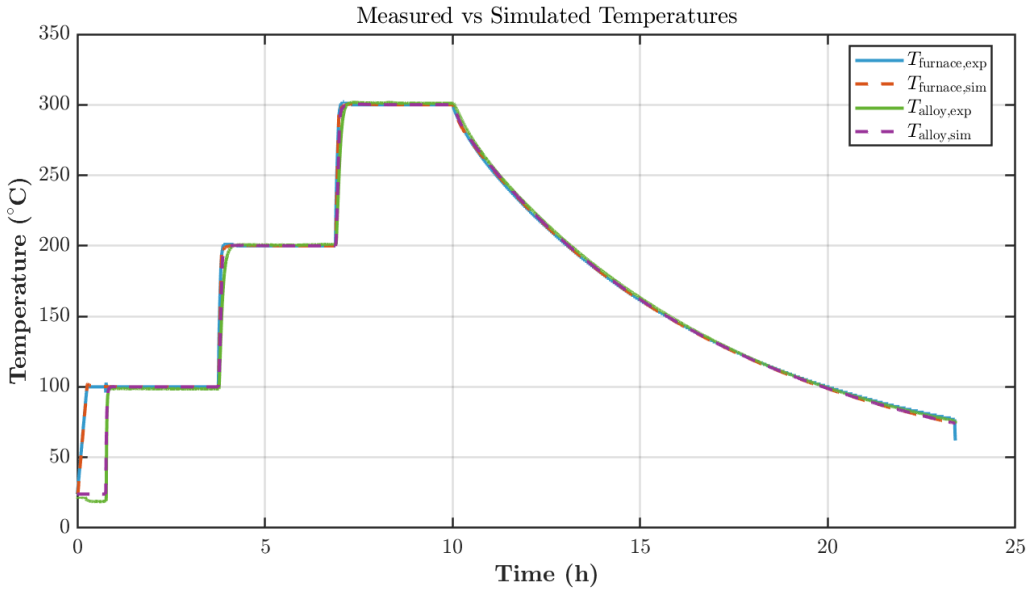

(a) “Step” dataset

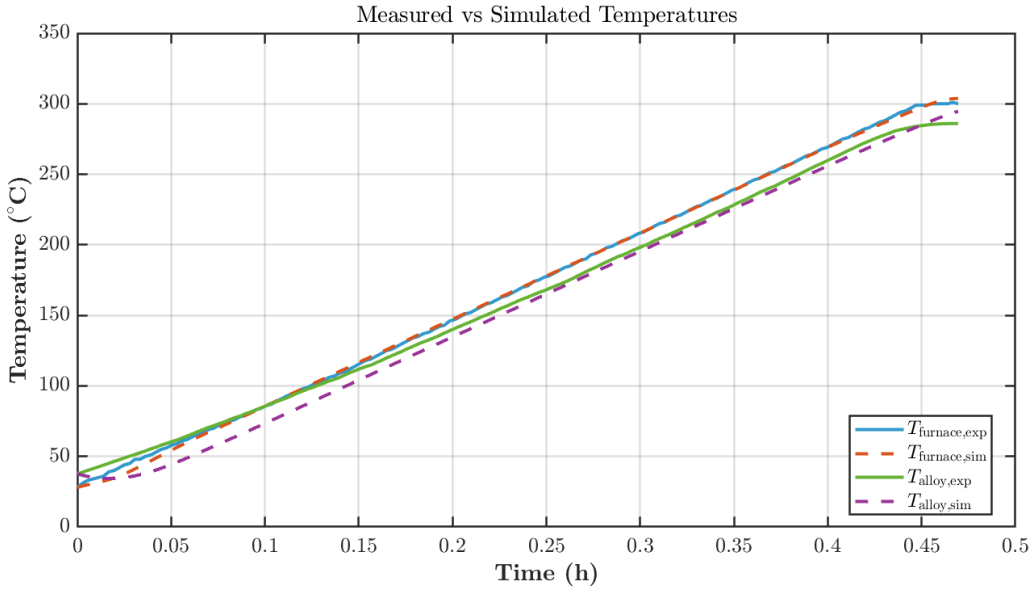

(b) “Ramp” dataset

Figure 6.7: Simulated vs. measured temperature profiles using PID-controlled power input.

Table 6.5: Optimized calibration parameters compared to initial estimates for both datasets

Parameter	Optimized Value	Initial Value	Units
Emissivity Parameters			
ϵ_w	0.006163	0.9	–
ϵ_{ms}	1	0.8	–
Convective Heat Transfer Coefficients			
$h_{f\to w}$	0.223288	10	$\mathrm{W\,m^{-2}\,K^{-1}}$
$h_{\text{ms}\to f}$	9.269076	20	$\mathrm{W\,m^{-2}\,K^{-1}}$
$h_{f\to \text{al}}$	37.901044	20	$\mathrm{W\,m^{-2}\,K^{-1}}$
$h_{h\to f}$	51.152402	25	$\mathrm{W\,m^{-2}\,K^{-1}}$
h_{out}	1.569847	10	$\mathrm{W\,m^{-2}\,K^{-1}}$
Heater-Related Parameters			
m_h	82.453652	10	kg
Material Properties			
$\lambda_{w,1}$	0.061096	0.05	$\mathrm{W\,m^{-1}\,K^{-1}}$
$\lambda_{w,2}$	55.509497	50	$\mathrm{W\,m^{-1}\,K^{-1}}$
Controller Parameters			
K_p	369.205	100	–
K_i	0.601	1	–
K_d	91.563	10	–
n_{heaters}	10	18	–

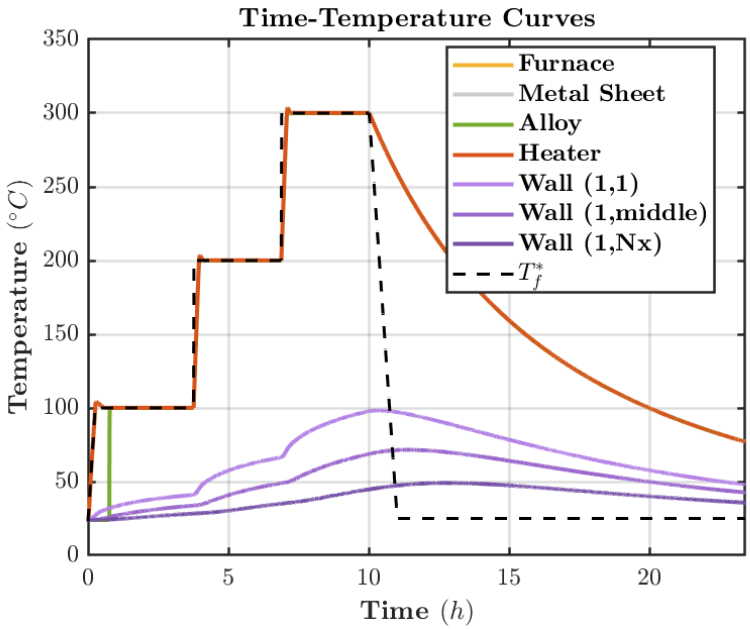

Figure 6.8: Thermal dynamics of the calibrated model for the "step" dataset using PID-controlled power input.

7 Optimization Framework

This chapter presents the complete optimization framework developed for the purpose of balancing competing objectives: minimizing energy consumption and processing time while ensuring the final mechanical properties meet stringent quality requirements. We begin by formulating the (non-convex) optimization problem in a mathematically rigorous manner, defining the control/design variables and the objective function, followed by the relevant constraints and bounds. These include both physical bounds and performance-based requirements derived from the desired mechanical properties of the treated alloy. To handle the inherent complexity and nonlinearity of the thermal process, surrogate models are used to efficiently approximate the relationship between thermal histories and resulting mechanical properties. We outline an optimization strategy, including MATLAB's `fmincon` and projected gradient descent (PGD) as solution algorithms.

7.1 Formulation of the Optimization Problem

The objective of the optimization is to determine a time-temperature curve $T_f(t)$ that minimizes the total energy consumption and process time while ensuring that the resulting mechanical properties of the aluminum alloy remain within acceptable ranges. The problem is formulated as a multi-objective optimization subject to physical, operational, and mechanical constraints.

7.1.1 Control Variables

The primary control (otherwise known as decision) variables are:

- $T_f(t)$: furnace temperature profile over the process duration $t \in [0, t_f]$,
- $\boldsymbol{\alpha}$: fractional vector of time allocations per segment,
- t_f: total process time.

7.1.2 Temperature Profile Parameterizations

While the optimization framework supports multiple types of temperature profiles, their suitability varies depending on practical constraints.

(1) N-Step Constant Temperature Profiles

The primary focus of this work is on piecewise-constant N-step profiles. This choice reflects the practical limitations of the industrial furnace controller, which is designed to accept a fixed sequence of target set-points. Each step corresponds to a constant target temperature maintained for a fixed fraction of the total process duration.

For an N-step profile, the control vector is:

$$\mathbf{x} = \left[T_{f,1}^{\text{set}}, \ldots, T_{f,N}^{\text{set}}, \alpha_1, \ldots, \alpha_N, t_f\right],$$

where:

- $\mathbf{T_f}^{\text{set}} \in \mathbb{R}^N$ are the target temperatures,
- $\boldsymbol{\alpha} \in \mathbb{R}^N$ are the fractional time allocation per segment ($\Sigma\, \alpha_i = 1$),
- t_f is the total process duration.

This results in $2N + 1$ decision variables.

(2) Linear Interpolation Profiles

Some industrial furnace controllers can also accommodate piecewise-linear ramps between temperature set-points rather than abrupt steps. In this parameterization, the control variables are still the set-points $\mathbf{T_f}^{\text{set}}$ and fractional time allocations $\boldsymbol{\alpha}$, but the resulting profile is piecewise linear and smooth. These profiles are more representative of continuous heating processes and can reduce thermal shock, but they require slightly more advanced control capabilities.

(3) Spline-Based Profiles

For theoretical completeness, we also allow the optimization framework to represent arbitrary smooth curves via cubic splines [13]. This provides maximum flexibility and may yield optimal solutions unattainable by stepwise or linear approximations. However, spline-based profiles are not currently reproducible on the industrial furnace we aim to test on. Thus, they are used primarily to explore theoretical limits rather than to propose immediately deployable solutions.

Although the optimization framework supports all three parameterizations, N-step profiles are emphasized throughout this thesis because they are the most practical to implement and test experimentally.

7.1.3 Objective/Loss Function

We define a composite objective function $\mathcal{L}$ that balances three competing goals:

1. **Energy consumption:** $\mathcal{P}(\mathbf{x})$,
2. **Process time:** $\mathcal{T}(\mathbf{x})$,
3. **Mechanical performance:** $\mathcal{M}(\mathbf{x})$,

$$\min_{\mathbf{x}} \quad \mathcal{L}(\mathbf{x}) = w_{\text{time}}\, \mathcal{T}(\mathbf{x}) + w_{\text{power}}\, \mathcal{P}(\mathbf{x}) + w_{\text{mech}}\, \mathcal{M}(\mathbf{x}), \tag{7.1}$$

where w_{time}, w_{power}, w_{mech} are scenario-dependent weighting factors reflecting the relative influence of each term on the total loss.

Time Loss

The total process duration is simply:

$$\mathcal{T}(\mathbf{x}) = \frac{1}{t_{\max}} t_f, \tag{7.2}$$

Here, $t_{\max}$ represents the longest permissible time duration, a normalization factor.

Power Loss

Energy consumption is penalized relative to the maximum possible power:

$$\mathcal{P}(\mathbf{x}) = \frac{1}{N_t \, P_{\max}} \Sigma_{k=1}^{N_t} P^{(k)}(T_f^{\text{set }(k)}, T_f^{(k)}), \tag{7.3}$$

where,

- $P_{\max}$: maximum power,
- N_t: total number of time-steps,
- $P^{(k)}$: power supplied at time step (k), computed via a PID controller.

Mechanical Loss

Mechanical performance is quantified through predicted yield strength (σ_y), ultimate tensile strength (σ_u), and uniform elongation (ε_u) at the final timestep of the procedure. Deviations from minimum acceptable thresholds are penalized quadratically, while exceeding the targets yields a tunable reward. This formulation encourages strong mechanical performance without imposing hard constraints:

$$\mathcal{M}(\mathbf{x}) = w_{\sigma_y} \cdot \ell(\delta_{\sigma_y}(\mathbf{x})) + w_{\sigma_u} \cdot \ell(\delta_{\sigma_u}(\mathbf{x})) + w_{\varepsilon_u} \cdot \ell(\delta_{\varepsilon_u}(\mathbf{x})), \tag{7.4}$$

where:

$$\ell(\delta) = \begin{cases} \delta^2, & \delta < 0 \quad \text{(penalty)}, \\ -\gamma \cdot \delta^2, & \delta \geq 0 \quad \text{(reward)}, \end{cases}$$

with normalized deviations from the minimum acceptable values:

$$\delta_{\sigma_y}(\mathbf{x}) = \frac{\sigma_y(\mathbf{x}) - \sigma_{y_{\min}}}{\sigma_{y_{\min}}}, \quad \delta_{\sigma_u}(\mathbf{x}) = \frac{\sigma_u(\mathbf{x}) - \sigma_{u_{\min}}}{\sigma_{u_{\min}}}, \quad \delta_{\varepsilon_u}(\mathbf{x}) = \frac{\varepsilon_u(\mathbf{x}) - \varepsilon_{u_{\min}}}{\varepsilon_{u_{\min}}},$$

where:

- $\sigma_{y_{\min}}$, $\sigma_{u_{\min}}$, $\varepsilon_{u_{\min}}$: minimum required values for the respective properties,
- w_{σ_y}, w_{σ_u}, w_{ε_u}: scenario-dependent weights reflecting relative importance,
- γ: reward gain factor $(0 < \gamma \leq 1)$.

These weights are normalized to create a more balanced and flexible evaluation of mechanical performance, offsetting the difference in scale of the different properties. Otherwise, the uniform elongation, which is numerically considerably smaller than the other two properties, contributes little to the loss, potentially neglecting its importance.

7.1.4 Scenario-Based Optimization

The choice of weights w_{time}, w_{power}, w_{mech}, as well as w_{σ_y}, w_{σ_u}, w_{ε_u}, fundamentally alters the optimization landscape. Instead of seeking a single universal optimum, we define a **scenario matrix** that captures different industrial priorities:

Scenario	w_{time}	w_{power}	w_{mech}	w_{σ_y}	w_{σ_u}	w_{ε_u}
S1	High	Low	Medium	2	1	0.5
S2	Low	High	Medium	2	1	0.5
S3	Medium	Low	High	2	1	0.5
S4	Medium	Low	High	100	0	0
S5	Medium	Low	High	0	100	0
S6	Medium	Low	High	0	0	100

Table 7.1: Scenario matrix defining different industrial optimization priorities.

This formulation and approach establishes a flexible optimization framework capable of handling multiple objectives and control strategies. The reader may wonder why the yield strength is assigned a weighting twice that of the ultimate tensile strength, which is again twice that of the uniform elongation. This is simply a reflection of the ranking of importance for these mechanical properties from an industrial perspective, which may be altered at will to suit the needs of the particular optimization setting. While our framework can in theory optimize arbitrarily smooth temperature curves, we focus primarily on N-step profiles for their ease of implementation and experimental validation. The use of a scenario matrix provides a set of optimal solutions tailored to different industrial priorities.

7.2 Constraints

The optimization problem is subject to the following constraints:

$$
\begin{aligned}
&T_{\min} \leq T_f^{\text{set}\,(k)} \leq T_{\max}, \quad \forall k, \\
&\dot{T}_{\min} \leq \frac{T_f^{\text{set}\,(k+1)} - T_f^{\text{set}\,(k)}}{\Delta t} \leq \dot{T}_{\max}, \quad \forall k, \\
&\sigma_y(\mathbf{x}) \geq \sigma_{y_{\min}}, \quad \sigma_u(\mathbf{x}) \geq \sigma_{u_{\min}}, \quad \varepsilon_u(\mathbf{x}) \geq \varepsilon_{u_{\min}}, \\
&0 < t_f \leq t_{\max}, \quad 0 \leq \alpha_i \leq 1, \quad \sum_{i=1}^{N} \alpha_i = 1.
\end{aligned}
$$

Here, $T_{\min}$ and $T_{\max}$ are the furnace temperature limits, while $\dot{T}_{\min}$ and $\dot{T}_{\max}$ impose maximum heating and cooling rates based on thermal dynamics as described below.

Maximum Heating and Cooling Rates

To ensure physically feasible time-temperature trajectories, bounds are imposed on the maximum heating and cooling rates of the furnace air. Based on simulation runs of the full furnace

model, conservative estimates were derived from the steepest average slopes of observed temperature curves:

$$\dot{T}_{\max} = R_{\text{heat,max}} = 0.1587 \text{ K/s},$$
$$\dot{T}_{\min} = R_{\text{cool,max}} = 0.0278 \text{ K/s}.$$

These values represent effective upper limits on the achievable rates within the operating range. Enforcing them prevents the optimizer from proposing trajectories the physical system cannot realize.

Limitations and Potential Improvements

Using fixed constants simplifies implementation but introduces several approximations. Actual heating and cooling rates vary with temperature, as radiative and convective heat transfer intensify at higher temperatures, and the achievable rates also depend on system-specific factors such as fan speed, alloy mass, and wall material properties. Furthermore, the estimates are based on average slopes rather than true instantaneous maxima. Future refinements could improve model fidelity by defining temperature-dependent rate functions, $R_{\text{heat}}(T)$ and $R_{\text{cool}}(T)$, employing lookup tables with interpolation to capture nonlinear furnace dynamics, or coupling achievable rates explicitly to heater power and boundary conditions. Despite these simplifications, the imposed bounds effectively limit physically unrealistic solutions.

7.3 Surrogate Model for Mechanical Properties

The furnace model described in this work outputs the time-temperature evolution of the aluminum alloy, $T_{\text{alloy}}(t)$, as a result of a candidate heating strategy. However, the optimization objective is not defined directly in terms of temperature but in terms of the final mechanical properties of the material: yield strength (σ_y), ultimate tensile strength (σ_u), and uniform elongation (ε_u). The mechanical model $\mathbf{M}$ and phase-fraction model $\mathbf{X}$ handles the transformation from temperature trajectiories into mechanical performance of the alloy element.

Overview of the Mechanical Model and Role in Optimization

AIT's in-house model predicts the mechanical properties of the alloy based on its thermal history. The workflow can be summarized as follows:

1. **Phase fraction computation.** Given the alloy temperature profile $T_{\text{alloy}}(t)$, the model computes the evolution of microstructural phase fractions ($\mathbf{X}(t)$) over time.
2. **Property prediction.** The final mechanical properties (σ_y, σ_u, ε_u) are then estimated from these phase fractions using empirical relations derived from experimental data.
3. **Penalize/reward mechanical performance.** Deviations from the desired minimum property thresholds. ($\sigma_{y,\min}$, $\sigma_{u,\min}$, $\varepsilon_{u,\min}$) are incorporated into the objective function as weighted penalty terms.

Thus, the mapping between a candidate temperature trajectory and the resulting mechanical performance is established through (See Figure 1.2):

$$T_{\text{alloy}}(t) \quad \longrightarrow \quad \mathbf{X}(t) \quad \longrightarrow \quad \{\sigma_y(t),\, \sigma_u(t),\, \varepsilon_u(t)\}.$$

Limitations and Considerations

While the surrogate mechanical model enables rapid evaluation of alloy properties, it has several inherent limitations. The model considers only the relative phase fractions and does not account for precipitate radii or size distributions, which can significantly influence mechanical behavior. Moreover, it relies on the full alloy temperature history as input, meaning that any error in the predicted $T_{\text{alloy}}(t)$ propagates directly into the computed phase fractions and resulting mechanical properties. Finally, the model has been developed and calibrated for a specific Al-Mg-Si alloy with a defined phase sequence, and its predictions cannot be directly generalized to other alloys with different compositions or precipitation behaviors.

7.4 Solvers and Strategy

This section describes the strategies and algorithms used to solve the optimization problem for determining the optimal furnace time-temperature profile. The aim is to find a control trajectory that minimizes energy consumption and total process time while ensuring that the resulting mechanical properties go above and beyond satisfaction of the desired mechanical performance thresholds.

7.4.1 Adaptive Time-Parameterized Temperature Control

To flexibly represent the furnace's desired temperature trajectory $T_f^{\text{set}}(t)$, we adopt a parameterization based on N control segments. This approach allows the optimizer to adjust both the **shape** of the time-temperature curve and the **total process duration** simultaneously. This avoids the rigidity of fixed time partitions and allows the optimizer to adjust both timing and temperature levels dynamically.

Parameterization

We define:

- N — number of control segments,
- $T_{f,i}^{\text{set}}$ — furnace set-point temperature during segment i,
- α_i — normalized time fraction for segment i, where $\alpha_i \geq 0$ and $\sum_{i=1}^{N} \alpha_i = 1$,
- t_f — total process duration.

The duration of each segment is:

$$\Delta t_i = \alpha_i \, t_f.$$

The cumulative time points are:

$$t_1 = 0, \qquad t_i = \sum_{j=1}^{i-1} \Delta t_j = \left(\sum_{j=1}^{i-1} \alpha_j \right) t_f.$$

Control profiles

We distinguish between three main temperature control profiles.

1. **Constant N-step profiles:**

$$T_f^{\text{set}}(t) = T_{f,i}^{\text{set}}, \quad \text{for } t_i \leq t < t_{i+1}, \quad i = 1, \ldots, N.$$

2. **Piecewise-linear interpolation:**

$$T_f^{\text{set}}(t) = T_{f,i}^{\text{set}} + \frac{T_{f,i+1}^{\text{set}} - T_{f,i}^{\text{set}}}{t_{i+1} - t_i}(t - t_i), \quad t_i \leq t < t_{i+1}, \quad i = 1, \ldots, N.$$

3. **Spline interpolation:** Smooth trajectories can be obtained via cubic splines through the $(t_i, T_{f,i}^{\text{set}})$ points, enabling more gradual temperature transitions and smoother derivatives. However, since the N 120/85 HA model furnace cannot handle higher order temperature set-point trajectories, these results cannot be validated and are therefore of less practical interest in this work.

7.4.2 Solver Strategy

The optimization landscape is likely nonconvex and nonlinear. See Appendix 9.4 for visualizations of the loss landscape under optimization scenario S3. We leverage two solver strategies in tandem to enhance reliability of the results.

(1) Black-Box Optimization with `fmincon`

MATLAB's `fmincon` from the Optimization Toolbox enables constrained nonlinear optimization using Sequential Quadratic Programming (SQP) or interior-point algorithms, supporting gradient-based solutions with optional analytical gradients [41]. Analytical gradients should improve convergence speed and solution accuracy, while finite-difference approximations are slower but still feasible; both approaches were tested for performance. Although `fmincon` is generally reliable in engineering applications, local minima remain a potential issue, which we mitigate by employing multiple random initializations and global search strategies such as MultiStart and GlobalSearch from MATLAB's Global Optimization Toolbox.

(2) Projected Gradient Descent (PGD)

To benchmark against black-box solvers, we implement a custom projected gradient descent (PGD) method. PGD iteratively updates the decision vector:

$$\boldsymbol{x}^{(k+1)} = \Pi_{\mathcal{C}}\Big(\boldsymbol{x}^{(k)} - \eta^{(k)} \nabla J(\boldsymbol{x}^{(k)})\Big),$$

where:

- $J(\boldsymbol{x})$ is the objective function,
- $\eta^{(k)}$ is the step size at iteration k,

- $\Pi_{\mathcal{C}}(\cdot)$ projects onto the feasible set $\mathcal{C}$.

We derive analytical expressions for ∇J in Section 7.5 to improve convergence speed and avoid numerical noise from finite differences. Modern frameworks such as PyTorch or TensorFlow offer automatic differentiation in other programming languages, and therefore we initially attempted to implement automatic differentiation in MATLAB using ADiMat and the Deep Learning Toolbox [5]. In practice, these approaches proved unreliable due to a combination of server instability (in the case of ADiMat), limited support for certain operations such as linear system solves, and incompatibilities with the highly customizable code pipeline, which relies on conditional `try/catch` statements and dynamic branching. Consequently, manually deriving the gradients not only ensured robust computation but also provided valuable insight into the sensitivity of the model to each parameter.

PGD is well-studied for nonconvex optimization and enjoys theoretical guarantees in several domains. For instance, in nonconvex low-rank tensor regression, PGD converges effectively under structured assumptions [10], particularly for suitably chosen step-sizes. In nonconvex matrix completion, PGD achieves linear convergence to stationary points in many practical settings [54].

Step-size selection significantly impacts convergence. In this work, we employ the adaptive step-size rule proposed by Malitsky and Mishchenko [26],

$$\eta^{(k)} = \min\left\{ \sqrt{1+\theta^{(k-1)}}\,\eta^{(k-1)},\ \frac{\|x^{(k)}-x^{(k-1)}\|}{2\,\|\nabla J(x^{(k)})-\nabla J(x^{(k-1)})\|} \right\}, \quad \text{with} \quad \theta^{(k)} = \frac{\eta^{(k)}}{\eta^{(k)}}, \tag{7.5}$$

This adaptive rule has been shown to perform robustly across a wide range of optimization problems without the need for extensive hyperparameter tuning. It automatically balances stability and progress by adjusting the step size based on recent iterate and gradient changes, often achieving fast convergence rates while remaining computationally inexpensive. For these reasons, only the adaptive step-size scheme was used in this study.

7.4.3 Projecting onto the feasible set

In order to ensure a consistent comparison between the proposed projected gradient descent (PGD) optimizer and the benchmark `fmincon` solver, the PGD framework was extended with a *proximal projection operator* that enforces all physical and mechanical feasibility constraints at each iteration. In theory, the projected gradient descent (PGD) algorithm requires that each tentative update

$$\mathbf{y}^{(k+1)} = \mathbf{x}^{(k)} - \eta^{(k)}\nabla J(\mathbf{x}^{(k)}) \tag{7.6}$$

be projected onto the feasible set $\mathcal{C}$, yielding

$$\mathbf{x}^{(k+1)} = \Pi_{\mathcal{C}}\left(\mathbf{y}^{(k+1)}\right) = \arg\min_{\mathbf{z}\in\mathcal{C}} \frac{1}{2}\|\mathbf{z}-\mathbf{y}^{(k+1)}\|_2^2. \tag{7.7}$$

where $\Pi_{\mathcal{C}}$ denotes the Euclidean projection operator. If all constraints were convex and easily expressed, Eq. (7.7) could be solved directly as a small quadratic program. However, in this work the feasible set combines several qualitatively different types of constraints such as simple convex bounds (temperature and time limits), affine equality constraints (time-fraction normalization),

nonlinear sequential constraints (heating and cooling rate limits), and nonconvex constraints that depend on the thermal history (mechanical properties). Solving a full projection over this composite set would therefore require a large nonlinear optimization problem at every iteration, defeating the efficiency advantage of PGD.

To make the projection step tractable, the feasible set is decomposed into simpler subsets,

$$\mathcal{C} \approx \mathcal{C}_t \cap \mathcal{C}_\alpha \cap \mathcal{C}_{T,\text{rate}} \cap \mathcal{C}_{\text{mech}},$$

and the projection operator is implemented as a sequence of partial projections:

$$\Pi_{\mathcal{C}} \approx \Pi_{\mathcal{C}_{\text{mech}}} \circ \Pi_{\mathcal{C}_{T,\text{rate}}} \circ \Pi_{\mathcal{C}_\alpha} \circ \Pi_{\mathcal{C}_t}.$$

Each sub-projection enforces a specific subset of constraints, with simple sets handled exactly and complex ones handled approximately. This approach corresponds to a *composite projection operator*, similar in spirit to alternating or block-coordinate projection methods such as Dykstra's algorithm [6]. While not an exact Euclidean projection, it ensures that each iterate remains feasible with respect to all physical and mechanical limits, while keeping the computational cost small compared to the thermal simulation.

Physical and process constraints

Projection for the physical parameters $\{T_f^{\text{set}}, \boldsymbol{\alpha}, t\}$ is handled through simple elementwise or sequential operations:

1. **Total time constraint:** The process time is clamped to its limits:

$$t_{\text{total}} \leftarrow \min(\max(t_{\text{total}}, t_{\min}), t_{\max}).$$

2. **Time fraction constraint:** The vector $\boldsymbol{\alpha}$ must remain nonnegative and sum to one. It is therefore projected onto the probability simplex using MATLAB's `projsplx` [11].

3. **Temperature and rate constraints:** The temperature profile T_f^{set} must satisfy both absolute temperature limits and heating/cooling rate limits:

$$T_{\min} \leq T_f^{\text{set}\,(k)} \leq T_{\max}, \quad \forall k,$$

and

$$\dot{T}_{\min} \leq \frac{T_f^{\text{set}\,(k+1)} - T_f^{\text{set}\,(k)}}{\Delta t} \leq \dot{T}_{\max}, \quad \forall k,$$

These constraints are enforced by a heuristic "slope-limiting" procedure that adjusts each temperature segment to the nearest value consistent with both bounds and allowable temperature changes. This provides a fast and numerically stable projection in most cases. If this heuristic correction fails (for example, due to accumulated inconsistencies in a highly nonlinear region), a small quadratic program (QP) is solved to restore feasibility more precisely.

Mechanical feasibility

Mechanical properties such as yield stress, ultimate strength, and elongation depend nonlinearly on the entire temperature history, making them more challenging to constrain directly. To make their enforcement computationally tractable, these constraints are linearized locally around the current iterate $\mathbf{x}^{(k)}$ using first-order Taylor expansions:

$$M_i(\mathbf{x}) \approx M_i(\mathbf{x}^{(k)}) + \nabla M_i(\mathbf{x}^{(k)})^\top (\mathbf{x} - \mathbf{x}^{(k)}), \qquad i \in \{\sigma_y, \sigma_u, \varepsilon_u\}. \tag{7.8}$$

Substituting these linearized forms into Eq. (7.7) yields a small quadratic program (QP) that determines the correction step $\boldsymbol{\delta}^{(k)} = \mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}$:

$$\begin{aligned} \min_{\boldsymbol{\delta}^{(k)}} \quad & \frac{1}{2}\|\boldsymbol{\delta}^{(k)}\|_2^2, \\ \text{s.t.} \quad & \nabla M_i\left(\mathbf{x}^{(k)}\right)^\top \boldsymbol{\delta}^{(k)} \geq M_{i,\min} - M_i\left(\mathbf{x}^{(k)}\right), \qquad i \in \{\sigma_y, \sigma_u, \varepsilon_u\}. \end{aligned} \tag{7.9}$$

This convex QP represents the minimal modification of the current iterate that restores feasibility with respect to the locally linearized mechanical constraints. It is solved efficiently using MATLAB's `quadprog` routine, introducing negligible overhead compared to the thermal simulation. If the QP fails to find a feasible step, for example due to poor linearization in a highly nonlinear region, a fallback repair procedure based on `fmincon` is used to recover feasibility.

By incorporating these proximal projection steps, the PGD optimizer enforces exactly the same feasibility criteria as `fmincon`. Consequently, any differences in the resulting objective value or convergence rate between the two solvers can be attributed purely to their algorithmic properties rather than inconsistencies in constraint enforcement. This dual-solver strategy allows cross-validation of results: `fmincon` provides a robust baseline, while PGD offers a transparent and customizable alternative to compare results with.

7.5 Gradient of the Loss Function with Respect to Control Variables

This section derives the analytical gradients of the overall loss function with respect to the control variables. These gradients are used for the projected gradient descent (PGD) solver and provide insights into the sensitivities of the system with respect to temperature set-points and segment durations.

Problem Setup

We parameterize the furnace temperature profile using a set of N control segments, each with a constant set-point temperature and an associated segment duration. The decision vector $\mathbf{x} \in \mathbb{R}^{2N+1}$ is defined as:

$$\boldsymbol{x} = \left[\, T_{f,1}^{\text{set}},\ \ldots,\ T_{f,N}^{\text{set}},\ \alpha_1,\ \ldots,\ \alpha_N,\ t_f \,\right].$$

The scalar loss function is defined as:

$$\mathcal{L}(\mathbf{x}) = w_{\text{time}}\, \mathcal{T}(\mathbf{x}) \ + \ w_{\text{power}}\, \mathcal{P}(\mathbf{x}) \ + \ w_{\text{mech}}\, \mathcal{M}(\mathbf{x}), \tag{7.10}$$

where w_{time}, w_{power}, and w_{mech} are scenario-dependent weights balancing time, energy, and mechanical-property objectives.

Analytical gradients $\nabla_{\mathbf{x}}\mathcal{L}$ are derived for all three terms by exploiting the differentiability of the mechanical model M. This approach avoids the numerical noise associated with finite-difference approximations due to temperature step discontinuities, enabling significantly faster optimization, and provides direct sensitivity information for interpreting optimizer behavior. The detailed derivations are presented in the following subsections.

Gradient Decomposition

The gradient of the loss function with respect to the control variables is a vector of size $[1 \times (2N+1)]$ and is given by:

$$\frac{\mathrm{d}\mathcal{L}(\mathbf{x})}{\mathrm{d}\mathbf{x}} = w_{\text{time}} \frac{\mathrm{d}\mathcal{T}(\mathbf{x})}{\mathrm{d}\mathbf{x}} + w_{\text{power}} \frac{\mathrm{d}\mathcal{P}(\mathbf{x})}{\mathrm{d}\mathbf{x}} + w_{\text{mech}} \frac{\mathrm{d}\mathcal{M}(\mathbf{x})}{\mathrm{d}\mathbf{x}}. \tag{7.11}$$

Derivative of the time term.

The derivative of the normalized process time $\mathcal{T}(\mathbf{x})$ is straightforward. Using e_i to denote the i-th standard basis vector in $\mathbb{R}^{2N+1}$, we have:

$$\frac{\mathrm{d}\mathcal{T}(\mathbf{x})}{\mathrm{d}\mathbf{x}} = \frac{1}{t_{\max}} \frac{\mathrm{d}t_f}{\mathrm{d}\mathbf{x}} = \frac{1}{t_{\max}} e_{2N+1} = \frac{1}{t_{\max}} \begin{bmatrix} 0 & \dots & 0 & 1 \end{bmatrix}. \tag{7.12}$$

Derivative of the power term.

The power-related contribution is more involved. From the definition of $\mathcal{P}(\mathbf{x})$, we have:

$$\frac{\mathrm{d}\mathcal{P}(\mathbf{x})}{\mathrm{d}\mathbf{x}} = \frac{1}{N_t\, P_{\max}} \sum_{k=1}^{N_t} \frac{\mathrm{d}P^{(k)}}{\mathrm{d}\mathbf{x}}. \tag{7.13}$$

Derivative of the instantaneous power. The power at time step k is given by:

$$P^{(k)} = \frac{1}{n_{\text{heaters}}} \min\Big(\max\Big(0,\, P_{\text{unsat}}^{(k)}\Big),\, P_{\max}\Big), \tag{7.14}$$

where $P_{\text{unsat}}^{(k)}$ is the unsaturated PID control output, defined in Eq. (3.7):

$$\begin{aligned} P_{\text{unsat}}^{(k)} = {} & K_p \left(T_f^{\text{set}\,(k)} - T_f^{(k)}\right) + K_i \Delta t \sum_{j=1}^{k-1} \left(T_f^{\text{set}\,(j)} - T_f^{(j)}\right) \\ & + \frac{K_d}{\Delta t} \Big[\left(T_f^{\text{set}\,(k)} - T_f^{(k)}\right) - \left(T_f^{\text{set}\,(k-1)} - T_f^{(k-1)}\right) \Big]. \end{aligned} \tag{7.15}$$

Differentiating $P^{(k)}$ with respect to $\mathbf{x}$ yields:

$$\frac{\mathrm{d}P^{(k)}}{\mathrm{d}\mathbf{x}} = \frac{1}{n_{\text{heaters}}} \frac{\mathrm{d}}{\mathrm{d}\mathbf{x}} \left[\min\Big(\max\Big(0,\, P_{\text{unsat}}^{(k)}\Big), P_{\max}\Big)\right]. \tag{7.16}$$

Because of the saturation limits, this derivative is piecewise-defined:

$$\frac{\mathrm{d}P^{(k)}}{\mathrm{d}\mathbf{x}} = \begin{cases} 0, & \text{if } P_{\text{unsat}}^{(k)} \leq 0 \text{ or } P_{\text{unsat}}^{(k)} \geq P_{\max}, \\ \dfrac{1}{n_{\text{heaters}}} \dfrac{\mathrm{d}P_{\text{unsat}}^{(k)}}{\mathrm{d}\mathbf{x}}, & \text{if } 0 < P_{\text{unsat}}^{(k)} < P_{\max}. \end{cases} \tag{7.17}$$

Derivative of the unsaturated power. Within the active region ($0 < P_{\text{unsat}}^{(k)} < P_{\max}$), we have:

$$\begin{aligned} \frac{\mathrm{d}P_{\text{unsat}}^{(k)}}{\mathrm{d}\mathbf{x}} = K_p \left(\frac{\mathrm{d}T_f^{\text{set}\,(k)}}{\mathrm{d}\mathbf{x}} - \frac{\mathrm{d}T_f^{(k)}}{\mathrm{d}\mathbf{x}} \right) + K_i \Delta t \sum_{j=1}^{k-1} \left(\frac{\mathrm{d}T_f^{\text{set}\,(j)}}{\mathrm{d}\mathbf{x}} - \frac{\mathrm{d}T_f^{(j)}}{\mathrm{d}\mathbf{x}} \right) \\ + \frac{K_d}{\Delta t} \left[\left(\frac{\mathrm{d}T_f^{\text{set}\,(k)}}{\mathrm{d}\mathbf{x}} - \frac{\mathrm{d}T_f^{(k)}}{\mathrm{d}\mathbf{x}} \right) - \left(\frac{\mathrm{d}T_f^{\text{set}\,(k-1)}}{\mathrm{d}\mathbf{x}} - \frac{\mathrm{d}T_f^{(k-1)}}{\mathrm{d}\mathbf{x}} \right) \right]. \end{aligned} \tag{7.18}$$

This expression accounts for the dependency of both the desired and actual furnace temperatures on the control vector $\mathbf{x}$. Substituting Eq. (7.18) into Eq. (7.17) yields a full analytical formula for $\frac{\mathrm{d}P^{(k)}}{\mathrm{d}\mathbf{x}}$, and thus for $\frac{\mathrm{d}\mathcal{P}(\mathbf{x})}{\mathrm{d}\mathbf{x}}$.

Chain rule for furnace temperature sensitivity. The derivative of the simulated furnace temperature $T_f^{(k)}$ with respect to the control variables follows from the implicit update scheme:

$$\frac{\mathrm{d}T_f^{(k)}}{\mathrm{d}\mathbf{x}} = \frac{\mathrm{d}T_f^{(k)}}{\mathrm{d}\mathbf{T}^{(k)}} \cdot \frac{\mathrm{d}\mathbf{T}^{(k)}}{\mathrm{d}\mathbf{x}}, \tag{7.19}$$

where $\mathbf{T}^{(k)} \in \mathbb{R}^4$ denotes the coupled temperature states of the system:

$$\mathbf{T}^{(k)} = \left[T_f^{(k)}, T_{\text{alloy}}^{(k)}, T_{\text{ms}}^{(k)}, T_h^{(k)} \right]^\top.$$

These sensitivities are obtained directly from the Jacobian of the implicit Euler update scheme, discussed in Section 5.2.3. The sensitivity of the furnace state with respect to the full temperature state vector is trivial by definition: $\frac{\mathrm{d}T_f^{(k)}}{\mathrm{d}\mathbf{T}^{(k)}} = [1, 0, 0, 0]^T$.

Set-point temperature derivative with respect to control variable. The derivative of the furnace set-point temperature depends on the chosen parametrization of the time-temperature curve.

1. **Constant N-step profile:** For a piecewise-constant profile over N steps, the derivative with respect to the control variable is simply the corresponding basis vector:

$$\frac{\mathrm{d}T_f^{\text{set}\,(k)}}{\mathrm{d}\mathbf{x}} = e_k,$$

where $e_k \in \mathbb{R}^{2N+1}$ is the standard basis vector selecting the k-th temperature step in the control vector $\mathbf{x}$.

2. **Linear interpolation profile:** For a linearly interpolated profile between two consecutive knots i and $i+1$, the set-point is

$$T_f^{\text{set}}(t) = T_{f,i}^{\text{set}} + \frac{T_{f,i+1}^{\text{set}} - T_{f,i}^{\text{set}}}{t_{i+1} - t_i}(t - t_i), \quad t_i \le t < t_{i+1},$$

where the time knots t_i themselves may depend on the normalized coefficients α_j and the total process time t_f.

- **Derivative with respect to temperatures $T_{f,i}^{\text{set}}$ and $T_{f,i+1}^{\text{set}}$:**

$$\frac{\mathrm{d}T_f^{\text{set}}(t)}{\mathrm{d}T_{f,i}^{\text{set}}} = 1 - \frac{t - t_i}{t_{i+1} - t_i}, \quad \frac{\mathrm{d}T_f^{\text{set}}(t)}{\mathrm{d}T_{f,i+1}^{\text{set}}} = \frac{t - t_i}{t_{i+1} - t_i}.$$

- **Derivative with respect to time coefficients α_j:**

$$\frac{\mathrm{d}T_f^{\text{set}}(t)}{\mathrm{d}\alpha_j} = \frac{\partial T_f^{\text{set}}}{\partial t_i}\frac{\partial t_i}{\partial \alpha_j} + \frac{\partial T_f^{\text{set}}}{\partial t_{i+1}}\frac{\partial t_{i+1}}{\partial \alpha_j},$$

with

$$\begin{cases} \frac{\partial t_i}{\partial \alpha_j} = t_f & \text{if } j < i, \\ \frac{\partial t_i}{\partial \alpha_j} = 0 & \text{otherwise.} \end{cases}$$

The partial derivatives with respect to the knots are

$$\frac{\partial T_f^{\text{set}}}{\partial t_i} = \frac{(T_{f,i} - T_{f,i+1})(t_{i+1} - t)}{(t_{i+1} - t_i)^2}, \quad \frac{\partial T_f^{\text{set}}}{\partial t_{i+1}} = \frac{(T_{f,i} - T_{f,i+1})(t - t_i)}{(t_{i+1} - t_i)^2}.$$

- **Derivative with respect to total process time t_f:**

$$\frac{\mathrm{d}T_f^{\text{set}}(t)}{\mathrm{d}t_f} = \frac{\partial T_f^{\text{set}}}{\partial t_i}\sum_{j=1}^{i-1}\alpha_j + \frac{\partial T_f^{\text{set}}}{\partial t_{i+1}}\sum_{j=1}^{i}\alpha_j.$$

These derivatives are combined into a single row vector for each time step t, yielding the full matrix $\mathrm{d}T_f^{\text{set}}/\mathrm{d}\mathbf{x}$ required for sensitivity analysis.

3. **Spline interpolation:** we rely on finite-difference approximation for simplicity here.

Residual-Based Sensitivity Formulation

Computing the derivative of the full temperature state $\mathbf{T}^{(k)}$ with respect to $\mathbf{x}$, warrants careful treatment. Recalling the nonlinear residual system from Eq. (5.29), the temperature states evolve according to

$$\mathbf{F}^{(k+1)}(\mathbf{T}^{(k+1)}, \mathbf{T}^{(k)}, \mathbf{T_{w,1}}^{(k+1)}, \mathbf{T_{w,1}}^{(k)}, P^{(k)}) = 0.$$

Differentiating the residual with respect to $\mathbf{x}$ yields:

$$\begin{aligned} 0 = &\frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T}^{(k+1)}}\frac{\mathrm{d}\mathbf{T}^{(k+1)}}{\mathrm{d}\mathbf{x}} + \frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T}^{(k)}}\frac{\mathrm{d}\mathbf{T}^{(k)}}{\mathrm{d}\mathbf{x}} + \frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T_{w,1}}^{(k+1)}}\frac{\mathrm{d}\mathbf{T_{w,1}}^{(k+1)}}{\mathrm{d}\mathbf{x}} \\ &+ \frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T_{w,1}}^{(k)}}\frac{\mathrm{d}\mathbf{T_{w,1}}^{(k)}}{\mathrm{d}\mathbf{x}} + \frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}P^{(k)}}\frac{\mathrm{d}P^{(k)}}{\mathrm{d}\mathbf{x}}. \end{aligned} \tag{7.20}$$

Jacobian Matrices. The term $\frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T}^{(k+1)}}$ is exactly the Jacobian already computed in Eq. (5.34), and $\frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T}^{(k)}} = -I_4$, where I_4 is the 4×4 identity. Rewriting Eq. (7.20) gives the recursive formula:

$$\frac{\mathrm{d}\mathbf{T}^{(k+1)}}{\mathrm{d}\mathbf{x}} = -\left(\frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T}^{(k+1)}}\right)^{-1}\left[\frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T}^{(k)}}\frac{\mathrm{d}\mathbf{T}^{(k)}}{\mathrm{d}\mathbf{x}} + \frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T}_{\mathbf{w},\mathbf{1}}{}^{(k+1)}}\frac{\mathrm{d}\mathbf{T}_{\mathbf{w},\mathbf{1}}{}^{(k+1)}}{\mathrm{d}\mathbf{x}} + \frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T}_{\mathbf{w},\mathbf{1}}{}^{(k)}}\frac{\mathrm{d}\mathbf{T}_{\mathbf{w},\mathbf{1}}{}^{(k)}}{\mathrm{d}\mathbf{x}} + \frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}P^{(k)}}\frac{\mathrm{d}P^{(k)}}{\mathrm{d}\mathbf{x}}\right]. \tag{7.21}$$

Starting from $\frac{\mathrm{d}\mathbf{T}^{(1)}}{\mathrm{d}\mathbf{x}} = 0$, this recurrence allows computing sensitivities for all time steps $k = 1, \ldots, N_t - 1$.

Power-dependent term. The contribution from the power input arises through the causal chain:

$$\mathbf{x} \to T_f^{\mathrm{set}\,(k)} \to P^{(k)} \to \mathbf{T}^{(k+1)}, \mathbf{T}_{\mathbf{w}}{}^{(k+1)} \to P^{(k+1)} \to \ldots$$

From the residual definition, we obtain:

$$\frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}P^{(k)}} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ -\frac{\Delta t}{m_h C_{p,h}} \end{bmatrix}. \tag{7.22}$$

Derivative with respect to wall temperatures. For the inner wall nodes, the residual sensitivities with respect to the wall temperatures can be expressed compactly as:

$$\frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}\mathbf{T}w,1^{(k+1)}} = \begin{bmatrix} \frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}Tw_1,1^{(k+1)}} & \cdots & \frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}T_{w_6,1}^{(k+1)}} \end{bmatrix}, \tag{7.23}$$

where each column corresponds to one of the six walls:

$$\frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}T_{w_i,1}^{(k+1)}} = \begin{bmatrix} -\frac{\Delta t}{m_f C_{p,f}} h_{f\to w_i} A_{w_i} \\ 0 \\ 0 \\ -\frac{\Delta t}{m_h C_{p,h}} 4\,\sigma\epsilon_{w_i} V F_{h\to w_i} A_{w_i} \left(T_{w_i,1}^{(k+1)}\right)^3 \end{bmatrix}. \tag{7.24}$$

Similary, the partial derivative with respect to the previous wall temperature is described column-wise by:

$$\frac{\mathrm{d}\mathbf{F}^{(k+1)}}{\mathrm{d}T_{w_i,1}^{(k)}} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ \frac{\Delta t}{m_h C_{p,h}} 12\,\sigma\epsilon_{w_i} V F_{h\to w_i} A_{w_i} \left(T_{w_i,1}^{(k)}\right)^3 \end{bmatrix}. \tag{7.25}$$

Wall Residual Formulation

Recall from Eq. (5.23) that the discretized heat diffusion equation for the furnace walls can be written as a linear system:

$$A\mathbf{T_w}^{(k+1)} = \mathbf{b}.$$

Differentiating both sides with respect to the control variable $\mathbf{x}$ and applying the product rule gives:

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}\mathbf{x}}\left(A\mathbf{T_w}^{(k+1)}\right) &= \frac{\mathrm{d}\mathbf{b}}{\mathrm{d}\mathbf{x}} \\ \implies A\frac{\mathrm{d}\mathbf{T_w}^{(k+1)}}{\mathrm{d}\mathbf{x}} + \frac{\mathrm{d}A}{\mathrm{d}\mathbf{x}}\mathbf{T_w}^{(k+1)} &= \frac{\mathrm{d}\mathbf{b}}{\mathrm{d}\mathbf{x}}. \end{aligned}$$

Rewriting, the derivative of the wall temperatures with respect to the control variables is

$$\frac{\mathrm{d}\mathbf{T_w}^{(k+1)}}{\mathrm{d}\mathbf{x}} = -A^{-1}\frac{\mathrm{d}A}{\mathrm{d}\mathbf{x}}\mathbf{T_w}^{(k+1)} + A^{-1}\frac{\mathrm{d}\mathbf{b}}{\mathrm{d}\mathbf{x}}. \tag{7.26}$$

Once $\frac{\mathrm{d}\mathbf{T_w}^{(k+1)}}{\mathrm{d}\mathbf{x}}$ is computed, the derivative of the first wall nodes, $\frac{\mathrm{d}\mathbf{T_{w,1}}^{(k+1)}}{\mathrm{d}\mathbf{x}}$, can be extracted by selecting the appropriate entries of this tensor of size $[N_{\text{walls}} \times N_x \times (2N+1)]$, yielding a $[N_{\text{walls}} \times (2N+1)]$ matrix.

Linear System Coefficients. The matrix A and vector $\mathbf{b}$ coefficients are known and depend on the thermal properties and boundary conditions.

For interior nodes $i = 2, \ldots, N_x - 1$:

$$\begin{aligned} \alpha_i &= \frac{(\rho_w C_{p,w} V_w)_i}{\Delta t} + \frac{(\lambda_w A_w)_{i+\frac{1}{2}}}{\Delta x_{i+1,i}} + \frac{(\lambda_w A_w)_{i-\frac{1}{2}}}{\Delta x_{i,i-1}}, \\ \beta_i &= -\frac{(\lambda_w A_w)_{i+\frac{1}{2}}}{\Delta x_{i+1,i}}, \\ \gamma_i &= -\frac{(\lambda_w A_w)_{i-\frac{1}{2}}}{\Delta x_{i,i-1}}, \\ b_i &= \frac{(\rho_w C_{p,w} V_w)_i}{\Delta t}\mathbf{T_{w,i}}^{(k)}. \end{aligned}$$

For the inner wall node $i = 1$:

$$\begin{aligned} \alpha_1 &= \frac{(\rho_w C_{p,w} V_w)_1}{\Delta t} + \frac{(\lambda_w A_w)_{1+\frac{1}{2}}}{\Delta x_{2,1}} + h_{f\to w}A_w + 4\,h_{\text{rad}}\,A_w\,(\mathbf{T_{w,1}}^{(k)})^3, \\ \beta_1 &= -\frac{(\lambda_w A_w)_{1+\frac{1}{2}}}{\Delta x_{2,1}}, \\ b_1 &= h_{f\to w}A_w T_f^{(k+1)} + \frac{(\rho_w C_{p,w} V_w)_1}{\Delta t}\mathbf{T_{w,1}}^{(k)} + h_{\text{rad}}A_w\left[(T_h^{(k+1)})^4 + 3(\mathbf{T_{w,1}}^{(k)})^4\right]. \end{aligned}$$

For the outer wall node $i = N_x$:

$$\alpha_{N_x} = \frac{(\rho_w C_{p,w} V_w)_{N_x}}{\Delta t} + \frac{(\lambda_w A_w)_{N_x - \frac{1}{2}}}{\Delta x_{N_x, N_x - 1}} + h_{w \to \infty} A_w,$$
$$\gamma_{N_x} = -\frac{(\lambda_w A_w)_{N_x - \frac{1}{2}}}{\Delta x_{N_x, N_x - 1}},$$
$$b_{N_x} = h_{w \to \infty} A_w T_\infty + \frac{(\rho_w C_{p,w} V_w)_{N_x}}{\Delta t} \mathbf{T_{w,N_x}}^{(k)}.$$

Derivative Contributions. Among the matrix coefficients, only α_1 depends on the control variables. Hence, $\mathrm{d}A/\mathrm{d}\mathbf{x}$ is zero everywhere except for entry $(1, 1)$:

$$\frac{\mathrm{d}\alpha_1}{\mathrm{d}\mathbf{x}} = 12 \sigma \epsilon_w V F_{h \to w} A_w (\mathbf{T_{w,1}}^{(k)})^2 \frac{\mathrm{d}\mathbf{T_{w,1}}^{(k)}}{\mathrm{d}\mathbf{x}}.$$

The derivative of the right-hand-side vector $\mathbf{b}$ is non-zero for all nodes. Applying the chain and product rules, noting that $\mathrm{d}T_\infty/\mathrm{d}\mathbf{x} = 0$:

For $i = 1$:

$$\frac{\mathrm{d}\mathbf{b}_1}{\mathrm{d}\mathbf{x}} = h_{f \to w} A_w \frac{\mathrm{d}T_f^{(k+1)}}{\mathrm{d}\mathbf{x}} + 4 h_{\mathrm{rad}} A_w (T_h^{(k+1)})^3 \frac{\mathrm{d}T_h^{(k+1)}}{\mathrm{d}\mathbf{x}}$$
$$+ \frac{(\rho_w C_{p,w} V_w)_1}{\Delta t} \frac{\mathrm{d}\mathbf{T_{w,1}}^{(k)}}{\mathrm{d}\mathbf{x}} + 12 h_{\mathrm{rad}} A_w (\mathbf{T_{w,1}}^{(k)})^3 \frac{\mathrm{d}\mathbf{T_{w,1}}^{(k)}}{\mathrm{d}\mathbf{x}}.$$

For $i = 2, \ldots, N_x$:

$$\frac{\mathrm{d}\mathbf{b}_i}{\mathrm{d}\mathbf{x}} = \frac{(\rho_w C_{p,w} V_w)_i}{\Delta t} \frac{\mathrm{d}\mathbf{T_{w,i}}^{(k)}}{\mathrm{d}\mathbf{x}}.$$

Circular Dependency. Because the derivative of the inner wall nodes depends on the furnace temperature states, and vice versa, a circular dependency arises. There are multiple approaches to handle this problem:

1. **Fixed-point iteration:** Initialize $\frac{\mathrm{d}\mathbf{T_{w,1}}^{(k+1)}}{\mathrm{d}\mathbf{x}}$, compute $\frac{\mathrm{d}\mathbf{T}^{(k+1)}}{\mathrm{d}\mathbf{x}}$, and iterate until convergence. This approximates the analytical derivative but requires a good initial guess.
2. **Monolithic sensitivity system:** Solve the coupled system of wall and furnace sensitivities simultaneously, yielding an accurate derivative without iterative guesses.

After trying both methods, the monolithic sensitivity system is preferred for its stability and higher accuracy.

Monolithic Sensitivity System. We define the stacked sensitivity vector

$$\mathbf{z} = \begin{bmatrix} \dfrac{\mathrm{d}\mathbf{T}^{(k+1)}}{\mathrm{d}\mathbf{x}} \\ \dfrac{\mathrm{d}\mathbf{T_w}^{(k+1)}}{\mathrm{d}\mathbf{x}} \end{bmatrix} \in \mathbb{R}^{4+6N_x},$$

where $\mathbf{T}^{(k+1)} \in \mathbb{R}^4$ contains the furnace, alloy, metal sheet, and heater temperature states, and $\mathbf{T_w}^{(k+1)} \in \mathbb{R}^{6N_x}$ concatenates all six walls.

We assemble and solve the system of equations:

$$J\,\mathbf{z} \;=\; \mathrm{RHS},$$

where the $(4+6N_x)\times(4+6N_x)$ Jacobian is given by

$$J = \begin{bmatrix} \dfrac{\mathrm{d}\mathbf{F}}{\mathrm{d}\mathbf{T}^{(k+1)}} & \underbrace{\Big[\dfrac{\mathrm{d}\mathbf{F}}{\mathrm{d}T_{w_1,1}^{(k+1)}} \otimes \mathbf{e}_1^T \quad \cdots \quad \dfrac{\mathrm{d}\mathbf{F}}{\mathrm{d}T_{w_6,1}^{(k+1)}} \otimes \mathbf{e}_1^T\Big]}_{4\times 6N_x} \\ -A^{-1}\,\dfrac{\mathrm{d}\mathbf{b}}{\mathrm{d}\mathbf{T}^{(k+1)}} & I_{6N_x} \end{bmatrix}.$$

Here, $\mathbf{e}_1^T = [1,0,\ldots,0] \in \mathbb{R}^{1\times N_x}$ is a standard basis vector that selects each wall's first node for the coupling to $\mathbf{F}$, and $\frac{\mathrm{d}\mathbf{b}}{\mathrm{d}\mathbf{T}^{(k+1)}} \in \mathbb{R}^{6N_x\times 4}$ captures every single wall node's dependence on the furnace temperature states.

The right-hand side is stacked as

$$\mathrm{RHS} = \begin{bmatrix} -\left(\frac{\mathrm{d}\mathbf{F}}{\mathrm{d}\mathbf{T}^{(k)}}\,\frac{\mathrm{d}\mathbf{T}^{(k)}}{\mathrm{d}\mathbf{x}} \;+\; \frac{\mathrm{d}\mathbf{F}}{\mathrm{d}\mathbf{T_{w,1}}^{(k)}}\,\frac{\mathrm{d}\mathbf{T_{w,1}}^{(k)}}{\mathrm{d}\mathbf{x}} \;+\; \frac{\mathrm{d}\mathbf{F}}{\mathrm{d}P^{(k)}}\,\frac{\mathrm{d}P^{(k)}}{\mathrm{d}\mathbf{x}}\right) \\ A^{-1}\left(\frac{\mathrm{d}\mathbf{b}}{\mathrm{d}\mathbf{T_w}^{(k)}}\,\frac{\mathrm{d}\mathbf{T_w}^{(k)}}{\mathrm{d}\mathbf{x}} - \frac{\mathrm{d}A}{\mathrm{d}\mathbf{x}}\,\mathbf{T_w}^{(k+1)}\right) \end{bmatrix}. \tag{7.27}$$

Once assembled, the sensitivity vector is obtained by solving

$$\mathbf{z} = J^{-1}\,\mathrm{RHS}.$$

This monolithic approach fully resolves the implicit coupling between all six walls and the furnace in a single solve.

Summary of Causal Chain. With $\mathbf{z}$ computed, almost all necessary ingredients for the full derivative $\frac{\mathrm{d}\mathcal{L}(\mathbf{x})}{\mathrm{d}\mathbf{x}}$ are assembled. The causal chain of the power derivative dependencies in the furnace model can be summarized as follows:

$$\frac{\mathrm{d}\mathcal{P}}{\mathrm{d}\mathbf{x}} \;\leftarrow\; \ldots\; \frac{\mathrm{d}P^{(k)}}{\mathrm{d}\mathbf{x}} \;\rightarrow\; \frac{\mathrm{d}\mathbf{T}^{(k+1)}}{\mathrm{d}\mathbf{x}},\; \frac{\mathrm{d}\mathbf{T_w}^{(k+1)}}{\mathrm{d}\mathbf{x}} \;\rightarrow\; \frac{\mathrm{d}P^{(k+1)}}{\mathrm{d}\mathbf{x}} \;\rightsquigarrow \text{recursive loop}.$$

Tracking and accumulating this chain yields the complete gradient of the loss function with respect to the control variable $\mathbf{x}$, after taking into account the derivative of the mechanical cost term.

Derivative of the mechanical-property term

The final mechanical properties are computed using the differentiable surrogate model $\mathbf{M}$ described in Section 7.3. Since the surrogate provides analytical derivatives with respect to the alloy temperature trajectory, and the temperature trajectory itself depends smoothly on $\mathbf{x}$, the derivative $\frac{\mathrm{d}\mathcal{M}(\mathbf{x})}{\mathrm{d}\mathbf{x}}$ is directly obtained via expansion of the chain rule:

$$\frac{\mathrm{d}\mathcal{M}(\mathbf{x})}{\mathrm{d}\mathbf{x}} = \frac{\mathrm{d}}{\mathrm{d}\mathbf{x}}\Big[w_{\sigma_y} \cdot \ell(\delta_{\sigma_y}(\mathbf{x}))\Big] + \frac{\mathrm{d}}{\mathrm{d}\mathbf{x}}\Big[w_{\sigma_u} \cdot \ell(\delta_{\sigma_u}(\mathbf{x}))\Big] + \frac{\mathrm{d}}{\mathrm{d}\mathbf{x}}\Big[w_{\varepsilon_u} \cdot \ell(\delta_{\varepsilon_u}(\mathbf{x}))\Big] \tag{7.28}$$

$$= \sum_{M \in \{\sigma_y, \sigma_u, \varepsilon_u\}} \frac{\mathrm{d}}{\mathrm{d}\mathbf{x}}\Big[w_M \cdot \ell(\delta_M(\mathbf{x}))\Big] \tag{7.29}$$

$$= \sum_{M \in \{\sigma_y, \sigma_u, \varepsilon_u\}} w_M \cdot \frac{\mathrm{d}\ell}{\mathrm{d}\delta_M} \cdot \frac{\mathrm{d}\delta_M(\mathbf{x})}{\mathrm{d}M} \cdot \frac{\mathrm{d}M}{\mathrm{d}\mathbf{x}}. \tag{7.30}$$

Derivative of the normalized deviations. This follows from the definition:

$$\delta_M(\mathbf{x}) = \frac{M(\mathbf{x}) - M_{\min}}{M_{\min}} \quad \Longrightarrow \quad \frac{\mathrm{d}\delta_M(\mathbf{x})}{\mathrm{d}M} = \frac{1}{M_{\min}}.$$

Derivative of the piecewise quadratic penalty-reward function. The derivative of ℓ with respect to δ_M is obtained directly:

$$\frac{\mathrm{d}\ell}{\mathrm{d}\delta_M} = \begin{cases} 2\delta_M, & \delta_M < 0, \\ -2\gamma\,\delta_M, & \delta_M \geq 0. \end{cases}$$

Derivative of the mechanical properties. Each property $M \in \{\sigma_y, \sigma_u, \varepsilon_u\}$ depends on the alloy trajectory through temperature T_m and final time t_f. Its derivative can therefore be decomposed as

$$\frac{\mathrm{d}M}{\mathrm{d}\mathbf{x}} = \sum_k \frac{\partial M^{(k)}}{\partial T_m^{(k)}} \cdot \frac{\mathrm{d}T_m^{(k)}}{\mathrm{d}\mathbf{x}} + \frac{\partial M}{\partial t_f} \cdot \frac{\mathrm{d}t_f}{\mathrm{d}\mathbf{x}}. \tag{7.31}$$

The temperature sensitivity $\frac{\partial M^{(k)}}{\partial T_m^{(k)}}$ follows from the phase-fraction model $\mathbf{X}$. Let $\mathbf{X}^{(k)}$ denote the phase fractions at the final timestep of the procedure. Then

$$\frac{\partial M^{(k)}}{\partial T_m^{(k)}} = \frac{\partial M^{(k)}}{\partial \mathbf{X}^{(k)}} \cdot \frac{\partial \mathbf{X}^{(k)}}{\partial T_m^{(k)}}. \tag{7.32}$$

The time sensitivity $\frac{\partial M}{\partial t_f}$ also follows from the surrogate model of phase fractions:

$$\frac{\partial M}{\partial t_f} = \frac{\partial M^{(N_t)}}{\partial \mathbf{X}^{(N_t)}} \cdot \frac{\partial \mathbf{X}^{(N_t)}}{\partial t_f}. \tag{7.33}$$

Deriving further expressions for $\frac{\partial M^{(N_t)}}{\partial \mathbf{X}^{(N_t)}}$, $\frac{\partial \mathbf{X}^{(k)}}{\partial T_m^{(k)}}$, and $\frac{\partial \mathbf{X}^{(N_t)}}{\partial t_f}$ require further insight into the mechanical model, which is beyond the scope of this work. It is therefore unnessecassary to go into further detail here.

8 Results and Discussion

In this chapter we present and critically examine the results obtained from the scenario-based optimization framework introduced in Chapter 7. The analysis begins with the industrially most relevant case: the maximization of overall mechanical performance. Here, different temperature-profile parametrizations and step configurations are compared to assess their impact on solution quality and efficiency. Building on these insights, we then broaden the perspective to the alternative scenarios outlined in the scenario matrix (Table 7.1), highlighting the trade-offs between time, energy consumption, and the individual mechanical properties. Finally, we investigate the convergence behavior of the employed optimizers and conclude with an experimental validation to demonstrate the practical relevance of the optimized solutions.

8.1 Main Scenario S3: Mechanical Performance Prioritization

Scenario S3 in Table 7.1 is considered the most relevant from an industrial perspective. It prioritizes the attainment of maximum mechanical performance, while still accounting for process duration and energy expenditure in a balanced manner. In this section, we present the optimal furnace temperature profiles obtained under this setting. Instead of reporting every parametrization and step count individually, we focus on the best-performing solutions for each curve type (constant, linear, spline) and compare them directly. A comprehensive overview of all tested parametrizations is provided in Table 8.5. The weighting factors applied in scenario S3 are summarized in Table 8.1, where w_{time}, w_{power}, and w_{mech} denote the relative importance of minimizing process duration, energy consumption, and deviations from target mechanical properties, respectively. The sub-weights w_{σ_y}, w_{σ_u}, w_{ε_u} further specify the prioritization among yield strength, ultimate tensile strength, and uniform elongation.

Scenario	w_{time}	w_{power}	w_{mech}	w_{σ_y}	w_{σ_u}	w_{ε_u}
S3	10	1	1000	2	1	0.5

Table 8.1: Weighting factors for scenario S3.

8.1.1 Optimal Solutions

Because of the flexibility and high dimensionality of the control variable, there is generally no unique global optimum. Instead, we identify representative solutions that achieve near-optimal performance for each profile type.

Constant Profiles

The resulting set-point trajectory does not represent a unique or global optimum. Similar outcomes would be obtained even if the first two set-point temperatures were slightly higher or

lower, or if the third set-point temperature were adjusted within a reasonable range. This is because the furnace's thermal response is governed by its physical and dynamic constraints—it can only heat up or cool down at rates permitted by its thermal inertia and heat transfer characteristics.

The optimal constant (three-step) solution thus serves as a practical baseline. Figure 8.1 and Table 8.2 show the best results obtained after repeated optimization runs initialized from different starting points. Although the set-point temperature profile T_f^* is parameterized as a three-step trajectory, the non-uniqueness of the solution effectively reduces it to a two-step trajectory. In practice, any number of set-point combinations with the same time parametrization would yield an equivalent outcome, as long as the resulting temperature evolution of the furnace remains unchanged.

Metric	Value
σ_y [MPa]	310.19
σ_u [MPa]	347.82
ε_u [%]	10.11
Loss (J)	-8.274665

Table 8.2: Mechanical properties for the optimal constant profile in scenario S3.

Linear Profiles

Allowing linearly interpolated set-points increases flexibility. The best-performing case is shown in Figure 8.2, with corresponding values summarized in Table 8.3. Compared with the constant set-point solution, the linear parametrization achieves slightly improved mechanical performance while maintaining a similar total process time. This improvement likely arises from the additional degrees of freedom in the parametrization, which enable the power profile to taper off more gradually rather than switching abruptly from full to zero power between time steps, as in the constant case. As before, the effective temperature trajectory ultimately collapses into two distinct phases—heating and cooling—reflecting the physical constraints of the furnace dynamics.

Metric	Value
σ_y [MPa]	310.00
σ_u [MPa]	347.74
ε_u [%]	10.12
Loss (J)	-8.320945

Table 8.3: Mechanical properties for the optimal linear profile in scenario S3.

Spline Profiles

Spline parametrizations provide even greater flexibility. The optimal solution (Figure 8.3, Table 8.4) once again reduces to what is effectively a two-step structure: rapid heating followed by passive cooling. The main benefit compared to constant profiles lies in smoother power modulation, which yields marginal improvements in mechanical performance.

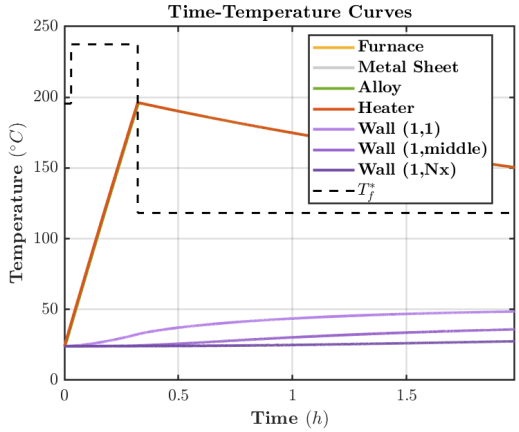

(a) Temperature profile

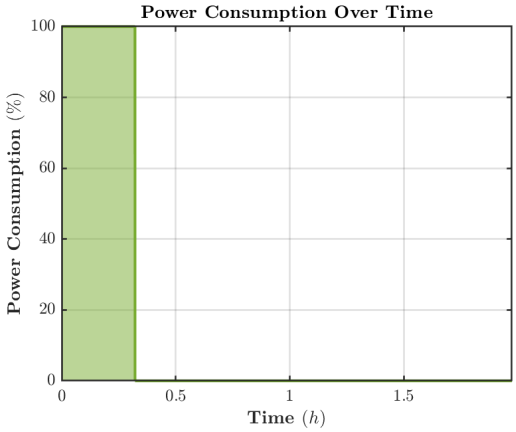

(b) Power input

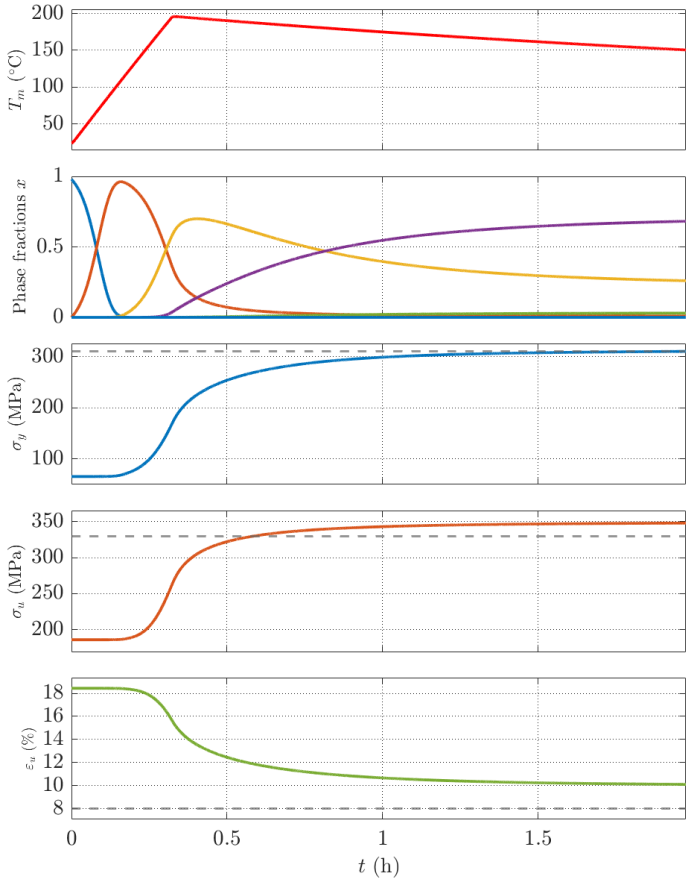

(c) Mechanical properties and phase fractions

Figure 8.1: Optimization results for scenario S3 using a constant temperature profile.

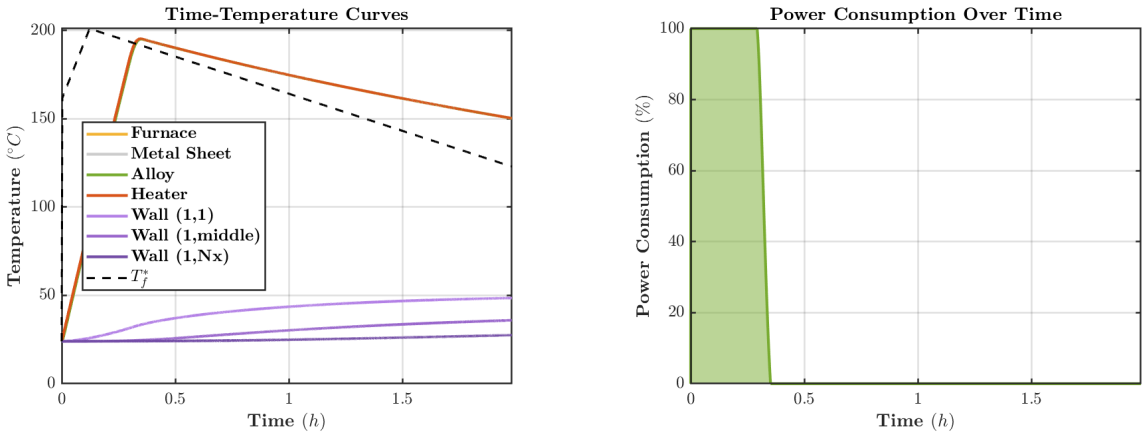

(a) Temperature profile (b) Power input

Figure 8.2: Optimization results for scenario S3 using a linear temperature profile.

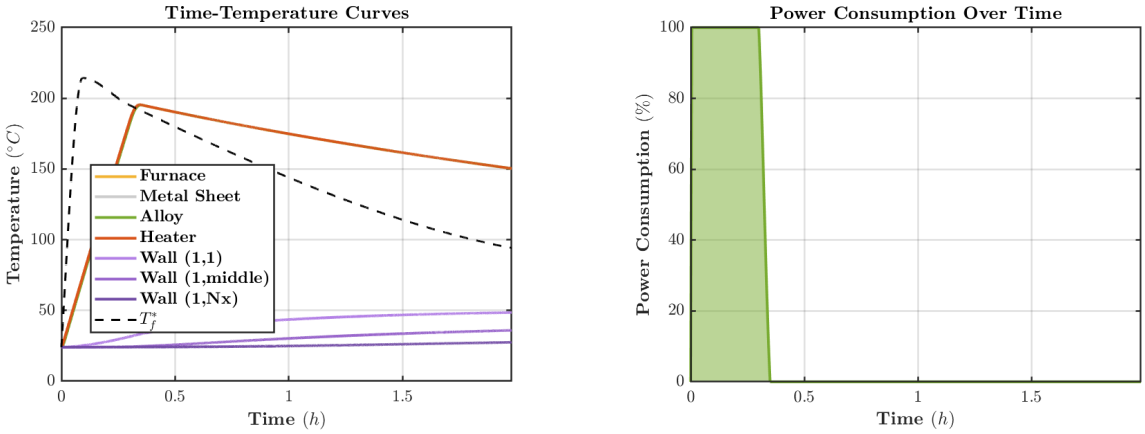

(a) Temperature profile (b) Power input

Figure 8.3: Optimization results for scenario S3 using a spline temperature profile.

Metric	Value
σ_y [MPa]	310.00
σ_u [MPa]	347.74
ε_u [%]	10.12
Loss	-8.321440

Table 8.4: Mechanical properties for the optimal spline profile in scenario S3.

Comparative Summary

Table 8.5 summarizes all tested parametrizations and step counts. In all cases, the optimizer converges to the same qualitative strategy: rapid heating followed by natural cooling. Introducing additional degrees of freedom (i.e., higher-order parametrizations) does not yield significant improvements, since the solutions effectively collapse to two-step trajectories.

Profile	Steps	Loss (J)	σ_y [MPa]	σ_u [MPa]	ε_u [%]
Spline	3	-8.321	310.00	347.74	10.12
Linear	3	-8.321	310.00	347.74	10.12
Constant	3	-8.275	310.19	347.82	10.11
Constant	4	-8.273	310.22	347.82	10.11
Constant	2	-8.164	310.00	347.98	10.14
Linear	2	-6.723	310.00	347.02	10.07
Spline	2	-6.631	310.01	347.01	10.07
Constant	1	-4.832	310.00	347.00	10.07

Table 8.5: Summary of optimal results for scenario S3 across different parametrizations.

8.1.2 Discussion

Across all parametrizations, the optimizer consistently selects a strategy of rapid furnace heating to the target temperature, followed by power-off cooling. Even when three-step or spline-based parametrizations are allowed, the resulting solutions reduce to effective two-step strategies, indicating that added flexibility does not translate to measurable performance gains. This observation supports the use of low-dimensional parametrizations for industrial application, which combine computational efficiency with practical implementation. The results also align with the concept of artificial aging, in which alloys are maintained at elevated temperatures for a controlled duration to maximize mechanical performance.

8.2 Alternative Scenarios

As emphasized in Chapter 7, in what sense a solution is considered "optimal" depends entirely on the relative weightings factors assigned to the different components of the objective function. In this section, we investigate how varying these weightings gives rise to distinct scenarios with correspondingly different solutions. It is important to note that when the weighting factors change, the objective function itself changes as well. Consequently, the absolute values of the

loss function evaluations are no longer directly comparable across scenarios. Instead, our analysis focuses on the resulting process characteristics and mechanical performance.

8.2.1 Scenario S1: Fast Cycle

In scenario S1, the primary objective is to minimize process duration, while still achieving at least the minimal acceptable levels of mechanical performance. The weighting factors for this case are adjusted accordingly, as summarized in Table 8.6.

Scenario	w_{time}	w_{power}	w_{mech}	w_{σ_y}	w_{σ_u}	w_{ε_u}
S3	1000	1	10	2	1	0.5

Table 8.6: Scenario S1 weighting factors.

Since the three-step temperature profile parametrizations yielded the most promising results in the previous section, we also adopt them here. The corresponding optimal solution is shown in Figure 8.4, while Table 8.7 summarizes the resulting mechanical properties.

Metric	**Value**
σ_y [MPa]	310.00
σ_u [MPa]	345.74
ε_u [%]	9.97

Table 8.7: Optimal mechanical properties for scenario S1.

Relative to the baseline optimization of scenario S3, scenario S1 demonstrates that the process can be completed in as little as half an hour while still satisfying the minimum mechanical requirements. Naturally, this comes at the expense of reduced strength and ductility compared to the more performance-oriented solutions.

8.2.2 Scenario S2: Energy Efficiency

In scenario S2, the focus shifts from time efficiency to energy conservation. While there may be some overlap between these two objectives, the weighting factors here are chosen to strongly penalize power consumption, as detailed in Table 8.8.

Scenario	w_{time}	w_{power}	w_{mech}	w_{σ_y}	w_{σ_u}	w_{ε_u}
S3	1	1000	10	2	1	0.5

Table 8.8: Scenario 2 weightings.

The optimization results are shown in Figure 8.5 and Table 8.9. This particular case exhibits somewhat unusual behavior that warrants additional explanation. To ensure consistency with the measurement data, the furnace temperature (including all domains) was initialized at 24°C, while the external environment was set to 25°C. As a result, the process initially exhibits a long phase in which the furnace very gradually warms from 24 to 25°C through natural heat diffusion. Only afterwards does a short but sharp power input occur to achieve the required mechanical

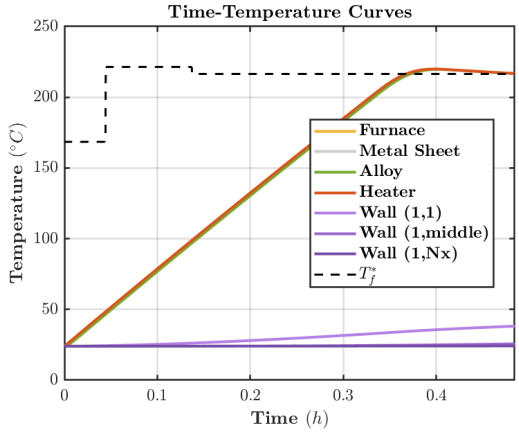

(a) Temperature profile

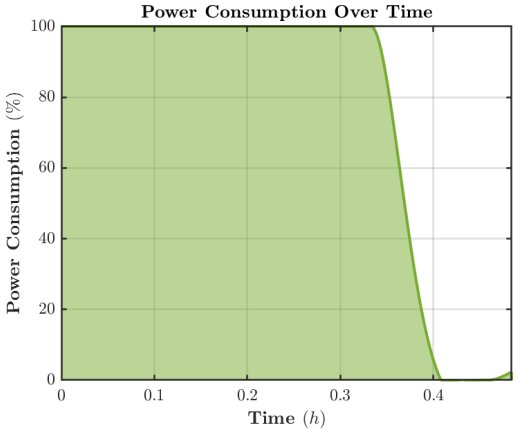

(b) Power input

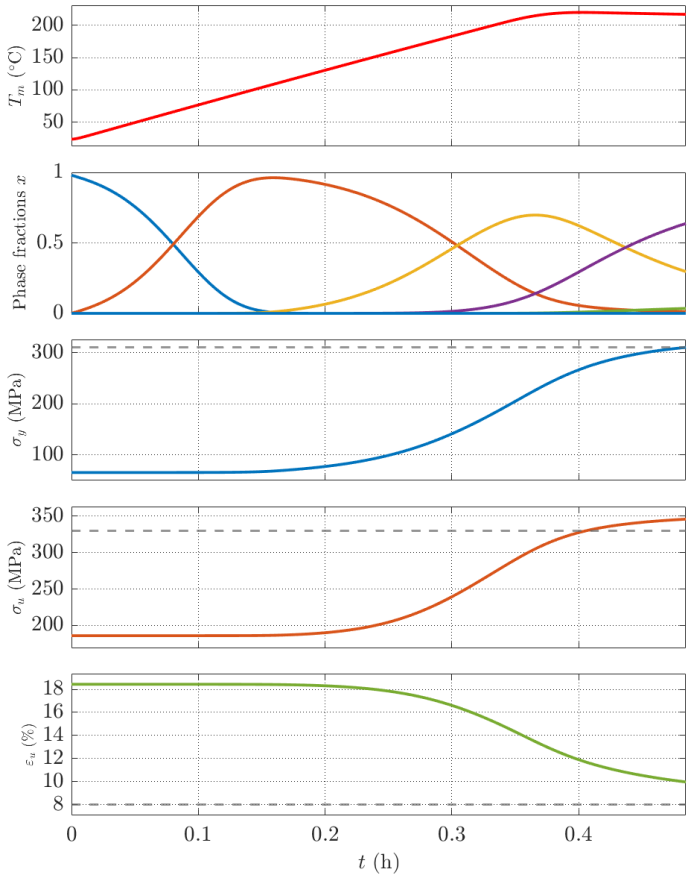

(c) Mechanical properties and phase fractions

Figure 8.4: Optimization results for scenario S1.

performance. Had the furnace been initialized directly at 25°C, the optimization outcome would have differed substantially. In this case, the relatively low weighting on process duration allows the optimizer to exploit the extended idle phase in order to minimize total power consumption.

Metric	Value
σ_y [MPa]	310.00
σ_u [MPa]	347.71
ε_u [%]	10.12

Table 8.9: Optimal mechanical properties for scenario S2.

8.2.3 Scenario S4–S6: Mechanical Property Emphasis

The final set of scenarios examine the effect of prioritizing individual mechanical properties within the loss function. The weighting factors follow the general structure of Table 8.1, with property-specific adjustments as defined in Table 7.1. A condensed overview is provided in Table 8.10.

Scenario	w_{time}	w_{power}	w_{mech}	w_{σ_y}	w_{σ_u}	w_{ε_u}
S4	10	1	1000	100	0	0
S5	10	1	1000	0	100	0
S6	10	1	1000	0	0	100

Table 8.10: Scenario S4-S6 weighting factors.

These results must be interpreted with care. Yield strength (σ_y) and ultimate tensile strength (σ_u) tend to evolve in parallel, while uniform elongation (ε_u) behaves in the opposite way, decreasing monotonically with time. As a result, we expect scenarios prioritizing σ_y and σ_u to yield broadly similar outcomes, whereas prioritizing ε_u leads to markedly different heating strategies and trade-offs.

Scenario S4: Yield Strength Priority

Figure 8.6 and Table 8.11 show the optimization results when σ_y is emphasized. As expected, this produces higher σ_y values relative to prior scenarios, accompanied by a marginal increase in σ_u compared to the baseline scenario S3. The improvement in strength, however, comes at the expense of ductility: uniform elongation drops noticeably relative to the balanced case. In fact, at 8% the uniform elongation is at its lowest permissible level compatible with the constraints.

Metric	Value
σ_y [MPa]	327.33
σ_u [MPa]	351.99
ε_u [%]	8.00

Table 8.11: Optimal mechanical properties for scenario S4.

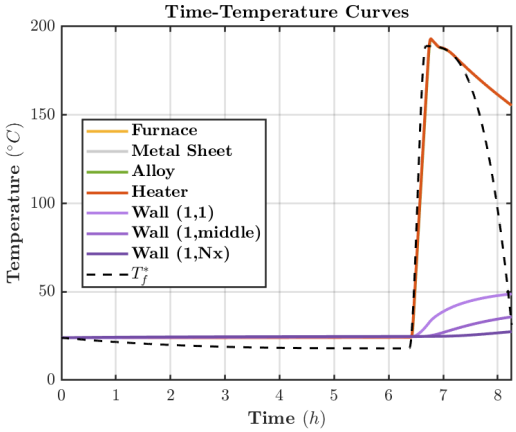

(a) Temperature profile

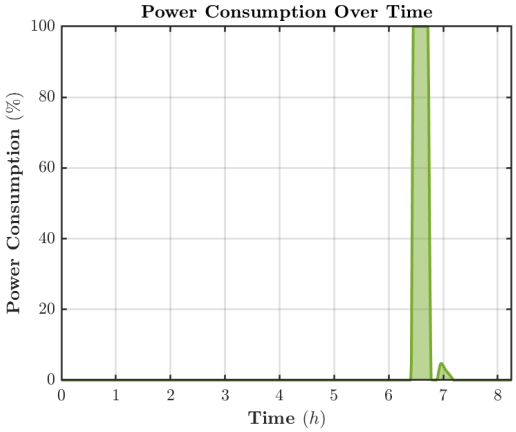

(b) Power input

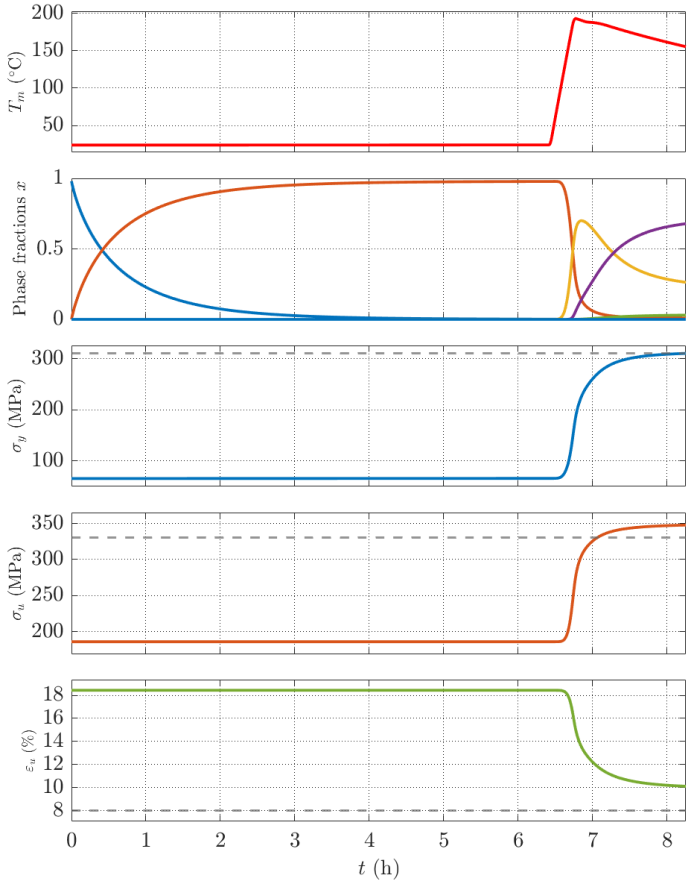

(c) Mechanical properties and phase fractions

Figure 8.5: Optimization results for scenario S2.

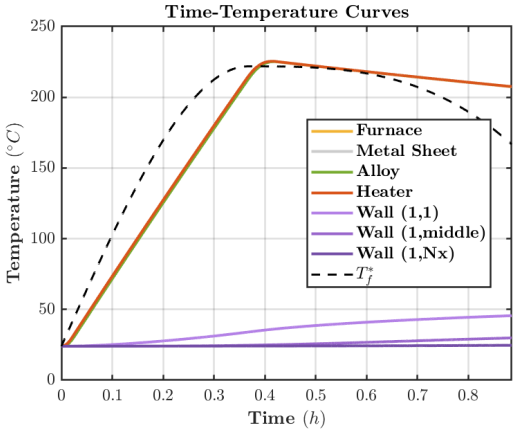

(a) Temperature profile

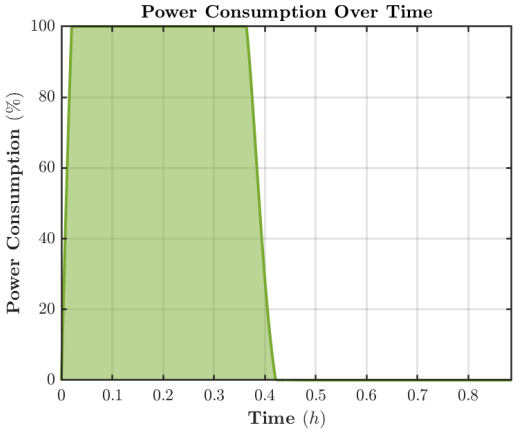

(b) Power input

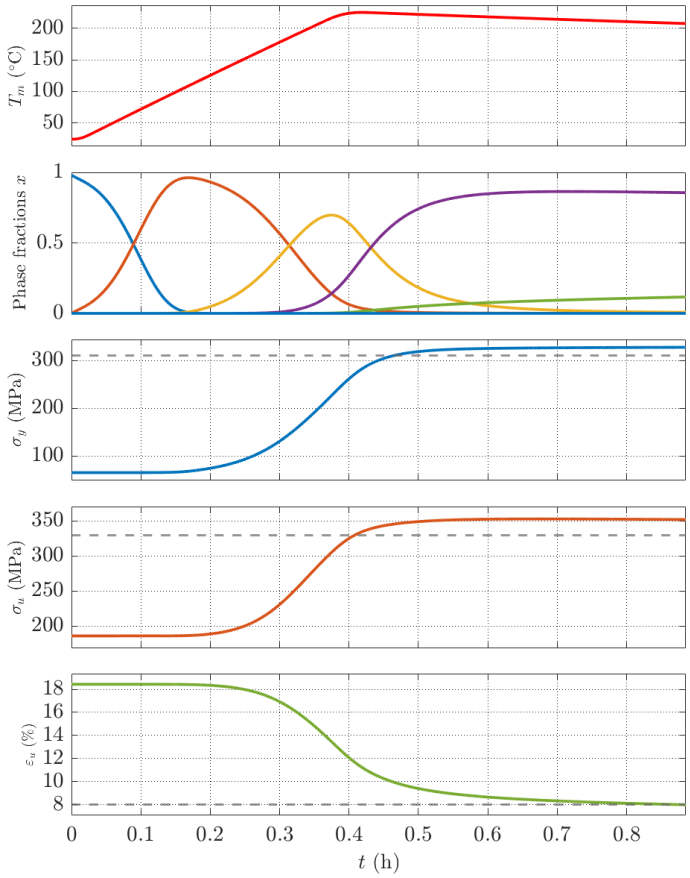

(c) Mechanical properties and phase fractions

Figure 8.6: Optimization results for scenario S4 (yield strength emphasis).

Scenario S5: Ultimate Tensile Strength Priority

In scenario S5, σ_u is prioritized above all other factors. As expected, the results (Figure 8.7, Table 8.12) still resemble those of scenario S4, reflecting the strong correlation between σ_y and σ_u. However, a closer look at the numerical results reveals that, while the ultimate tensile strength σ_u and elongation ε_u increase in S5 compared to S4, the yield strength σ_y is lower. This indicates a subtle trade-off: optimizing primarily for σ_u can yield improvements in ductility and ultimate strength at the expense of yield strength. Even strongly correlated properties do not always peak simultaneously, emphasizing the importance of carefully choosing priority objectives when using local gradient-based solvers.

Metric	Value
σ_y [MPa]	323.81
σ_u [MPa]	354.33
ε_u [%]	8.74

Table 8.12: Optimal mechanical properties for scenario S5.

Scenario S6: Uniform Elongation Priority

Finally, scenario S6 focuses on maximizing ductility. In contrast to S4 and S5, improving ε_u requires operating in conditions where strength is reduced. The results (Figure 8.8, Table 8.13) confirm this: ε_u is substantially higher than in the previous two scenarios, but σ_y and σ_u decline accordingly. This underscores the inherent trade-off between strength and ductility in the process. Notably, it appears infeasible to achieve ε_u values above 10.14% while maintaining σ_y above 310 MPa, suggesting a hard constraint imposed by the underlying material behavior.

Metric	Value
σ_y [MPa]	310.00
σ_u [MPa]	348.01
ε_u [%]	10.14

Table 8.13: Optimal mechanical properties for scenario S6.

Taken together, scenarios S4-S6 demonstrate how shifting emphasis within the objective function shapes the optimization outcome. Prioritizing σ_y or σ_u leads to comparable solutions due to their strong correlation, though solver sensitivity can produce variations in the balance between strength and ductility. In contrast, emphasizing ε_u results in a fundamentally different strategy, where ductility is enhanced only at the expense of strength. These findings underline the importance of explicitly defining performance priorities in industrial practice, as the optimization inherently reflects the trade-offs dictated by the material's physical limits.

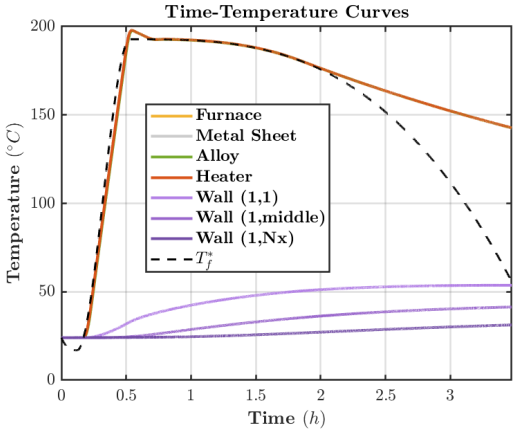

(a) Temperature profile

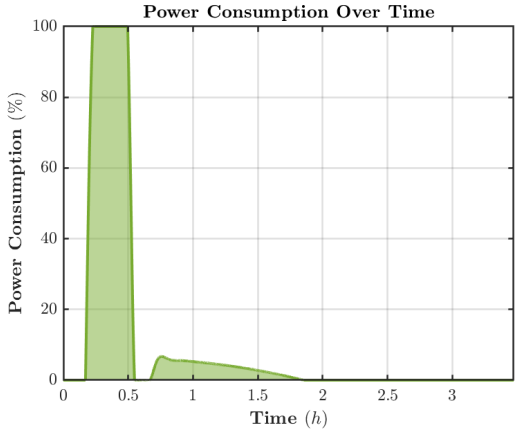

(b) Power input

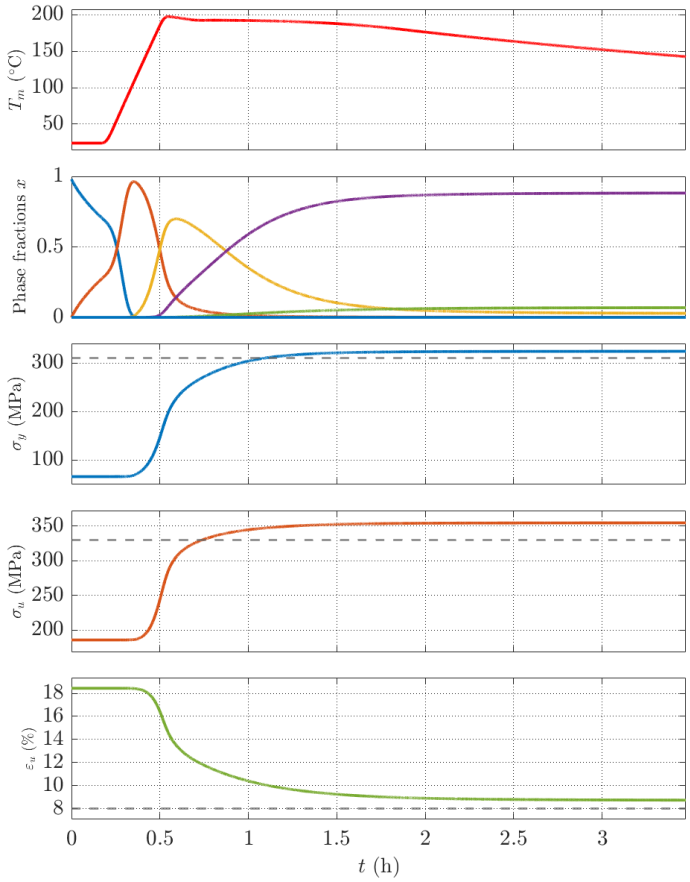

(c) Mechanical properties and phase fractions

Figure 8.7: Optimization results for scenario S5 (ultimate tensile strength emphasis).

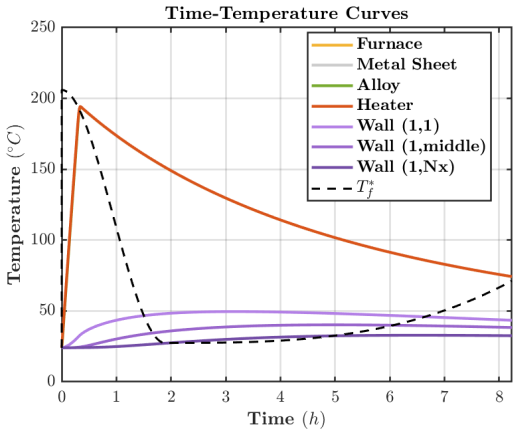

(a) Temperature profile

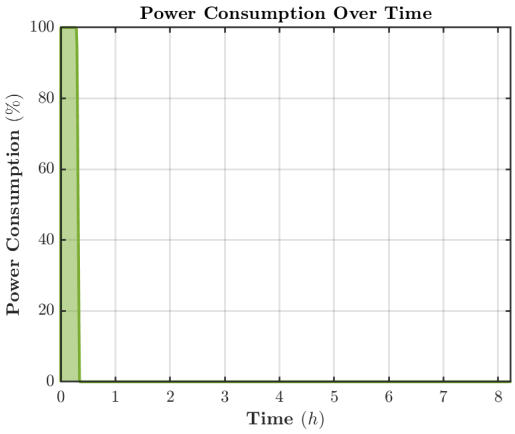

(b) Power input

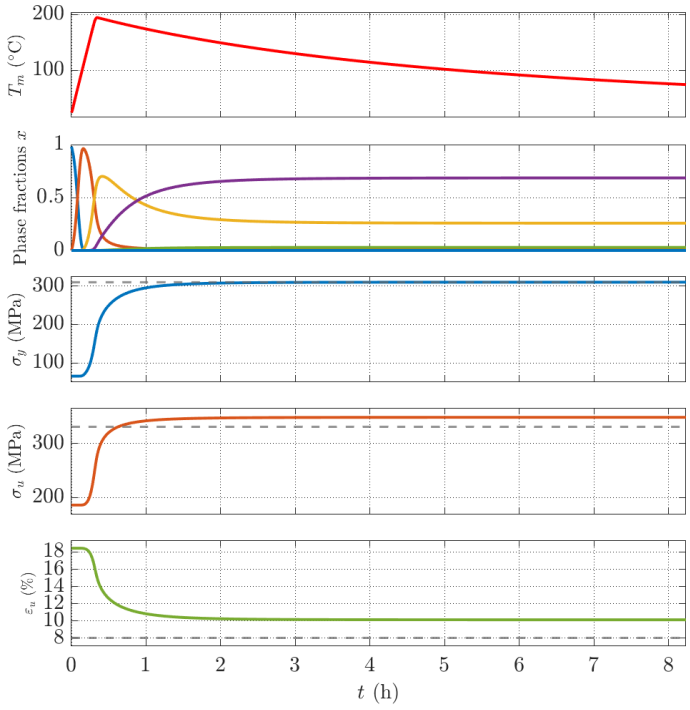

(c) Mechanical properties and phase fractions

Figure 8.8: Optimization results for scenario S6 (uniform elongation emphasis).

8.3 Extended Model: Alloy Temperature Distribution

In the results presented thus far, the alloy element has been treated as a lumped mass with a uniform temperature. In reality, the size and geometry of the alloy influence the spatial and temporal temperature gradients within the material. This is especially prevalent in larger materials such as industrial steel or aluminum coils. In this section, we extend the furnace model to compute and track the temperature distribution of an aluminum coil (modeled as a simple cylindrical alloy element) over time, and explore how these gradients may affect the predicted mechanical performance.

It is important to note that these results cannot be experimentally validated, and the mechanical model—originally calibrated for the specific alloy specimen as depicted in Figure 6.1—is unlikely to be fully reliable for this theoretical cylindrical element. Nonetheless, this exercise demonstrates that the furnace model can be readily extended to account for distributed temperature fields within alloy elements.

Following the methodology of finite-volume spatial discretization described in Chapter 5, the cylindrical alloy is modeled as a one-dimensional heat conduction problem. The optimization routine is adjusted to enforce mechanical performance constraints across the entire element, rather than only at the surface. Specifically, both the inner and outer nodes must satisfy the mechanical constraints, ensuring that intermediate nodes also remain feasible.

Figure 8.9 illustrates that, as proof-of-concept, the current optimization framework can identify a furnace trajectory that meets minimal mechanical performance requirements across the alloy's temperature distribution. This example serves purely as a conceptual demonstration; no experimental validation has been performed.

8.4 Convergence Analysis

While the primary focus of this work lies in obtaining reliable and physically meaningful optimization results, the computational efficiency of the solvers used is also of practical relevance. This section therefore provides a comparative discussion of solver behavior and convergence characteristics. The intent here is not to perform a rigorous benchmarking study, but rather to offer a qualitative assessment of how solver design choices affect performance and solution quality. Specifically, we examine the performance of MATLAB's `fmincon` solver in comparison with a custom projected gradient descent (PGD) implementation, and assess the impact of different step-size strategies on convergence rates.

Solver Performance

As noted earlier, `fmincon` can be executed either with or without a supplied analytical gradient. Figure 8.10 illustrates captures a consistent trend: supplying the analytical gradient reduces the number of iterations, but increases the computational cost of every iteration, resulting in an overall longer computation time. Furthermore, running with the symbolic gradient tended to yield marginally worse solutions. This suggests that the accumulated numerical errors in

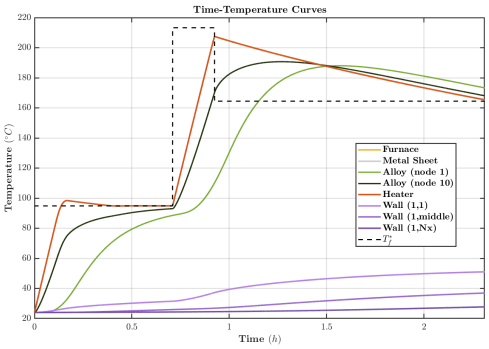

(a) Temperature profile

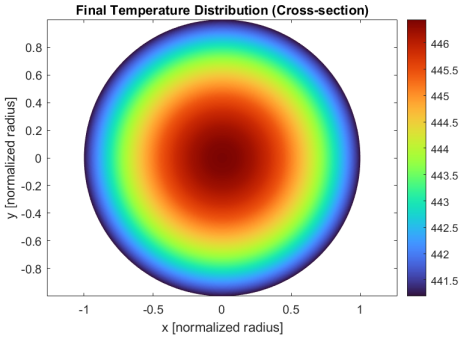

(b) Final temperature distribution heatmap

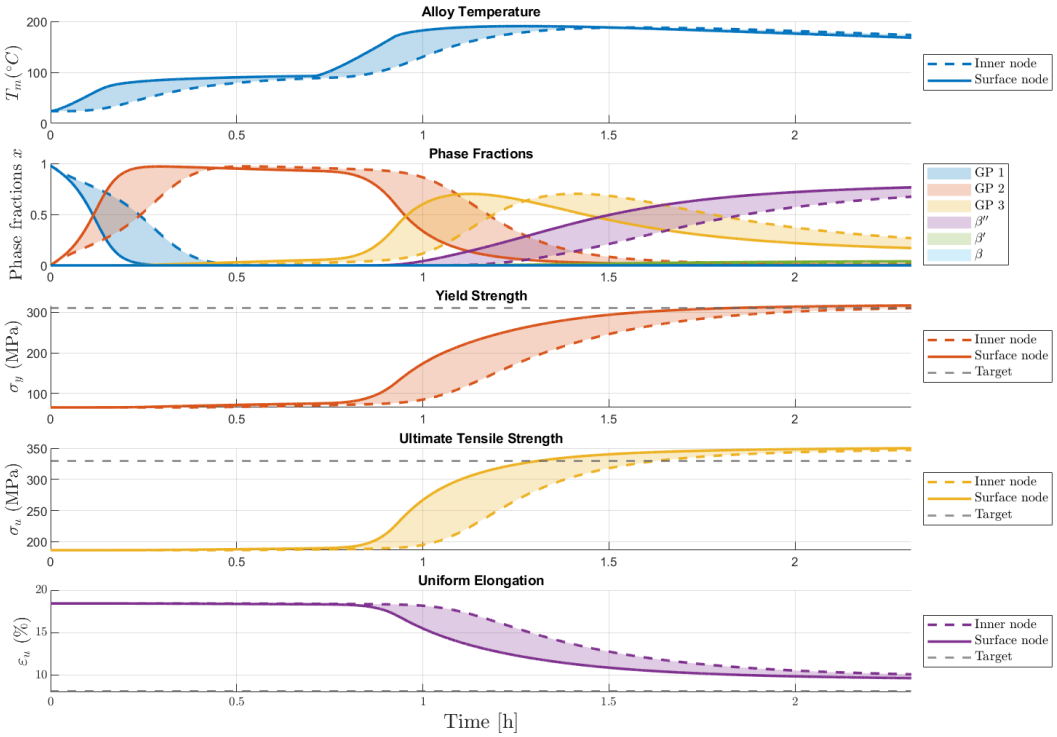

(c) Mechanical properties and phase fractions across the alloy

Figure 8.9: Optimization results for scenario S3 accounting for spatial temperature distribution within the alloy.

the thermal and mechanical property models can guide the optimization into sub-optimal local minima when the gradient is enforced too strictly. This may be linked to scaling issues or nonsmooth behavior in the mechanical model, suggesting that future work could explore rescaling strategies or hybrid gradient formulations to mitigate this effect.

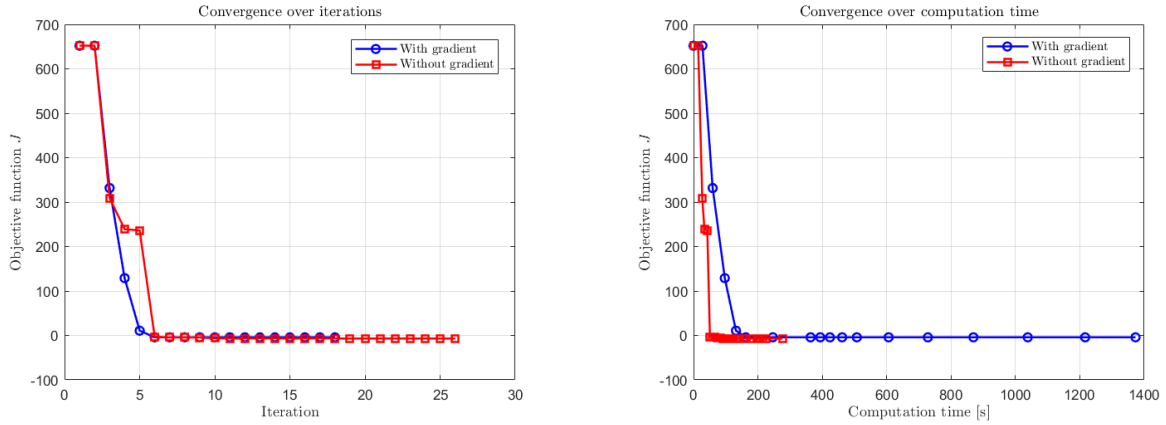

(a) Convergence over iterations

(b) Convergence over run-time

Figure 8.10: Convergence of `fmincon` with and without supplied analytical gradient for scenario S3.

We next compare `fmincon` with the PGD approach, using identical stopping criteria where possible. For scenario S3, Figure 8.11 exemplifies another trend we observed. When applying the PGD algorithm, both the total and per iteration computational cost were larger. In addition, the solution often exhibited a zigzagging convergence pattern. This oscillatory convergence pattern is characteristic of projection-based methods when the gradient update repeatedly drives the solution outside the feasible set, followed by a projection step that restores feasibility. Near active constraints, the projection operator can dominate the update, effectively canceling progress along the gradient direction and causing the solution to alternate between feasible and infeasible points. The result is a stagnation of the objective value and slow apparent convergence, particularly in regions where the feasible space has sharp or non-smooth boundaries. Even with adaptive step-sizing [26], this effect remained pronounced due the sharpness of the projection operator.

This limitation illustrates a well-known drawback of first-order projection-based methods, particularly when compared with more sophisticated constrained optimizers like `fmincon`, which handle active constraints more smoothly through interior-point or SQP frameworks. While this effect can often be mitigated by introducing smoother regularization or reformulating the constraints, no such modifications were applied here. The purpose of including PGD in this study was to provide a baseline for comparison with higher-level constrained solvers such as `fmincon`, rather than to refine PGD's performance itself. Maintaining the algorithm in its un-modified form therefore allows for a clearer illustration of the inherent limitations of basic projection-based methods—specifically, their sensitivity to active constraints and a tendency to stagnate near local minima. This behavior therefore serves as an instructive reference point rather than a shortcoming of the implementation.

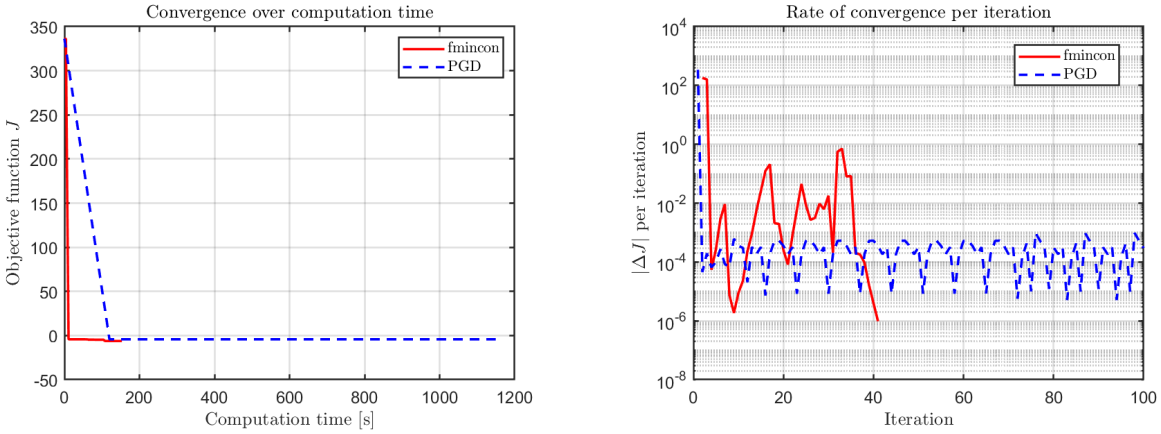

(a) Convergence over run-time

(b) Comparison per iteration

Figure 8.11: Comparison of `fmincon` and PGD convergence for scenario S3.

8.5 Optimization Validation

An experimental validation of the proposed optimization approach is currently being conducted to compare the predicted and measured alloy properties and temperature profile. The results of this validation will be discussed in future work once available.

8.6 Discussion

The results presented in this chapter highlight both the strengths and the limitations of the proposed optimization framework. By combining a physics-based furnace model with a flexible scenario-based optimization strategy, we demonstrated that the design of time-temperature curves can be systematically tailored to different industrial priorities. The framework produced solutions that were physically reasonable and applicable in practice. At the same time, the observed discrepancies underscore the importance of addressing the simplifications inherent in the furnace and material models.

A central finding is that optimization outcomes strongly depend on the weighting of objectives. Scenario S3, which emphasized a balanced improvement of mechanical properties, yielded the most industrially relevant solutions, but not the highest mechanical performance for individual properties. The analysis of alternative scenarios made it clear that prioritizing yield strength (σ_y) or ultimate tensile strength (σ_u) leads to similar optimal curves due to the correlation between the two, while emphasizing uniform elongation (ε_u) generates qualitatively different heating strategies that enhance ductility at the cost of strength. These results underline the value of the scenario matrix: by adjusting the relative weighting factors, industrial decision-makers can transparently explore the trade-offs between process duration, energy expenditure, and mechanical performance, rather than relying on heuristic trial-and-error.

The comparison of different curve parametrizations further showed how solution quality and

computational burden are interlinked. Increasing the flexibility of the parametrization (e.g., spline vs. linear vs. constant) generally improved mechanical outcomes but also made the optimization problem more difficult to solve. In practical settings, this trade-off may guide the choice of parametrization: while splines can provide fine-grained control over heating schedules, simpler representations may be preferable when computational efficiency is critical or when solutions must be easily interpretable by furnace operators.

The convergence analysis highlighted key differences between solvers. MATLAB's `fmincon` proved robust across scenarios, particularly in enforcing feasibility through its constraint-handling capabilities. However, supplying analytical gradients actually degraded performance, suggesting that small inconsistencies between the gradient expressions and the discretized models of furnace dynamics and mechanical properties may lead the optimizer toward suboptimal local minima. This finding indicates a need for more careful integration of analytical derivatives or, alternatively, the development of support for automatic differentiation. In contrast, projected gradient descent (PGD) exhibited slower convergence in terms of both iterations and computation time, and its tendency to stagnate near local minima produced solutions of inferior quality.

From an industrial perspective, the framework demonstrates a pathway toward data- and model-driven furnace operation. Instead of relying on operator experience or overly conservative process windows, optimal time-temperature curves can be computed for a given set of priorities, such as maximizing ductility, minimizing cycle time, or balancing multiple objectives. While the absolute solutions may not transfer directly to all furnace designs or alloy systems, the observed trade-offs and optimization strategies remain highly relevant. For instance, the consistent tension between maximizing mechanical strength and ductility reflects a fundamental physical limitation, not a modeling artifact, and thus provides valuable guidance for process planning.

Nevertheless, several limitations should be acknowledged. The furnace model employed a lumped air temperature, alloy system, and heating mass assumptions, as well as one-dimensional wall conduction, all of which simplify the actual three-dimensional heat transfer processes in the industrial furnace. The physical parameter choices, while sufficient to capture general trends, rely on limited calibration data. The solver comparisons were based on representative runs rather than statistically rigorous benchmarks, and computation time was not recorded systematically across all experiments. These limitations temper the strength of quantitative conclusions but do not diminish the qualitative insights gained.

Overall, the study shows that optimization of time-temperature curves is both feasible and informative. The scenario-based framework enables structured exploration of trade-offs, solver comparisons reveal strengths and weaknesses of different approaches, and the proposed solutions are practical. The findings collectively point toward a future where heat treatment processes can be tuned more intelligently, balancing productivity and performance with a level of transparency and efficiency not easily achievable through manual trial-and-error.

In summary, the discussion has highlighted the most important trends, trade-offs, and limitations of the optimization framework. While the results demonstrate both practical relevance and methodological promise, they also point to open questions regarding model fidelity, solver

design, and integration into industrial workflows. The following chapter draws these threads together by presenting the overall conclusions of the thesis, outlining the main contributions, acknowledging limitations, and suggesting directions for future research.

9 Conclusion

This thesis has developed and applied a numerical framework for modeling and optimizing the thermal behavior of a radiatively heated furnace system containing a heat-treatable aluminum alloy. The furnace model captured the coupled effects of conduction, convection, and radiation, enabling approximate prediction of time-temperature curves and their influence on mechanical properties. Building on this foundation, an optimization framework was proposed to design furnace trajectories that balance industrial performance goals such as process time, power expenditure, and mechanical outcomes.

9.1 Summary of Findings

The results demonstrated that the model can reproduce realistic heating dynamics and predict trends in yield strength, ultimate tensile strength, and uniform elongation as functions of time-temperature exposure. The scenario-based optimization study highlighted how different weighting strategies alter the optimal furnace trajectory: balanced optimization (scenario S3) produced mechanically robust outcomes, while single-property emphasis (scenarios S4-S6) revealed the fundamental trade-offs between strength and ductility. Solver comparisons showed that `fmincon` consistently produced higher-quality solutions than projected gradient descent, and at a lower computational cost. Although experimental validations of the proposed solutions were not yet available at the time of writing, they are expected to confirm the trends observed in simulation.

9.2 Contributions

This work makes several key contributions to the modeling and optimization of industrial aluminum heat treatment. First, a furnace thermal model has been developed that incorporates conduction, convection, and radiation in a tractable form suitable for optimization. Second, this work illustrates the decision-making process and the trade-offs encountered during the calibration of the furnace model using experimental measurement data. Third, a scenario-based optimization framework is introduced, explicitly accounting for the industrial trade-offs between process time, power consumption, and resulting mechanical properties. Finally, a short comparative analysis of different solver strategies is presented, highlighting the advantages and limitations of both black-box solvers and custom gradient-based methods.

9.3 Limitations

The presented model and optimization framework are subject to several simplifying assumptions and methodological limitations. For instance, radiative heat transfer was treated under a gray-body assumption without taking into account spectral effects, the furnace wall geometry was

reduced to one dimension, and lateral heat conduction was ignored. Heating elements and the thin metal sheets were modeled as lumped domains rather than spatially resolved structures, and the alloy itself was modeled as a uniform mass without internal temperature distribution. These assumptions improve computational efficiency but may omit local gradients relevant to larger or more complex components.

Furthermore, calibration and validation of the physical model were performed on a limited experimental dataset with sparse sensor coverage, which captured representative but not exhaustive operating scenarios. This limited coverage constrains the identifiability of the thermophysical parameters, potentially allowing parameter compensations that reproduce the observed data but are not unique. Consequently, the predictive accuracy of the resulting model outside the calibrated range remains uncertain.

In addition, the calibration occasionally produced extreme parameter values near physical or imposed bounds, indicating possible non-identifiability or parameter coupling. A sensitivity analysis—such as finite-difference or global Sobol indices [43]—could help clarify which parameters are well-informed by the data. Uncertainty quantification through confidence intervals, or Bayesian posteriors would further strengthen the model's predictive capacity [29]. Moreover, while upper and lower bounds were imposed during calibration, no additional penalties or prior distributions were used to favor physically realistic values. Incorporating such soft constraints (e.g., L2 regularization or prior-based penalties [58]) can help reduce overfitting and avoid parameter drift in poorly informed regions of the parameter space.

While the PGD algorithm was employed primarily as a reference for comparison with `fmincon`, further improvements could be achieved by exploring adaptive projection strategies into feasible regions, or other modern gradient-based optimizers such as Adam or L-BFGS [24]. Additionally, convergence across multiple initializations was not systematically analyzed and could represent an avenue for future refinement. Finally, classical numerical methods were deliberately chosen for transparency and interpretability—specifically, the finite volume method (FVM) for spatial discretization and implicit Euler for time integration. More recently developed machine-learning methods—such as extreme learning machines (ELMs) or physics-informed neural networks (PINNs)—may offer opportunities for reduced computational cost [21, 39].

Beyond algorithmic choices, computational efficiency could also benefit from the use of automatic differentiation (AD) for gradient evaluation. While MATLAB does not natively provide full AD support, integrating such capabilities (e.g., by tailoring code design for AD) could substantially reduce the run-time cost of gradient-based optimization. In contrast, languages such as Python or Julia, which offer more mature AD frameworks, may be better suited for large-scale implementations of this approach.

9.4 Future Research Directions

Several avenues for future work emerge from this study. One direction involves refined temperature modeling by incorporating a spatially resolved temperature distribution for alloy elements, which would capture gradients between the surface and interior and allow more accurate predic-

tions of local mechanical properties and potential surface effects. Improved radiation modeling could also be pursued, extending the current model to account for spectral absorption characteristics and gas participation, thereby increasing the physical fidelity of the furnace simulation. Additionally, simulating complex alloy geometries, including various shapes and sizes of workpieces, would better reflect industrial diversity and capture their influence on heating rates and mechanical outcomes.

Future studies could also integrate detailed models of dynamic furnace control, such as door operations or slit openings, to emulate adaptive cooling strategies that may improve mechanical performance. Alternative furnace model strategies could be explored as well, including multiple furnaces for rapid alloy transfer or accounting for mechanical contributions to specific heat, broadening the operational modeling scope. Advanced curve parametrization methods, such as adaptive time discretization for temperature profiles, could refine the grid near regions of high non-uniformity to achieve smoother, more precise optimal trajectories, with temperature-dependent constraints on heating and cooling rates to avoid non-unique local solutions.

Enhancing optimization methods is another promising avenue. Although the projected gradient-based approach used in this study provides a clear and interpretable baseline, modern optimization techniques offer substantial opportunities for improvement. Incorporating advanced algorithms—such as quasi-Newton methods, stochastic or adaptive gradient schemes, and machine-learning-driven optimizers—could reduce computational cost while improving convergence reliability. Furthermore, integrating automatic differentiation (AD) frameworks would enable efficient and accurate gradient evaluations, facilitating the application of more sophisticated solvers and enabling large-scale optimization with minimal manual derivation effort. Finally, comprehensive benchmarking of alternative solvers, assessed through metrics such as computation time, iteration count, and solution robustness, would provide valuable guidance for future industrial implementations and practical deployment.

References

[1] Karl J. Åström and Tore Hägglund. *PID Controllers: Theory, Design, and Tuning*. Instrument Society of America, 2nd edition, 1995. 18, 54, 58

[2] Kendall E. Atkinson and Weimin Han. *Elementary Numerical Analysis*. John Wiley & Sons, Hoboken, NJ, 3rd edition, 2003. 35

[3] Adrian Bejan. *Convection Heat Transfer*. John Wiley & Sons, Hoboken, NJ, 4th edition, 2013. 23

[4] Theodore L. Bergman and Adrienne S. Lavine. *Fundamentals of Heat and Mass Transfer*. John Wiley & Sons, Hoboken, NJ, 8th edition, 2017. 17, 24

[5] C. H. Bischof, H. M. Bücker, B. Lang, A. Rasch, and A. Vehreschild. Combining source transformation and operator overloading techniques to compute derivatives for matlab programs. In *Proceedings of the Second IEEE International Workshop on Source Code Analysis and Manipulation (SCAM 2002)*, Los Alamitos, CA, 2002. IEEE Computer Society. 75

[6] James P. Boyle and Richard L. Dykstra. A method for finding projections onto the intersection of convex sets in hilbert spaces. In *Advances in Order Restricted Statistical Inference*, volume 37 of *Lecture Notes in Statistics*, pages 28–47. Springer, New York, NY, 1986. 76

[7] Luca Cacciotti et al. Effect of bulk density on flame resistance of rockwool in combined fire-resistant facings. *Journal of Thermal Analysis and Calorimetry*, 147:11693–11704, June 2022. Open access. 53

[8] Mauro Carta, Leila Aydi, Pasquale Buonadonna, Donato Morea, and Mohamad El Mehtedi. Effect of post heat treatment on microstructure and mechanical properties of hot-rolled AA2017 aluminum alloy. *Heliyon*, 10(23):e40922, 2024. 2, 6

[9] M. W. Chase. *NIST-JANAF Thermochemical Tables*. American Chemical Society and American Institute of Physics for the National Institute of Standards and Technology, Woodbury, NY, 4 edition, 1998. Journal of Physical and Chemical Reference Data, Monograph No. 9. 52

[10] Han Chen, Garvesh Raskutti, and Ming Yuan. Non-convex projected gradient descent for generalized low-rank tensor regression. *Journal of Machine Learning Research*, 20:1–37, 2019. 75

[11] Yunmei Chen and Xiaojing Ye. Projection onto a simplex. *arXiv preprint arXiv:1910.09529*, November 2024. 76

[12] Samuel D. Conte and Carl de Boor. *Elementary Numerical Analysis: An Algorithmic Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 3rd edition, December 2018. 35

[13] Carl de Boor. *A Practical Guide to Splines*, volume 27 of *Applied Mathematical Sciences*. Springer-Verlag, New York, NY, 1st edition, 1978. 69

[14] R. C. Eberhart and J. Kennedy. A new optimizer using particle swarm theory. In *Proceedings of the Sixth International Symposium on Micro Machine and Human Science (MHS '95)*, pages 39–43, Nagoya, Japan, October 1995. IEEE. 61

[15] Engineering Toolbox. Emissivity coefficients of materials. https://www.engineeringtoolbox.com/emissivity-coefficients-d_447.html, 2025. Accessed: May 2025. 52, 53

[16] Enze Mfg Co. Ltd. 6082 aluminum alloy — everything you need to know. https://www.enzemfg.com/6082-aluminum-alloy-everything-you-need-to-know/, January 2024. Accessed: 2025-06-29. 49

[17] David R. Gaskell and David E. Laughlin. *Introduction to the Thermodynamics of Materials*. CRC Press, Boca Raton, FL, 6th edition, 2017. 27

[18] Alaa A. Hassan and Mohamed S. Hamed. Modeling of heat treatment of randomly distributed loads in multi-zone continuous furnaces. In *THERMEC 2011*, volume 706 of *Materials Science Forum*, pages 289–294. Trans Tech Publications Ltd, January 2012. 2

[19] H. Z. Hassan and N. M. Saeed. Advancements and applications of lightweight structures: A comprehensive review. *Discover Civil Engineering*, 1:47, August 2024. 6

[20] Joe D. Hoffman. *Numerical Methods for Engineers and Scientists*. Marcel Dekker, New York, NY, 2nd edition, 2001. 35

[21] Guang-Bin Huang, Qin-Yu Zhu, and Chee-Kheong Siew. Extreme learning machine: Theory and applications. *Neurocomputing*, 70(1–3):489–501, 2006. 106

[22] Frank P. Incropera, David P. DeWitt, Theodore L. Bergman, and Adrienne S. Lavine. *Fundamentals of Heat and Mass Transfer*. John Wiley & Sons, Hoboken, NJ, 7th edition, 2011. 15, 17, 23, 30, 52

[23] E. Kaschnitz, W. Funk, and T. Pabel. Electrical resistivity measured by millisecond pulse-heating in comparison to thermal conductivity of the aluminium alloy Al-7Si-0.3Mg at elevated temperature. *High Temperatures – High Pressures*, 43:175–191, January 2014. 49

[24] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, December 2014. 106

[25] Adem Korkmaz, Tuğçe Yağcı, and Osman Çulha. Optimization of T6 heat treatments for AlSi5Mgx alloys via computational materials engineering and experimental validation for automotive applications. *Journal of Materials Engineering and Performance*, 34:14157–14166, September 2024. 2

[26] Yura Malitsky and Konstantin Mishchenko. Adaptive gradient descent without descent. *arXiv preprint arXiv:1910.09529*, 2019. 75, 101

[27] MathWorks. Particleswarm. MATLAB R2025a Global Optimization Toolbox, 2025. 61

[28] Michael F Modest. *Radiative Heat Transfer*. Academic Press, 3rd edition, 2013. Copyright: Elsevier. 52

[29] Ali Mohammad-Djafari. Regularization, bayesian inference, and machine learning methods for inverse problems. *Entropy*, 23(12):1673, December 2021. 106

[30] Michael J. Moran, Howard N. Shapiro, Daisie D. Boettner, and Margaret B. Bailey. *Fundamentals of Engineering Thermodynamics*. John Wiley & Sons, Hoboken, NJ, 9th edition, 2018. 26

[31] Alexandru M. Morega, Marin Petre, Yelda Veli, and Alin A. Dobre. Flow and heat transfer in an aging heat furnace for aluminum alloy plates. *Revue Roumaine des Sciences Techniques – Série Électrotechnique et Énergétique*, 67(1):79–84, March 2022. 2

[32] F. Moukalled, L. Mangani, and M. Darwish. *The Finite Volume Method in Computational Fluid Dynamics: An Advanced Introduction with OpenFOAM and Matlab*, volume 113 of *Fluid Mechanics and Its Applications*. Springer International Publishing, 2016. 22

[33] Adrian P Mouritz. *Introduction to aerospace materials*. Woodland Publishing Limited, Cambridge, UK, 2012. 6

[34] National Institute of Standards and Technology. NIST chemistry webbook, 2025. Accessed May 2025. 27, 52

[35] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer, New York, NY, 2nd edition, 2006. 61

[36] Riccardo Poli. Analysis of publications on particle swarm optimization applications. *IEEE Transactions on Evolutionary Computation*, 15(6):892–902, 2010. 61

[37] Ian Polmear, David StJohn, J.F. Nie, and Ma Qian. *Light Alloys : Metallurgy of the Light Metals*. Butterworth-Heinemann, Oxford, UK, 5th edition, March 2017. 8

[38] S. Prathap Singh, D.X. Tittu George, A. Anitto Joe Xavier, and G. Abinicks Raja. Effect of heat treatment on the hardness behaviour of the aluminium 6061 alloy. *Materials Today: Proceedings*, 2023. 2, 7

[39] Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. 106

[40] Jose Immanuel Rajan. *Development of ultrafine grained A356 aluminum alloy by severe plastic deformation and studies on its deformation behavior and machinability*. PhD thesis, Indian Institute of Technology Bhilai, March 2017. 6

[41] Singiresu Rao. *Engineering Optimization: Theory and Practice*. John Wiley & Sons, Hoboken, NJ, 5th edition, October 2020. 18, 74

[42] ROCKWOOL Technical Insulation. Comfort: Basic theory. https://rti.rockwool.com/applications/marine-and-offshore/comfort/basic-theory/, 2025. Accessed: 2025-09-01. 53

[43] Andrea Saltelli, Marco Ratto, Terry Andres, Francesca Campolongo, Jessica Cariboni, Debora Gatelli, Michaela Saisana, and Stefano Tarantola. *Global Sensitivity Analysis: The Primer*. John Wiley & Sons, 1st edition, December 2007. 106

[44] Yuhui Shi and Russell C. Eberhart. A modified particle swarm optimizer. In *1998 IEEE International Conference on Evolutionary Computation*, pages 69–73. IEEE, 1998. 61

[45] N.M. Siddesh Kumar, Dhruthi, G.K. Pramod, P. Samrat, and M. Sadashiva. A critical review on heat treatment of aluminium alloys. *Materials Today: Proceedings*, 58:71–79, 2022. International Conference on Artificial Intelligence & Energy Systems. 2, 6

[46] Chandra Veer Singh and Derek H. Warner. An atomistic-based hierarchical multiscale examination of age hardening in an Al-Cu alloy. *Metallurgical and Materials Transactions A*, 44(6):2625–2644, February 2013. 4, 7, 8

[47] Pushpraj Singh, Raj Kumar Singh, and Anil Kumar Das. Optimization of heat treatment cycle for cast-Al6082 alloy to enhance the mechanical properties. *Engineering Research Express*, 6(1):015036, January 2024. 2

[48] Yao Sun. The use of aluminum alloys in structures: Review and outlook. *Structures*, 57:105290, November 2023. 2, 6

[49] Ferenc Szidarovszky and Sidney J. Yakowitz. *Principles and Procedures of Numerical Analysis*, volume 14 of *Mathematical Concepts and Methods in Science and Engineering*. Springer, New York, NY, 1978. 36

[50] Engineering Toolbox. Understanding convective heat transfer: Coefficients, formulas & examples. https://www.engineeringtoolbox.com/convective-heat-transfer-d_430.html, 2003. Accessed: 2025-05-20. 51

[51] Stephen R. Turns. *An Introduction to Combustion: Concepts and Applications*. McGraw-Hill, 2nd edition, 2000. 52

[52] Christian Vargel. *Corrosion of Aluminium*. Elsevier Science, Amsterdam, The Netherlands, 2nd edition, May 2020. 6

[53] H.K. Versteeg and W. Malalasekera. *An Introduction to Computational Fluid Dynamics: The Finite Volume Method*. Pearson Education Ltd, Harlow, England, 2nd edition, 2007. 31, 45

[54] Fei Wen, Rendong Ying, Peilin Liu, and Trieu-Kien Truong. Matrix completion via nonconvex regularization: Convergence of the proximal gradient algorithm. *arXiv preprint arXiv:1910.09529*, 2019. 75

[55] Philipp Wiechmann, Hannes Panwitt, Horst Heyer, Michael Reich, Manuela Sander, and Olaf Kessler. Combined calorimetry, thermo-mechanical analysis and tensile test on welded EN AW-6082 joints. *Materials*, 11(8):1396, August 2018. 27, 51

[56] Jiannan Yang, Yunqiang Ni, Hui Li, Xuewei Fang, and Bingheng Lu. Heat treatment optimization of 2219 aluminum alloy fabricated by wire-arc additive manufacturing. *Coatings*, 13(3):610, 2023. 2

[57] Magdalena Zawada-Michałowska and Paweł Pieśko. Post-machining deformations of thin-walled elements made of EN AW-2024 T351 aluminum alloy as regards the mechanical properties of the applied, rolled semi-finished products. *Materials*, 14(24):7591, 2021. 7

[58] Jun Zhu, Ning Chen, and Eric P. Xing. Bayesian inference with posterior regularization and applications to infinite latent svms. *Journal of Machine Learning Research*, 15:1799–1847, 2014. 106

[59] Yunus A. Çengel. *Heat and Mass Transfer: Fundamentals and Applications*. McGraw-Hill Education, New York, NY, 5th edition, 2014. 22

Appendix A: Parameter Table

Table 1: Calibratable parameters used in the furnace model, including optimized values, initial guesses, bounds, and units.

Parameter	Description	Optimized	Initial	Bounds	Units
Convective and Radiative Parameters					
$h_{f\to\text{alloy}}$	Conv. coeff. furnace → alloy	37.901	20	$[0, 50]$	$\mathrm{W\,m^{-2}\,K^{-1}}$
$h_{\text{ms}\to f}$	Conv. coeff. sheet → furnace	9.269	20	$[0, 50]$	$\mathrm{W\,m^{-2}\,K^{-1}}$
$h_{f\to w}$	Conv. coeff. furnace → walls	0.223	10	$[0, 50]$	$\mathrm{W\,m^{-2}\,K^{-1}}$
$h_{h\to f}$	Conv. coeff. heater → furnace	51.152	25	$[0, 100]$	$\mathrm{W\,m^{-2}\,K^{-1}}$
h_{out}	Heat loss to ambient	1.570	10	$[0, 50]$	$\mathrm{W\,m^{-2}\,K^{-1}}$
ϵ_{ms}	Emissivity of metal sheet	1	0.8	$[0.01, 1]$	–
ϵ_w	Emissivity of furnace walls	0.006	0.9	$[0.01, 1]$	–
Material Properties					
$C_{p,f}$	Furnace air heat capacity	–	1150	Fixed	$\mathrm{J\,kg^{-1}\,K^{-1}}$
$C_{p,\text{ms}}$	Metal sheet heat capacity	–	500	Fixed	$\mathrm{J\,kg^{-1}\,K^{-1}}$
$C_{p,w,1}$	Heat capacity of wall layer 1	–	900	Fixed	$\mathrm{J\,kg^{-1}\,K^{-1}}$
$C_{p,w,2}$	Heat capacity of wall layer 2	–	500	Fixed	$\mathrm{J\,kg^{-1}\,K^{-1}}$
$\lambda_{w,1}$	Thermal conductivity of wall 1	0.0611	0.05	$[0.01, 1]$	$\mathrm{W\,m^{-1}\,K^{-1}}$
$\lambda_{w,2}$	Thermal conductivity of wall 2	55.509	50	$[1, 100]$	$\mathrm{W\,m^{-1}\,K^{-1}}$
$\rho_{w,1}$	Density of wall layer 1	–	200	Fixed	$\mathrm{kg\,m^{-3}}$
$\rho_{w,2}$	Density of wall layer 2	–	7800	Fixed	$\mathrm{kg\,m^{-3}}$
m_{alloy}	Mass of alloy	–	1.2×10^{-5}	Fixed	kg
m_{ms}	Mass of metal sheet	–	33.2	Fixed	kg
m_h	Mass of heater	82.454	10	$[0.01, 100]$	kg
A_{alloy}	Surface area of alloy	–	2.9×10^{-5}	Fixed	$\mathrm{m^2}$
A_h	Heating surface area	–	1	Fixed	$\mathrm{m^2}$
Controller Parameters					
K_p	PID proportional gain	369.205	100	$[0, 1000]$	–
K_i	PID integral gain	0.601	1	$[0, 10]$	–
K_d	PID derivative gain	91.563	10	$[0, 100]$	–
n_{heaters}	Number of heaters	10	18	$[8, 20]$	–

Appendix B: Loss Landscape

Figure 1 displays a two-dimensional contour plot (left), and a three-dimensional surface plot (right) of the loss landscape under optimization scenario S3, respectively. The red-dotted line borders the feasible region, while the red asterisk highlights the lowest point identified in 900 loss function evaluations.

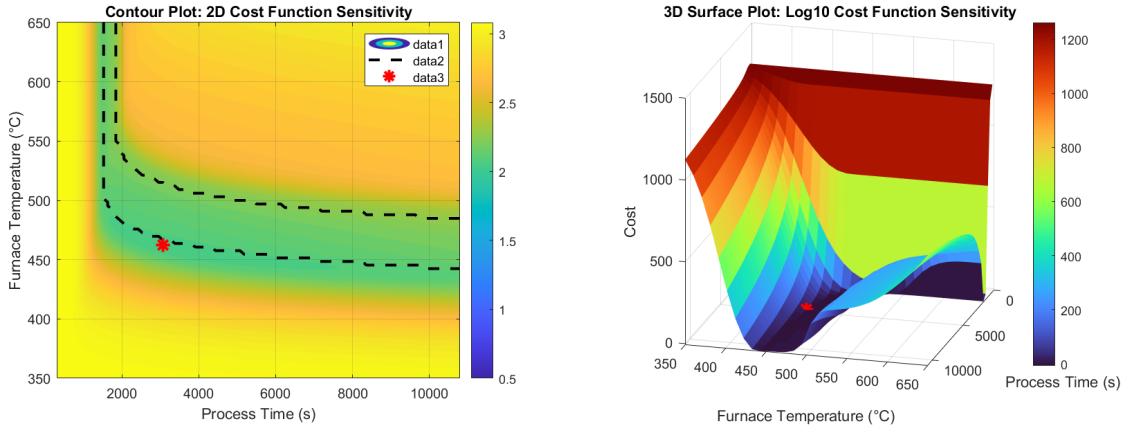

(a) 2D Contour plot

(b) 3D Surface plot

Figure 1: Loss Landscape Analysis

Appendix C: Code Availability

The full set of MATLAB scripts developed during this thesis, including the furnace model, calibration routines, and optimization tools, is available online at:

https://github.com/BramBastiaansenNL/Furnace-Model

This repository contains only the code authored by the thesis author. Certain components of the furnace simulation framework depend on internal AIT source code. These modules cannot be made public due to confidentiality restrictions and are therefore not included in the repository.