



CHALMERS
UNIVERSITY OF TECHNOLOGY



Simulation Analysis of Ceramic Shell Representation and Process Parameters in IN718 Investment Casting

Master's Thesis in Applied Mechanics

Joel Andersson

Department of Industrial and Materials Science

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

MASTER'S THESIS 2026

**Simulation Analysis of Ceramic Shell
Representation and Process Parameters
in IN718 Investment Casting**

Joel Andersson



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Industrial and Materials Science
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026

Simulation Analysis of Ceramic Shell
Representation and Process Parameters
in IN718 Investment Casting
Joel Andersson

© Joel Andersson, 2026.

Supervisor: Christer Persson, Department of Industrial and Materials Science
Supervisor: Liene Zaikovska, GKN Aerospace Sweden AB
Examiner: Christer Persson, Department of Industrial and Materials Science

Master's Thesis 2026
Department of Industrial and Materials Science
Chalmers University of Technology
SE-412 96 Gothenburg
Sweden
Telephone +46 31 772 1000

Cover:

Typeset in L^AT_EX
Gothenburg, Sweden 2026

Simulation Analysis of Ceramic Shell Representation and Process Parameters in IN718 Investment Casting

Joel Andersson

Department of Industrial and Materials Science

Chalmers University of Technology

Abstract

This thesis investigates how ceramic shell representation and selected process parameters influence the simulated solidification behaviour and predicted defect tendency of an investment-cast IN718 aerospace component. Casting simulations are used to compare different thermal conditions around the ceramic shell, with focus on temperature history, heat-flux response, filling and solidification behaviour, shell representation and porosity-related defect indicators. The study treats the ceramic shell as part of the thermal system, rather than only as a geometric boundary, since shell thickness and surrounding boundary conditions affect heat extraction, feeding conditions and last-to-freeze regions. The cast material is not varied as an independent parameter. IN718 is treated as the selected alloy system through which the thermal effects of shell representation and process parameters are interpreted.

The simulations show that insulation wrapping has a strong influence on the thermal response of the mould and casting system. Among the investigated process parameters, shell preheat temperature has the strongest influence on the evaluated outputs. Higher shell preheat gives preferable porosity indications, while lower shell preheat gives less favourable indications, especially when combined with lower mass flow. The mass-flow-rate variation has a smaller influence, while the enclosure emissivity has negligible influence for the investigated setup.

The comparison between constant and non-uniform shell representations shows that shell representation can affect the local thermal response and predicted defect tendency. However, this effect is smaller than the influence of shell preheat temperature. A limited qualitative comparison with available process observations shows some agreement between simulated defect-indicator regions and observed porosity indications, but the comparison is not sufficient for quantitative validation. The results should therefore be interpreted as simulation-based sensitivity trends that support process understanding, rather than as verified predictions of exact casting quality or as a final production specification.

Keywords: IN718, investment casting, ceramic shell, solidification, porosity, simulation.

Acknowledgements

I would like to thank my supervisor at GKN Aerospace, Liene Zaikovska, for her guidance and support throughout this project. I would also like to thank Martin Risberg at RISE and Ronan Cronin at Keysight Technologies for their assistance with questions related to the casting simulation software. I am grateful to the members of GAIC in Hallstahammar for their help with the shell measurement data and for providing validation material. Finally, I would like to thank Christer Persson, my supervisor and examiner at Chalmers, for his guidance and feedback during the work.

List of Acronyms

Below is the list of acronyms and abbreviations that have been used throughout this thesis, listed in alphabetical order:

AI	Artificial Intelligence
CAD	Computer-Aided Design
GAIC	GKN Aerospace Investment Casting
GKN	GKN Aerospace
GTC	Global Technology Centres
IN713C	Inconel 713C
IN718	Inconel 718
IN792	Inconel 792
SDAS	Secondary Dendrite Arm Spacing

Nomenclature

Below is the nomenclature of indices, sets, parameters, and variables that have been used throughout this thesis.

Indices

i Index for ceramic shell layer

Sets and domains

Ω Selected casting region used when evaluating local minimum Niyama values

Parameters

A Surface area used in the casting modulus V/A

B Mould/process constant in the Chvorinov-type solidification relation

C_λ Alloy- and process-dependent coefficient in the SDAS-cooling-rate relation

c_p Specific heat capacity

F Radiation view factor

$f_{s,\text{crit}}$ Critical solid-fraction threshold used to identify the late mushy region

$g_{l,\text{cr}}$ Critical liquid fraction at pore formation

h_{m-s} Metal-shell interfacial heat-transfer coefficient

$h_{s-\infty}$ Effective heat-transfer coefficient from the shell outer surface to the surroundings

h_{rad} Equivalent radiative heat-transfer coefficient

k Thermal conductivity

k_i	Thermal conductivity of shell layer i
k_s	Effective ceramic shell thermal conductivity
$k_{s,eq}$	Equivalent through-thickness thermal conductivity of a layered shell
L	Latent heat of solidification
L_i	Thickness of shell layer i
L_s	Ceramic shell thickness
$L_{s,0}$	Baseline ceramic shell thickness
L_{tot}	Total ceramic shell thickness
n	Empirical exponent in the SDAS–cooling-rate relation or Chvorinov-type relation
N	Number of ceramic shell layers
ΔP_{cr}	Critical pressure drop for pore formation
R''_{tot}	Total thermal resistance per unit area
T_0	Reference temperature for enthalpy calculation
T_L	Liquidus temperature
T_S	Solidus temperature
T_∞	Ambient or far-field surrounding temperature
T_{sur}	Surrounding surface or ambient temperature used in radiation calculation
β	Solidification shrinkage factor
β'	Effective shrinkage factor used in the pore-volume relation
ε	Emissivity
μ_l	Liquid viscosity
ρ_l	Liquid density
ρ_s	Solid density
σ	Stefan–Boltzmann constant

Variables

ΔT_f	Freezing range
f_l	Liquid fraction
f_s	Solid fraction
G	Thermal gradient
g_p	Shrinkage pore volume fraction

H	Specific enthalpy
K	Mush permeability
Ny	Niyama criterion value
Ny^*	Niyama star microporosity indicator
$Ny_{\min,\Omega}$	Minimum Niyama value in the selected casting region Ω
p	Pressure
q''	Heat flux per unit area
q''_{rad}	Radiative heat flux per unit area
$\dot{q}$	Internal or numerical source term in the energy equation
S_{y,L_s}	Normalised sensitivity of output y to shell thickness L_s
t	Time
t_s	Local solidification time
T	Temperature
T_m	Metal temperature
T_s	Shell surface temperature
$\dot{T}_c$	Cooling rate
$\mathbf{u}_l$	Intrinsic interdendritic liquid velocity
V	Volume used in the casting modulus V/A
y	Generic simulation output
y_0	Baseline value of simulation output y
λ_2	Secondary dendrite arm spacing
ρ	Density

Contents

List of Acronyms	ix
Nomenclature	xi
List of Figures	xvii
List of Tables	xix
1 Introduction	1
1.1 Background	1
1.2 Purpose	4
1.3 Goals	6
1.4 Limitations and demarcations	7
1.5 Ethical and sustainability aspects	8
2 Literature Review	11
2.1 Purpose and scope of the literature review	11
2.2 Investment casting and aerospace quality requirements	12
2.3 Ceramic shell systems and thickness variation	12
2.4 Shell thickness, heat extraction and material response	13
2.5 Cast material effects and IN718 solidification	15
2.6 Solidification defects and simulation indicators	16
2.7 Simulation workflows and validation practices	17
2.8 Research gap and implications for the thesis	18
2.9 Chapter conclusion	19
3 Theory	21
3.1 Purpose and scope of the theory chapter	21
3.2 The ceramic shell system	22
3.3 IN718 as the cast material	22
3.4 Fraction solid and the meaning of the 0.7 cutoff	23
3.5 Heat-transfer theory for investment casting	25
3.6 Cooling rate, thermal gradient and solidification time	26
3.7 Shell thickness as a thermal and mechanical variable	27
3.8 Feeding theory and shrinkage porosity	28
3.9 Niyama criterion	30
3.10 Niyama star microporosity	30

3.11	Relation between shell thickness and Niyama star microporosity . . .	32
3.12	Simulation assumptions and credibility checks	33
3.13	Implications for the thesis simulation plan	33
3.14	Traceability of formulas and symbols	35
3.15	Summary of key theoretical quantities	39
4	Methods	41
4.1	Approach	41
4.2	Simulation software	41
4.3	CAD model and casting system	42
4.4	Ceramic shell modelling	42
4.5	Insulation wrapping	43
4.6	Simulation case structure	44
4.7	Filling and solidification simulation	44
4.8	Post-processing and evaluated outputs	45
4.9	Comparison between simulation cases	46
4.10	Validation and comparison with process observations	46
4.11	Confidentiality and reporting of parameters	47
5	Results & Discussion	49
5.1	Presentation of results	49
5.2	Overview of simulation cases	50
5.3	Overall trends from the simulation cases	51
5.4	Influence of insulation wrapping	53
5.5	Influence of shell preheat temperature	55
5.6	Influence of mass flow rate	57
5.7	Influence of enclosure emissivity	59
5.8	Influence of shell representation	60
5.9	Temperature–heat-flux relationship	61
5.10	Defect-indicator results	62
5.11	Comparison with observed porosity indications	67
5.12	Summary of main result trends	67
6	Conclusion	69
6.1	Recommendations for future work	72
6.2	Final words	73
	Bibliography	75

List of Figures

1.1	Examples of ceramic-shell manufacture in investment casting: (a) wax pattern being dipped in ceramic slurry, and (b) schematic build-up of the ceramic shell by repeated coating layers [8, 15].	1
1.2	Example of shell-thickness variation around a complex investment-casting geometry [8].	3
5.1	Temperature histories for the initial wrapped and unwrapped model comparison at Location A. The wrapped cases show that insulation wrapping strongly affects the thermal response of the mould and casting system. Numerical axis values are omitted due to confidentiality.	54
5.2	Heat-flux histories for the initial wrapped and unwrapped model comparison at Location A. The difference between wrapped and unwrapped cases is larger than the difference between the two wrapped shell representations. Numerical axis values are omitted due to confidentiality.	55
5.3	Temperature histories at Location A for the non-uniform wrapped shell cases. The comparison shows the influence of shell preheat temperature at both standard and lower measured-average mass-flow conditions. Numerical axis values are omitted due to confidentiality. .	56
5.4	Heat-flux histories at Location A for the non-uniform wrapped shell cases. The heat-flux response supports the interpretation that shell preheat temperature was the dominant investigated process parameter. Numerical axis values are omitted due to confidentiality.	57
5.5	Temperature histories at Location B. The comparison between Tests 2 and 5 shows the influence of reducing mass flow at lower shell preheat temperature, while the comparison between Tests 3 and 6 shows the corresponding effect at higher shell preheat temperature. Numerical axis values are omitted due to confidentiality.	59
5.6	Comparison between constant and non-uniform wrapped shell representations at Location A. The shell representation affects the local thermal response, but the difference is smaller than the influence of shell preheat temperature. Numerical axis values are omitted due to confidentiality.	61

5.7	Representative temperature–heat-flux plots for selected non-uniform wrapped shell cases at Location A. These plots compare local heat extraction at similar thermal states and support the trends observed in the time-based plots. Numerical axis values are omitted due to confidentiality.	62
5.8	Niyama star microporosity indicator for the initial wrapped and unwrapped model comparison. The wrapped cases give a more favourable indication pattern than the corresponding unwrapped cases. The numerical colour scale has been removed due to confidentiality.	64
5.9	Niyama star microporosity indicator for Tests 2 and 3. Each row compares the constant-shell and non-uniform-shell representations for the same process condition. The lower shell-temperature case shows stronger indications than the higher shell-temperature case. The numerical colour scale has been removed due to confidentiality.	65
5.10	Niyama star microporosity indicator for Tests 5 and 6. Each row compares the constant-shell and non-uniform-shell representations for the same process condition at the lower measured-average mass-flow condition. The numerical colour scale has been removed due to confidentiality.	66

List of Tables

3.1	Traceability of the main formulas used in the theory chapter.	35
3.2	Key theoretical quantities for interpreting shell-thickness and IN718 casting simulations.	39
5.1	Anonymised monitoring locations used for representative thermal-history plots.	49
5.2	Overview of the non-uniform wrapped shell sensitivity cases. Exact numerical values are omitted due to confidentiality.	50
5.3	Qualitative summary of the main simulation trends.	52

1

Introduction

1.1 Background

Investment casting is a manufacturing process used to produce components with high dimensional accuracy, good surface finish and complex geometries that would be difficult or costly to manufacture using conventional machining or forming methods. The process is based on producing a disposable wax pattern, which defines the geometry of the final component. For many smaller components, such as turbine blades, wax patterns are assembled into a casting tree before being repeatedly dipped into ceramic slurry, stuccoed and dried until a ceramic shell has been formed. The Hub-sector studied in this work, a structural segment of the inner, hub-side casing of a Turbine Rear Frame (TRF), is not cast using a casting tree. The wax is then removed, leaving a cavity inside the shell. After pattern removal, the shell is fired and preheated before molten alloy is poured into the cavity. Once the metal has solidified and cooled, the ceramic shell is removed and the casting is cleaned, heat treated if required, inspected and prepared for use [12].

The repeated slurry-dipping and layer-building steps are illustrated in Figure 1.1. In practice, the wax pattern is first coated by slurry and then successive stucco and slurry layers are built up until the required shell thickness and strength are obtained.




Figure 1.1: Examples of ceramic-shell manufacture in investment casting: (a) wax pattern being dipped in ceramic slurry, and (b) schematic build-up of the ceramic shell by repeated coating layers [8, 15].

In aerospace applications, investment casting is particularly important because many components combine complex geometry with strict quality requirements. Components may contain thin walls, internal passages, section transitions, local mass

concentrations and other features that make filling and solidification difficult to control. At the same time, aerospace castings are often exposed to harsh mechanical, thermal and environmental conditions during service. Defects such as shrinkage porosity, misruns, cold shuts, inclusions, geometry deviations, shell-related surface defects or unfavourable microstructure can therefore lead to rejection, expensive rework or reduced component reliability.

This work is carried out in connection with GKN Aerospace, where investment casting is being introduced as a manufacturing capability at the GTC Trollhättan site in collaboration with the existing facility in Hallstahammar. The industrial motivation is to improve process understanding, support production development and strengthen the ability to manufacture advanced aerospace components through simulation-based methods. Physical trial casting is expensive, time-consuming and dependent on available material, tooling, production capacity and inspection resources. Simulation therefore offers an efficient way to study how selected process variables influence the casting before they are applied in production.

A central part of the investment casting process is the ceramic shell. It defines the geometry, provides mechanical support during handling and pouring, resists thermal shock and chemical attack, allows gas transport through itself and is one of the control mechanisms for the heat extraction rate from the molten metal. The shell therefore not only determines the shape of the component. It is also part of the thermal system that affects how the casting fills, cools and solidifies.

In practice the ceramic shell thickness is difficult to control. The final thickness depends on the number of coats, slurry pick-up, drainage, stucco size, drying conditions, pattern orientation, local geometry and process control. Local features such as corners, concave surfaces, thin passages and curved regions may lead to uneven coating. As a result, a nominal shell thickness may not fully represent the real shell surrounding a casting. This is important because local thickness can influence local heat extraction, especially in thin-walled or geometrically complex components.

An example of this type of shell-thickness variation is shown in Figure 1.2, where the local shell thickness changes around different regions of the same geometry.




Figure 1.2: Example of shell-thickness variation around a complex investment-casting geometry [8].

Shell thickness affects casting behaviour mainly through its influence on heat transfer. A thinner shell usually reduces thermal resistance and can increase cooling intensity, while a thicker shell can retain heat and slow down cooling. However, the effect is not universal. Faster cooling may reduce local hot spots in one region, but may also freeze certain feeding paths too early. A thicker shell may keep thin sections open for longer and improve feeding in some cases, but it may also delay solidification and increase the duration of local hot spots. The result depends on component geometry, shell properties, mould preheat, interface heat transfer, external heat exchange, gating conditions, feeding conditions and material.

The selected cast material is also important when interpreting shell-thickness effects. Nickel-based superalloys such as IN718 are used in aerospace applications because of their high-temperature strength, corrosion resistance and mechanical performance. At the same time, IN718 has a complex solidification behaviour, including a temperature interval where liquid and solid coexist. This means that changes in cooling history may influence feeding conditions, last-to-freeze regions and the risk of shrinkage defects. In this paper, the material properties are not experimentally measured or varied as an independent study. They are treated as input assumptions from available software data, literature or company-defined sources. Its role is to explain why the simulated response to shell-thickness variation must be interpreted through the behaviour of the selected alloy.

Casting simulation makes it possible to study these relationships in a controlled and consistent way. ProCAST [14], the commercial casting simulation software used in this study, can calculate filling, heat transfer, solidification behaviour and quality indicators. Outputs such as temperature evolution, cooling rate, thermal gradient, solidification time, fraction-solid development, last-to-freeze regions and porosity indicators can be compared between simulation cases. However, simulation

results are not direct measurements of reality. They depend on geometry, mesh quality, material databases, shell properties, heat-transfer coefficients, radiation and convection assumptions, boundary conditions, porosity models and post-processing criteria. A simulation study should as such document its assumptions and distinguish between sensitivity analysis and validation.

Sensitivity analysis shows how the casting response changes when selected input parameters are varied. Validation meanwhile shows whether the model represents the real casting process with any real accuracy. Validation against material data, temperature measurements, inspection results or trial castings is valuable if available, but it is not treated as a requirement. Even without validation a study can provide some understanding of the process by identifying critical regions, comparing trends and showing which outputs are most affected by a variation in the shell thickness.

The problem addressed in this thesis is not about whether a ceramic shell should be made thicker or thinner because no single correct answer exists and varies between cases. The purpose is to understand how shell-thickness variation affects the predicted solidification behaviour and material quality of the studied casting and how this effect should be interpreted. The work is meant to contribute towards the development of casting simulation procedures by connecting process simulation with the physical mechanisms that control casting quality.

1.2 Purpose

The purpose is to investigate how variations in ceramic shell thickness influence the simulated solidification behaviour and predicted material quality of an investment-cast aerospace component. The aim is not to compare several shell-thickness values, but to understand why the casting response changes or why it does not when it's varied.

Casting simulation is utilised as the investigative method. A baseline simulation model is created using available CAD geometry, process parameters, material data and shell assumptions. From this baseline, a study is performed where shell thickness is varied. The resulting simulation cases are compared using outputs such as temperature evolution, cooling rate, thermal gradient, solidification time, fraction-solid development, last-to-freeze regions, Niyama criterion, Niyama star microporosity or other porosity indicators.

An additional purpose is to connect the simulation results to the behaviour of the selected cast material. Since IN718 has a freezing range and is sensitive to late-stage feeding conditions, shell-thickness effects must be interpreted through the material response rather than only through how it shapes the final cast part. This means that it should be explained how changes in heat extraction influence the solidification sequence, feeding paths, hot spots and shrinkage-porosity risk.

The industrial purpose is to support process development for investment casting at

GKN Aerospace. By identifying how shell-thickness variation influences predicted casting performance, the hope is that the work can support future decisions regarding the company's casting process. The expected contribution is therefore improved process understanding rather than a final specification for shell thickness.

1.3 Goals

The primary goal is to evaluate the influence of ceramic shell-thickness variation on the simulated solidification behaviour and predicted defect risk of a selected investment-cast aerospace component under different process variables.

To achieve this the following sub-goals are defined:

1. Establish a theoretical and literature-based foundation for the study by describing investment casting, ceramic shell behaviour, IN718 solidification, heat transfer, feeding, shrinkage porosity and relevant simulation indicators.
2. Create a baseline casting simulation model using the available CAD geometry, process parameters, shell assumptions and cast material data.
3. Identify the shell-thickness values or shell-thickness variations that should be utilised.
4. Perform a study where shell thickness is varied while other parameters are kept controlled.
5. Compare the simulation cases using relevant thermal and quality-related outputs.
6. Perform similar studies, now varying one process parameter at a time to see what, if any, results change, both between cases and shell thickness.
7. Identify which regions of the component are most sensitive to shell-thickness variation.
8. If suitable process or material data from trial castings are available, compare simulation results with these to evaluate whether the simulated trends are physically reasonable.
9. Formulate conclusions that support the development of robust investment-casting procedures for their aerospace components.

The intended result is a documented simulation-based assessment of shell-thickness sensitivity. It should be attempted to both state which simulation case gives higher or lower software output values, but also try to explain why. For example, if a thicker shell were to reduce the cooling rate, but also lower the thermal gradient then the effects on potential porosity indicators should be considered carefully. If a thinner shell reduces a hot spot but in turn freezes a feeding path too early, then the improvement in solidification time need not necessarily correspond to improved casting quality. The goal is therefore to interpret the results as a connected thermal and flow problem.

1.4 Limitations and demarcations

This study is limited to a simulation-based investigation of investment casting. The primary focus is the effect of ceramic shell-thickness variation on simulated solidification behaviour and defect indicators. The work does not aim to develop a new casting solver, a new porosity model, a new material model or at all study the material properties of the cast component. Commercial simulation software is used as the computational tool, and the focus is on model setup, shell-thickness variation and interpretation of results.

The study is limited to the selected component geometry and the process conditions made available for the project. The results should therefore not be interpreted as a general recommendation for investment casting. Shell-thickness effects are heavily dependent on component geometry, shell material, preheat condition, boundary conditions, gating system, feeding system and simulation assumptions. A shell thickness that might work for the specific component and system in this study might be wholly unsuitable for another.

The cast alloy is not investigated as a variable. No experimental material-property measurements are performed, and no comparison between different cast materials is included. The material definition used in the simulation is treated as a fixed input required to run the model, not as an object of study. Consequently, the thesis does not evaluate the influence of liquidus temperature, solidus temperature, latent heat, density, thermal conductivity, heat capacity, viscosity or fraction-solid curve as separate parameters. These quantities are only relevant as they are part of the predefined simulation setup.

The ceramic shell is the main investigated variable. The study focuses on how changes in shell thickness influence heat extraction, solidification sequence, hot-spot development, feeding behaviour and predicted shrinkage-related indicators. Detailed shell manufacturing processes such as slurry rheology, stucco distribution, drying behaviour, binder chemistry, shell permeability, shell cracking during dewaxing and local coat-by-coat formation are outside the scope. If the simulation model uses a homogenised shell material, the shell is treated as a representation rather than as a fully accurate multilayer ceramic system.

The simulation results are limited by the accuracy of the boundary conditions and modelling assumptions. Metal-shell interface heat transfer, shell preheat condition, radiation, convection, external cooling and contact conditions can strongly affect predicted cooling rate, thermal gradient and solidification time. If these parameters are not calibrated against measured data, the results should primarily be interpreted as comparative sensitivity trends rather than correct predictions of casting quality.

The study focuses mainly on solidification-related simulation outputs. Shrinkage porosity, hot spots, feeding conditions, last-to-freeze regions, Niyama criterion and Niyama star microporosity indicators are central to the analysis. Other defects such

as gas porosity, oxide films, inclusions, shell cracking, metal penetration, surface reactions, dimensional distortion, residual stress and hot tearing may be discussed if they are ever relevant. However they are not part of the main modelling targets.

Filling-related behaviour is included only to the extent necessary for the selected simulation workflow. Complete filling simulation is performed and the results may be used to identify possible filling-related tendencies. However, the main emphasis is shell-thickness influence on heat transfer, solidification and shrinkage-related indicators rather than fluid-flow optimisation. The gating and feeding will not be redesigned, instead they are used as part of the simulation throughput.

Experimental validation is included only if suitable trial-casting data, inspection results or temperature measurements are available within the project time frame. Validation is therefore considered desirable but not mandatory. If no direct validation data are available, the credibility of the simulation is assessed through documentation of assumptions, comparison with expected physical behaviour, consistency between outputs and support from literature. In that case, the conclusions are limited to simulation-based sensitivity and should not be presented as verified production behaviour.

The project is also limited by the available time. The work is planned as a 30-credit master's thesis over approximately 20 weeks. This limits the number of simulation cases, the extent of parameter calibration, the possibility of repeated mesh studies and the depth of experimental comparison. For this reason, the simulation matrix should be focused on the most important quality indicators.

Finally it is not meant to define an approved production shell specification. A production specification would require broader process qualification, repeated trials, statistical studies, inspection evidence and approval within the relevant industrial quality system. The intended contribution is instead to provide a simulation workflow, identify shell-thickness sensitivities and support further development of investment-casting procedures.

1.5 Ethical and sustainability aspects

Due to the project being carried out in an aerospace manufacturing context this means ethical and sustainability aspects are present even though they are not the main subject of the work. The project itself is a technical simulation study focused on ceramic shell-thickness variation in investment casting. It does not include a separate ethical analysis, environmental assessment or life-cycle study. However, the industrial context affects how the work should be understood.

One ethical aspect is the safety-critical nature of aviation. Aerospace components are manufactured under strict quality requirements because defects, process variation or incorrect technical conclusions can have consequences for reliability and service life. It is not meant to qualify a component or define an approved production route. The

simulation results are therefore treated as support for process understanding rather than as direct evidence of component acceptability.

Additionally, the project is carried out at a company with activities in both civil and defence-related aerospace. GKN Aerospace has direct involvement in defence-sector manufacturing, which makes the industrial context of the thesis ethically relevant even though the specific component and simulations in this work are not presented as a defence application. Improvements in manufacturing simulation, process understanding and casting may contribute to general aerospace production capability, and such capability can be relevant to both civil and military supply chains. Therefore it cannot be considered entirely separated from the defence context in which the company operates. At the same time, the work itself is limited to investment-casting process simulation, shell-thickness variation and solidification-related quality indicators. It does not address weapon systems, operational use, military performance requirements or defence strategy the context should however still be kept in mind.

Sustainability is also relevant due to its connection to the aviation industry, a sector closely connected to fuel consumption, emissions, resource use and long-term environmental impact. Although this thesis does not investigate aircraft operation, engine efficiency or environmental performance directly, the manufacturing of aerospace components is part of the wider aviation value chain. Process development in this sector therefore exists within a broader context where improved reliability, reduced scrap, efficient material use and manufacturing methods are connected to sustainability concerns.

The sustainability aspect is mainly contextual. The work does not include a life-cycle assessment, carbon-footprint calculation or comparison between different manufacturing routes. It also does not quantify how shell-thickness variation affects environmental impact. However, simulation-based process development can still be related to sustainability because it may support better process understanding before physical trials are performed, reduce unnecessary scrap or rework, and contribute to more controlled manufacturing of aerospace components. These aspects are acknowledged, but they are not evaluated further.

Finally, artificial intelligence tools were used to support the writing process of this report. Their use was limited to improving language consistency, structure, and readability. The workflow, simulation setup, interpretation of results, and final content remain the work of the author.

Overall, the ethical and sustainability aspects of this thesis are mainly connected to its industrial setting: aviation safety, defence-adjacent manufacturing, confidentiality, responsible interpretation of simulation results, resource use in casting development and transparent use of AI as a writing aid. They are acknowledged as a part of the project, but they are not investigated further.

2

Literature Review

2.1 Purpose and scope of the literature review

The purpose of this literature review is to study the published work needed for the problem. How ceramic shell characteristics, cast material response and simulation practice interact when predicting solidification behaviour and defects in complex aerospace components.

A recurring theme in the literature is that shell thickness is not merely a geometric property of the mould. It affects how heat transfers out of the liquid metal, how feeding paths develop during solidification, how long regions remain vulnerable to shrinkage porosity and how stresses appear in the shell and casting. Simultaneously the cast alloy controls the response to a given shell condition through its freezing range, thermophysical properties, shrinkage behaviour, flow behaviour and solidification morphology. For these reasons a shell thickness that is suitable for one alloy, a specific geometry or a gating arrangement can therefore be not be suited for another.

For these reasons, the review is focused on five areas. The investment casting process and its aerospace quality demands. Ceramic shell structure and thickness variation. The influence of shell thickness on heat extraction and material response. Solidification defects and simulation indicators. Simulation workflows for sensitivity analysis and validation. Meanwhile detailed mathematical definitions, formulas and numerical post-processing criteria are instead left for the theory and methodology chapters.

The sources used in this review have deliberately been mixed. Peer-reviewed articles and validated modelling studies give the strongest basis for mechanism and method selection. Industrial or software sources are primarily used to describe practical simulation outputs or foundry terminology. This difference is important because the thesis should not treat all references as equally strong. Studies that includes thermal calibration, trial-casting comparison or measured defect validation are naturally more likely to contain correct information than sources that only use software functionalities or a single unvalidated simulation.

2.2 Investment casting and aerospace quality requirements

Aerospace investment castings commonly involve nickel-based superalloys, stainless steels, titanium alloys and other high-performance materials. These castings are often manufactured for parts with demanding service environments and harsh acceptance criteria. Shrinkage, cold shuts, shell inclusions, dimensional distortion or uncontrolled microstructure can lead to scrapping or costly reworking. Thus investment casting is not only treated as a shape-making process but as a process where tooling, shell manufacture, alloy behaviour, heat transfer and inspection requirements must all be controlled.

In their article "An overview of ceramic moulds for investment casting of nickel superalloys" Kanyo et al. [15] describes ceramic moulds as central to the near-net-shape casting of nickel superalloys. This is due to their great resistance to heat and chemical compatibility at high temperature. Their review also describes practical limitations of ceramic shells, including thermal shock resistance, permeability, dimensional stability, strength and leachability. These limitations are directly relevant to casting within our applications because the shell must survive dewaxing, firing, preheating, metal pressure and cooling without cracking, changing the shape of the cavity or reacting with the alloy.

For thin-walled or complex components, solidification control is especially difficult. Thin sections can freeze early and cut off feeding to other regions. On the other hand insulated or larger regions can remain hot for too long and become isolated hot spots. Rakoczy and Cygan [20] showed that temperature distribution in shell moulds during thin-wall superalloy casting will influence features in both its macrostructure and microstructure. This gives credence to the view that the shell is part of the metallurgical system rather than a passive container. Shell thickness variation should therefore be treated as a material-quality issue as well as a mould-design issue

The aerospace context also changes how simulation results should be analysed. A case that seems adequate from a thermal perspective may still be unacceptable if it predicts localised shrinkage, hot spots, large thermal gradients or unfavourable solidification in critical regions. Therefore it is useful to connect predicted shell-thickness effects to quality requirements, not simply to contour plots. This means looking at outputs that can be related to rejection mechanisms, such as last-to-freeze regions, porosity indicators, solidification time, local cooling rate and feeding isolation.

2.3 Ceramic shell systems and thickness variation

Ceramic investment shells usually consist of a primary coat and several backup coats. The primary coat controls the surface in contact with the metal and influences surface finish, chemical interaction and local heat transfer. The backup coats build mechanical strength and thickness. Final shell thickness depends on the number of

coats, slurry pick-up, stucco size, drainage, drying conditions, local geometry, pattern orientation and process control. These factors mean that its practically impossible to achieve a perfectly uniform shell. This is especially true around features with complex curvature or restricted drainage.

It must also satisfy several requirements. The shell must be strong enough for handling and pouring, permeable enough to release gases, refractory enough to withstand metal temperature and stable enough to preserve component geometry. Increasing shell thickness can improve robustness, but it can also increase shell weight, process time, heat capacity and insulation. Jones and Yuan [13] showed that the performance of the ceramic shell depends on the combined effects of slurry, binder, stucco, processing route and the mechanical and physical properties which came thereof. Thus for this work this means that nominal thickness should not be varied while simultaneously ignoring the isolation from shell material properties.

Rui-zhe Dong et al. [8] investigated ceramic shell thickness uniformity in turbine blade investment casting and connected non-uniform shell build-up with precision outcomes. Their work is relevant because it links a measurable shell-manufacturing variable to the later casting accuracy. For a simulation workflow, it suggests that global shell-thickness values are a useful initial step however it may not entirely represent the full industrial process. If measured shell thickness maps become available, local thickness variation could provide a more realistic basis for later simulation cases.

The literature also shows that shell thickness does not have a universal effect. A thicker shell usually behaves as a more insulating mould, but the real effect of it depends on preheat temperature, shell composition, radiation, convection, interface behaviour, cluster arrangement and local geometry. For example Flow Science [10] used parametric simulations to show that thickness changes can alter shrinkage behaviour in ways that depend on which casting system was used. This means that a sensitivity study should evaluate a realistic interval of thickness values rather than assume in advance that thicker or thinner is better.

With these things in mind the shell thickness should therefore be looked at from two angles. The first is the idealised modelling, where thickness is varied globally to identify the general thermal direction of the response. The second level is from manufacturing-realism where local thickening or thinning is considered in areas where coating build-up, drainage or restricted access could create non-uniform shell geometry. The first is easier to implement. The second level on the other hand is more representative of the industrial problem.

2.4 Shell thickness, heat extraction and material response

Several studies show that shell thickness affects casting quality. Primarily through its influence on thermal history. Shell thickness will change the rate at which heat

leaves the metal, the temperature distribution through the shell and the duration in which certain regions solidifies. This affects feeding, hot-spot formation, porosity risk and the development of the microstructure. However the same thickness change can have different consequences depending on alloy, geometry and boundary conditions.

Alvandi et al. [2] studied IN792 turbine blade investment casting using ProCAST and compared simulations with castings. They reported that a shell that was too thin increased cooling intensity and hindered feeding. While a shell that was too thick increased the duration of local solidification and led to an increase in distributed shrinkage. In their blade geometry, an intermediate shell thickness gave the best results. This study is useful for the present thesis because it shows why shell thickness should be treated as an optimisation and robustness problem rather than as a simple variable. At the same time, it should be read as a specific result for their case. Not as evidence for a universally ideal thickness.

Raza et al. [21] looked at the combined effects of shell thickness, insulation and casting temperature in thin-wall nickel-base turbine blade casting. Their results showed that increased shell thickness could improve feeding distance in thin sections. While insulation and the metal-mould thermal condition impacted shrinkage resistance and how it managed to fill the cavities. This is important because it shows that shell thickness interacts with other process controls. A thickness change may either harm or benefit the process depending on the boundary conditions.

Szeliga [24] investigated different types of insulation strategies for thin-walled nickel-based superalloy castings and showed that changing the thermal boundary around the ceramic shell had an impact on the cooling rate, how the shrinkage behaved and the distribution of porosity. Although the work used IN713C rather than IN718, the underlying properties are relevant. External insulation and shell thermal behaviour influence thermal gradients, feeding conditions and shrinkage risk. Together these three studies supports a design-of-experiments approach rather than a single comparison between one thin and one thick shell case.

The articles are consistent in one area but differs in another. They all agree that shell thickness matters because it changes thermal history and feeding conditions. They differ in the direction and size of the effect because each study uses a particular alloy, geometry, process condition and boundary assumption. This is important for the project at hand. Shell-thickness sensitivity must be established for the selected casting system instead of being imported directly from published turbine-blade studies.

The literature therefore supports treating shell thickness as one part of a coupled thermal system. For simulation purposes, shell thickness should be considered together with shell conductivity, shell heat capacity, preheat condition, interface heat transfer, external heat exchange and component geometry. Not varied along with other parameters. This is important for the thesis because the stated objective is to identify parameters influencing casting performance, not simply to generate isolated

contour plots for a few thickness values.

2.5 Cast material effects and IN718 solidification

The cast material determines how strongly a given shell condition influences quality. Nickel-based superalloys generally have complex solidification behaviour, strong segregation tendencies and strict microstructural requirements. Their response to shell thickness is therefore not the same as that of aluminium alloys, steels or simpler casting alloys. A material with a broad freezing interval and significant interdendritic segregation can remain vulnerable to issues with feeding and microporosity over a wider temperature range than an alloy with a smaller freezing range.

IN718 is a precipitation-strengthened Ni–Fe–Cr alloy with important Nb, Mo, Ti and Al additions. The challenges with its castability are connected to segregation during solidification. Mitchell and Wang [17] describes IN718 as a well-studied superalloy in relation to solidification and precipitation, making it suitable for simulation analysis. At the same time its behaviour means that material data in commercial software should not be used unless identical to the one from the industrial application. The predicted freezing interval, liquid fraction behaviour and segregation sensitive outputs should be checked against literature or available process data where possible.

Shi et al. [22] measured the slow-cooling solidification behaviour of IN718 as-cast and described a solidification path including primary gamma formation, carbide formation and terminal Laves phase formation. Their work shows the persistence of interdendritic liquid and the importance of Nb-rich segregation in the final stages of solidification. For a shell-thickness study, this matters because thickness changes the thermal path through which late-stage liquid feeding and segregation-sensitive phases develop.

From this the literature on IN718 supports including material specific discussion in the thesis. A literature review can explain why IN718 is sensitive to cooling history, feeding and segregation. The detailed treatment of solid-fraction thresholds, permeability arguments, liquid-fraction definitions and post-processing cut-offs belongs more naturally in the theory or methodology chapter, where those criteria can be defined, justified and linked to the selected simulation outputs.

Here the key point is that shell thickness affects material quality indirectly. It has an impact on cooling rate, thermal gradient, solidification time and the connectivity of remaining liquid. The final effect depends on the alloy's freezing behaviour, shrinkage behaviour and microstructural sensitivity. Thus presenting shell thickness as an isolated mould variable is incorrect and should instead be connected to cast material response. This is in turn consistent with the thesis title and project brief.

2.6 Solidification defects and simulation indicators

The main defect in a shell-thickness sensitivity study is shrinkage porosity. Shrinkage forms when solidification contraction is not offset by liquid feeding. In complex castings, shrinkage may appear in isolated hot spots, section transitions, large regions, centreline locations or areas where early solidification hinders feeding. The risk is influenced by thermal gradients, cooling rate, solidification time, mushy-zone feeding, gas content and geometry.

Casting simulation packages often report the Niyama criterion or related microporosity indicators. In their study “Use of the Niyama Criterion to Predict Shrinkage-Related Leaks in High-Nickel Steel and Nickel-Based Alloy Castings.” [6] Carlson and Beckermann describe the Niyama criterion as a thermal indicator for shrinkage-related defects and Altair Engineering [1] presents it as a practical foundry result for identifying solidification shrinkage risk. The important part here is not the mathematical form of the criterion but how it is interpreted. It should be used as a risk indicator where the critical value depends on alloy, process and inspection requirements.

In “Prediction of Shrinkage Pore Volume Fraction,” [5] Carlson and Beckermann developed a dimensionless form of the Niyama approach to account for alloy properties and solidification behaviour. Guo et al. [11] utilised a dimensionless Niyama model with a ProCAST porosity model for nickel-based superalloy castings and validated against experimental microporosity measurements. These studies are useful since they link simulation post-processing with measured defect behaviour compared to only using things such as temperature plots.

For this literature review it is enough to show that Niyama-type indicators combines local thermal conditions with feeding risk and are commonly used for shrinkage-related porosity assessment. In the thesis results, these indicators should be interpreted comparatively. A useful question is whether shell-thickness changes move predicted risk regions, reduce or increase risk volumes, or at all changes the connection between last-to-freeze regions and the feeding paths.

This matters because an indication from simulation is not the same as an actual defect in the casted product. A low or high indicator value should not be considered as proof of casting acceptability unless it is supported by calibration, inspection evidence or literature-based thresholds that are appropriate for the alloy and casting conditions. Thus overclaiming from a single defect criterion should be avoided. A better method is to compare several physically connected outputs, such as solidification time, thermal gradient, last-liquid location, porosity indicator and feeding-path continuity.

Other possible defects should also be acknowledged. Misruns and cold shuts are likely to happen when metal loses viscosity before filling is complete. This is especially true for thin regions. Hot tearing and residual stress can be influenced by shell

stiffness, restraint and the contraction or solidification path of the alloy. A shell cracking can introduce inclusions, flash or local surface defects. Gas-related porosity depends on melt quality, shell permeability and the local pressure conditions. These issues may not all be modelled in the simulation phase, but they matter because an ideal thickness for shrinkage alone may not be the same as that for the total quality requirement.

2.7 Simulation workflows and validation practices

ProCAST [14], which is used in the present study, is one of several commercial casting simulation tools. Other examples include MAGMASOFT, FLOW-3D CAST and Altair Inspire Cast, which are used to model combinations of filling, heat transfer, solidification, stress and defect prediction. The thesis proposal identifies work with commercial simulation software, implementation of process parameters and CAD models, sensitivity analysis on shell thickness, parameter calibration, documentation and possible validation against real process cases. This is consistent with the literature where simulation is used to reduce trial-and-error and to focus on the most influential mechanisms.

A recurring theme in the literature is that simulation quality depends strongly on input data. Alloy properties, shell properties, preheat temperature, heat-transfer coefficients, radiation assumptions, convection conditions, mesh quality and geometry simplifications can influence predicted solidification behaviour. Behera et al. [3] showed the importance of thermo-mechanical behaviour in investment casting shells, while Rafique and Iqbal [19] noted that investment casting heat transfer depends on changing boundary conditions, surface exposure and geometry. These studies support documentation of the simulation assumptions.

The metal-shell interface is especially important since contact is not ideal during casting. Air gaps, surface roughness, shell temperature, metal pressure, oxide formation and shrinkage can have an impact on the effective heat-transfer. Dong, Yivey et al. [9] used thermocouples to investigate the heat-transfer coefficients of the interfaces during single-crystal blade casting. On the other hand Szeliga et al. [25] used inverse modelling to look at the interface thermal resistance between a nickel superalloy and a ceramic shell. These shows that shell-thickness effects cannot be modelled reliably by changing shell geometry alone if the interface condition is not properly studied and modelled.

Torroba et al. [26] provide a useful example of an industrial simulation workflow for nickel-based superalloy nozzle guide vanes. Their Part I study combined casting trials, temperature measurements, inverse modelling, thermal calibration and porosity prediction. They then compared simulation predictions with observed porosity. In Part II [27], Torroba et al. extended their work to also look at grain-structure prediction. Here the first part of their work is relevant because it shows how thermal calibration and validation can strengthen defect predictions. The second part would be relevant if in future work shell-thickness-driven cooling differences are connected

to the grain-structure.

Upadhyaya et al. [28] also show that staged simulation can be useful. They note that complete investment-casting simulation, including view-factor calculations, filling and solidification, are computationally demanding for complex geometries. Thus due to the limited time available this supports first establishing a baseline model. Followed by a limited but well-documented sensitivity study. More detailed calibration and post-processing can then be focused on the most important cases.

From this the workflow can be constructed into different stages. First, a baseline model should be created using available CAD geometry, nominal shell thickness, selected alloy data and process settings. Second, uncertain boundary conditions should be calibrated within physically reasonable ranges if measured temperatures, known solidification behaviour or trial data are not available. Third, shell thickness should be varied either globally or locally depending on available data. Fourth, simulation cases should be compared through outputs such as solidification time, hot-spot location, porosity indicators, temperature gradients, cooling-rate trends and connection to feeding paths.

The workflow should also separate sensitivity analysis from validation. Sensitivity analysis shows how much the model changes when shell thickness or other inputs are varied. Validation instead shows whether the model represents the casting behaviour with reasonable accuracy. While these are related they are clearly not identical. If trial data is limited it can still give valuable insight into the process through its sensitivity study. However the strength of any quality prediction should be stated according to the level of available validation.

2.8 Research gap and implications for the thesis

The literature confirms that ceramic shell thickness affects investment casting through heat extraction, feeding conditions, defect formation and dimensional behaviour. However it also shows that no universal shell thickness can be transferred directly from one component or alloy to another. The effect depends on component geometry, gating, alloy, shell material, shell build quality, mould preheat, external heat exchange and interface behaviour. This creates a research gap for the proposed thesis. The task is not to repeat general investment casting theory, but to build a documented simulation workflow that identifies shell-thickness sensitivity for a specific aerospace component.

The first gap is that there is no way to transfer things specifically between components. Published studies provide useful mechanisms, especially for nickel-based superalloy turbine-blade castings, but these are literature cases rather than the component used in this study. The present study instead uses the Hub-sector introduced in Chapter 1, rather than a turbine blade. The published turbine-blade studies therefore do not directly define the shell-thickness response of the selected geometry, alloy data, process settings or quality targets. Thus the literature should be used to to

justify the variables and outputs, while using simulation to establish the response of the specific industrial case.

A second gap concerns non-uniform shell thickness. Industrial shell building does not produce a perfectly constant layer everywhere. Local geometry, pattern orientation, drainage and coating practice can create thicker and thinner regions. Rui-zhe Dong et al. [8] show that shell thickness uniformity is relevant to precision in turbine blade casting. If measured shell-thickness data become available from the GOM measurement system (GOM-scanned data), sectioning or other process measurements, it would be possible to compare an idealised uniform-thickness model with a more realistic thickness representation. If such data are not available, local sensitivity cases can still identify regions where thickness variation is most likely to influence quality.

A third gap is validation. Simulation studies become more relevant when they reproduce measured temperatures, known hot spots, porosity locations, dimensional trends or microstructural observations. The thesis proposal states that validation against material data from trial castings is possible but not required. Even if full validation is not achievable, the literature suggests that assumptions should be documented and predictions compared against any available evidence. If no defect data exist, the model can still be compared to expected behaviour and published case studies.

The thesis contribution should therefore be used to understand the process and develop a simulation workflow. Instead of an analysis to find the ideal shell thickness. A significant contribution to this would if possible be to rank the sensitivity of selected outputs to shell thickness, identify safer and more dangerous thickness ranges for the component, material and explain the mechanisms behind the observed trends. This would support a casting procedure development by connecting process modelling with the actual casted objects and their quality.

Thus for a completed study there are three questions that the thesis should try to answer. First, how does shell thickness change the thermal and feeding behaviour of the selected casting? Second, which predicted quality indicators are most sensitive to shell thickness and where are the critical regions located? Third, how reliable are these predictions given the available material data, boundary-condition assumptions and validation evidence? If these questions are answered the work would be useful even if the production of final shell specifications are outside the scope of the thesis.

2.9 Chapter conclusion

The literature supports the focus on shell-thickness sensitivity in investment casting. Ceramic shell thickness influences thermal history, feeding, shrinkage risk, dimensional behaviour and potentially microstructure. These effects are shown by alloy response, shell material properties, interface heat transfer, process conditions and component geometry. For IN718 and other nickel-based superalloys, this coupling is

important because solidification behaviour, segregation and late-stage feeding have a large impact on the final material quality.

The literature review therefore views shell thickness as a process variable rather than an isolated design input. It also separates literature review from theory and methodology. The review establishes why shell thickness matters, what previous studies have shown, why material-specific behaviour must be considered and why simulation with calibration and validation is appropriate. The later theory chapter can then define the governing equations, porosity criteria, solid-fraction thresholds and sensitivity measures in detail.

The most important conclusion is that shell-thickness analysis should be treated as a coupled thermal, metallurgical and modelling problem. The literature does not support a simple rule that thicker or thinner shells are always better. Instead, it supports a systematic simulation workflow in which shell thickness is varied, the resulting thermal and defect indicators are compared and the predictions are judged against material behaviour and any available validation evidence.

3

Theory

3.1 Purpose and scope of the theory chapter

The purpose of this chapter is to establish the physical and modelling basis for a simulation study of how IN718 cast material behaviour and ceramic shell thickness influence investment casting results. The literature review identifies what previous authors have reported. The theory chapter instead defines the mechanisms and quantities that should be understood before the simulations are designed, evaluated and discussed. Here the most important parts are heat transfer through the ceramic shell, solidification of IN718, feeding of interdendritic liquid, formation of shrinkage porosity, and the use of criteria such as Niyama and Niyama star microporosity indicators.

It is important to separate theory from potential validation and experiments. The thesis does not include direct measurement or verification of IN718 material properties. These are instead treated as simulation inputs taken from the selected software database, literature, or company-defined data. The theory chapter instead explains why properties such as liquidus temperature, solidus temperature, latent heat, density, thermal conductivity and the fraction-solid curve matter. This however does not mean that these properties were studied as part of the project. The main role of the study is instead to investigate how the chosen and documented inputs affect the simulated sensitivity to ceramic shell thickness.

Investment casting is a process which couples many different parts together. The final casting quality is not merely controlled by shell thickness nor alloy data alone. A change in shell thickness modifies the thermal resistance between metal and surroundings. This changes the cooling rate, thermal gradient, solidification time, fraction-solid evolution and local ability of remaining liquid to feed solidification shrinkage. The material then determines how sensitive the casting is to that thermal history. Because each alloy has its own freezing range, latent heat, density change, viscosity, conductivity, dendritic permeability and solid-fraction curve, the theory has to connect geometry, material data and numerical post-processing.

The chapter is written for a study using casting simulation software rather than for the derivation of a complete casting solver. The equations are therefore presented as the governing parts and interpretation tools. They define the quantities that appear in simulation outputs such as: temperature, fraction solid, cooling rate, thermal

gradient, solidification time, Niyama value, Niyama star microporosity and shrinkage-porosity volume fraction. Exactly how these are implemented within commercial software solvers can differ drastically. This is especially true for things such as turbulence, filling, radiation, contact resistance and porosity models. However their views remains the same throughout all of them. A simulation is only credible if the predicted data is consistent with the heat-transfer and solidification behaviour described in this chapter.

3.2 The ceramic shell system

The ceramic shell is the main way through which heat is transferred from the metal to the surrounding environment. Its behaviour is therefore not limited to geometry. Shell thickness, porosity, permeability, thermal conductivity and heat capacity all influence the thermal the cast object. For this reason, the ceramic shell should be treated as a material system rather than only as a boundary. Reviews of ceramic moulds for nickel-superalloy investment casting emphasise that refractory choice, binder chemistry, porosity, permeability, thermal expansion, strength and chemical reactivity all influence mould performance [15, 13]. Open porosity is especially important due to its effect on permeability and transfer of gas through the shell. Closed porosity on the other hand may affect strength and thermal properties, but does not provide the same gas-transport function. Within a simulation these details are often indirectly represented through effective thermal conductivity, density, heat capacity, emissivity and interface heat-transfer coefficients. The simplification is necessary, but it should be stated clearly because it means that a simulated "shell thickness" is in reality a combined representation of shell geometry and assumed material properties.

Here shell thickness is important for two reasons which interact with one another. First of all it changes the thermal resistance between the casting and its surroundings. Secondly it can change where the casting can be fed during solidification. A thin shell usually has faster heat removal and thus a reduced time for hot spots to remain liquid. However if local regions freeze too quickly their feeding channels can be cut off. A thicker shell may keep thin sections open for longer and may raise local Niyama values in some cases, but it can also retain heat and delay solidification in already large sections. The correct view of this is therefore not that thin or thick shells are universally better. The correct question is whether a given shell thickness creates a better combination of temperature gradient, cooling rate and feeding path for the specific alloy and component.

3.3 IN718 as the cast material

IN718 is a nickel-based superalloy used where strength, corrosion resistance and high-temperature capability are required. Its composition is based on a nickel matrix with substantial chromium and iron, and strengthening elements such as niobium, molybdenum, titanium and aluminium. In wrought or heat-treated form, the alloy

obtains much of its strength from precipitation of metastable γ'' and γ' phases. However in a casting simulation the most important features are not the final heat-treated precipitates but the solidification behaviour, liquidus and solidus temperatures, latent heat, thermal conductivity, heat capacity, density, viscosity, fraction-solid curve and segregation behaviour.

During solidification of IN718, the alloy does not freeze at a single temperature. It solidifies over a temperature interval, usually called the freezing range,

$$\Delta T_f = T_L - T_S, \quad (3.1)$$

where T_L is the liquidus temperature and T_S is the solidus temperature. Within this interval, liquid and solid coexist in the mushy zone. The solid fraction f_s increases from 0 to 1 while the liquid fraction f_l decreases according to

$$f_l = 1 - f_s. \quad (3.2)$$

The shape of the $f_s(T)$ curve is central to the simulation because it determines how latent heat is released and how quickly the remaining liquid becomes difficult to feed.

IN718 is particularly sensitive to segregation of niobium during solidification. Niobium partitions to the interdendritic liquid, which can cause formation of Nb-rich Laves phase and other interdendritic constituents late in solidification [17, 22]. This is important for porosity because the last liquid is located in the same interdendritic regions where feeding becomes increasingly difficult. Therefore the alloy has two issues: a thermal challenge, because the freezing range creates a mushy zone, and a metallurgical challenge, because the remaining liquid becomes chemically different from the primary dendrites. A simulation that only reports final solidification time without considering fraction-solid evolution and feeding risk can miss what caused the defect.

IN718 material data must be handled carefully. The thermophysical properties used in casting software are often temperature dependent, and values may differ between databases or alloy variants. For example, thermal conductivity $k(T)$, specific heat $c_p(T)$ and density $\rho(T)$ change during heating and cooling, and latent heat is released across the freezing interval rather than at one exact temperature. These properties should be documented as input assumptions rather than presented as experimentally verified values. The most important part here is consistency. If shell-thickness cases are compared, the same IN718 dataset, fraction-solid curve and porosity-model settings must be used unless the explicit purpose is to test material-data sensitivity. Any change in liquidus, solidus, latent heat or fraction-solid curve can affect the predicted hot spots and porosity indicators just as much if not more than a change in the thickness of the shell.

3.4 Fraction solid and the meaning of the 0.7 cutoff

The solid fraction f_s is one of the most useful variables in casting simulation because it shows how far the alloy has progressed through solidification. It is also easier to

interpret physically than temperature alone. Two locations at the same temperature may not have the same feeding state if the local alloy composition, cooling path or material data differ. The fraction-solid field gives a direct indication of where liquid remains, where feeding channels may still exist and where the mush has become too rigid or impermeable for effective flow.

A commonly used practical cutoff is $f_s = 0.7$. It should however not be presented as a universal physical constant. It is instead better used as an engineering threshold to mark the regions where the alloy is mostly solid but still contains enough liquid that feeding, shrinkage and microporosity is still developing. Above this value, the dendritic network is usually sufficiently developed that the remaining liquid is restricted to interdendritic channels. Flow no longer behaves like bulk liquid filling of an open cavity. It instead behaves more like pressure-driven seepage through a porous mush. Literature using solid-fraction analysis for shrinkage prediction reports that shrinkage porosity is mainly associated with late two-phase solidification and that when the solid fraction is larger than about 0.7 the remaining melt feeds through the mushy region rather than as freely flowing liquid [29, 16].

The value can be expressed in the thesis as

$$f_{s,\text{crit}} \approx 0.7, \quad (3.3)$$

where $f_{s,\text{crit}}$ is a chosen post-processing threshold for identifying the late mushy region. The practical meaning is that the simulation can be interrogated for the time, location and connectivity of regions where

$$f_s \geq f_{s,\text{crit}}. \quad (3.4)$$

These regions are relevant because they are close to the point where feeding becomes difficult. If a casting section reaches $f_s \geq 0.7$ while it is still isolated from a portion of the larger liquid sample the risk for shrinkage porosity within that area increases. If a feeder or neighbouring section remains sufficiently liquid and connected, feeding may still compensate for contraction.

The cutoff is also useful for comparing simulation cases. For example, shell-thickness cases can be compared using the time required for a critical region to reach $f_s = 0.7$, the volume of material above $f_s = 0.7$ while connected to the riser, or the last region to pass through $f_s = 0.7$. This provides more information than final solidification time alone. Two cases may have similar total solidification times, but one may lose feeding connectivity earlier. That case would be more vulnerable even if it appears acceptable from a simple hot-spot plot.

The limitation is that 0.7 is not the same for every alloy, grain structure or simulation objective. Dendrite coherency, rigidity and permeability depend on morphology, cooling rate and alloy chemistry. If a software manual or company standard defines a different value for the specific IN718 database, that value should be used in its stead. In the absence of such a calibrated value, $f_s = 0.7$ is a legitimate and primarily transparent threshold for identifying late-stage mushy solidification.

3.5 Heat-transfer theory for investment casting

The thermal part of investment casting simulation is governed by transient heat transfer with phase change. In enthalpy form, the energy equation can be written as

$$\rho \frac{\partial H}{\partial t} = \nabla \cdot (k \nabla T) + \dot{q}, \quad (3.5)$$

where ρ is density, H is specific enthalpy, t is time, k is thermal conductivity, T is temperature and $\dot{q}$ represents internal or numerical source terms if they are present within the solver. The enthalpy is commonly written as a heat part plus a latent-heat contribution,

$$H(T) = H(T_0) + \int_{T_0}^T c_p(\theta) d\theta + L f_l(T), \quad (3.6)$$

where c_p is specific heat, L is latent heat and $f_l(T)$ is liquid fraction. This form is useful because it accounts for heat release during solidification without requiring the latent heat to be released at a single temperature [23, 19].

For shell-thickness sensitivity, the casting-shell-environment path can be approximated by a thermal-resistance model. If the heat flux is simplified as one-dimensional through a shell of thickness L_s , the total thermal resistance per unit area is

$$R''_{\text{tot}} = \frac{1}{h_{m-s}} + \frac{L_s}{k_s} + \frac{1}{h_{s-\infty}}, \quad (3.7)$$

where h_{m-s} is the metal-shell interfacial heat-transfer coefficient, k_s is effective shell thermal conductivity and $h_{s-\infty}$ is the effective heat-transfer coefficient from the shell outer surface to the surroundings. The corresponding heat flux can be approximated as

$$q'' = \frac{T_m - T_\infty}{R''_{\text{tot}}}. \quad (3.8)$$

This equation shows the basic role of shell thickness. Increasing L_s increases the conductive resistance L_s/k_s and usually reduces heat flux if other terms are unchanged.

For a layered ceramic shell, the effective conductivity in the thickness direction can be written from the series-resistance analogy as

$$k_{s,\text{eq}} = \frac{L_{\text{tot}}}{\sum_{i=1}^N L_i/k_i}, \quad (3.9)$$

where $k_{s,\text{eq}}$ is the effective through-thickness conductivity, L_{tot} is the total shell thickness, L_i is the thickness of layer i , k_i is the thermal conductivity of that layer and N is the number of shell layers. This expression comes from the same composite-wall resistance approach used by Behera et al. for layered investment-casting shells. It is useful when primary and backup coats are represented separately, but if the simulation uses one homogenised shell material then k_s in Equation (3.7) should be interpreted as this effective value.

The thermal-resistance model is however only a simplified aid. Real castings are three-dimensional, the shell is curved and layered, radiation can dominate at high temperature, and contact resistance changes with time. However Equation (3.7) explains why the shell thickness has an effect on solidification. A thicker shell slows down the heat extraction because the shell stores heat and adds resistance. A thinner shell removes heat faster but may lead to rapid freezing of thin sections. The effect is therefore geometry dependent. In a thin wall, more heat extraction may prevent hot spots but may also cause premature freezing. In a thick boss or hub, reduced heat extraction may extend an already risky hot spot.

At investment casting temperatures, radiation from the shell surface can be important. A simplified radiative heat flux is

$$q''_{\text{rad}} = \varepsilon \sigma F (T_s^4 - T_{\text{sur}}^4), \quad (3.10)$$

where ε is emissivity, σ is the Stefan-Boltzmann constant, F is a view factor, T_s is shell surface temperature and T_{sur} is surrounding temperature. For convenience, radiation is sometimes expressed through an equivalent radiative heat-transfer coefficient,

$$h_{\text{rad}} = \varepsilon \sigma F (T_s + T_{\text{sur}})(T_s^2 + T_{\text{sur}}^2), \quad (3.11)$$

so that $q''_{\text{rad}} = h_{\text{rad}}(T_s - T_{\text{sur}})$. View factors are important in investment casting because parts of a complex shell can radiate toward other hot shell surfaces rather than toward a cold environment. Upadhyaya et al. [28] showed that view-factor effects and radiation shielding can move hot spots away from simply thick geometric regions, which is directly relevant when interpreting shell-thickness simulations.

The metal-shell interface heat-transfer coefficient comes with another major uncertainty. The interface is not a perfect bond. Roughness, air gaps, shell preheat, metallostatic pressure, local contraction and solidification shrinkage affect the real contact. Studies on nickel-superalloy casting simulations show that interface heat-transfer calibration can be necessary to reproduce measured temperatures and solidification behaviour [25, 26]. It is therefore important not to overanalyse minor differences between shell-thickness cases if the interface coefficient is not entirely correct. A sensitivity study can still be useful, but the uncertainty in boundary conditions should be considered.

3.6 Cooling rate, thermal gradient and solidification time

Three thermal outputs are especially important when discussing casting quality: cooling rate, thermal gradient and solidification time. The local cooling rate is the magnitude of the time derivative of temperature,

$$\dot{T}_c = \left| \frac{\partial T}{\partial t} \right|. \quad (3.12)$$

The local thermal gradient is

$$G = |\nabla T|. \quad (3.13)$$

The solidification time can be defined in different ways depending on the chosen temperature or fraction-solid criterion. A common definition is the time between reaching liquidus and reaching solidus,

$$t_s = t(T = T_S) - t(T = T_L), \quad (3.14)$$

although in practical post-processing it may also be defined by fraction-solid thresholds.

These variables are interconnected and should as such be treated together. A high cooling rate may refine microstructure, but it can also increase feeding pressure drops. A high thermal gradient tends to cause directional solidification and feeding toward hotter liquid, but it may also create thermal stresses. A long local solidification time often marks a hot spot, but the risk depends on whether the region remains connected to feeding liquid. The most useful simulation discussion is therefore not simply that one case has higher or lower cooling rate, but how cooling rate, gradient and solidification sequence combine.

The secondary dendrite arm spacing (SDAS) is commonly related to cooling rate by a power law,

$$\lambda_2 = C_\lambda \dot{T}_c^{-n}, \quad (3.15)$$

where λ_2 is SDAS, C_λ is an alloy-dependent constant and n is often near 1/3 for many cast alloys. However the exact constants are alloy specific. The important theoretical point is that faster cooling tends to reduce SDAS, while slower cooling tends to increase it. This has two consequences. First, microstructure becomes finer with faster cooling. Second, dendrite-arm spacing influences mushy-zone permeability, so it affects feeding and microporosity. As Carlson and Beckerman [5] showed the Niyama star microporosity framework explicitly accounts for this link through the role of λ_2 .

For a simple casting, Chvorinov's [7] rule provides a useful reason for why larger sections solidify slowly:

$$t_s = B \left(\frac{V}{A} \right)^n, \quad (3.16)$$

where V/A is the casting modulus, B is a mould/process constant and n is commonly close to 2 in idealised cases [7, 4]. Investment castings are too complex to design only from this rule, but the rule explains why thick sections, low-surface-area regions and thermally shielded areas tend to become last-to-freeze locations.

3.7 Shell thickness as a thermal and mechanical variable

Shell thickness influences heat transfer primarily through conduction resistance and thermal mass. Increasing shell thickness increases the term L_s/k_s in Equation (3.7).

It also increases the amount of ceramic that must be heated during pouring and cooling. Thus the shell acts as both a barrier and a source of heat. If the shell is preheated, a thicker shell can supply heat back to the casting early in the process. Later on it can instead slow down the heat loss to the surroundings. This is why the effect of shell thickness cannot be reduced to a single statement such as "thicker means slower cooling". The time-dependent preheat state and surface boundary conditions also matter.

A useful way to express thickness sensitivity is through a normalised response. For a simulation output y , such as minimum Niyama value or total solidification time, the local sensitivity to shell thickness may be approximated by

$$S_{y,L_s} = \frac{\Delta y/y_0}{\Delta L_s/L_{s,0}}, \quad (3.17)$$

where y_0 and $L_{s,0}$ are baseline values. This allows for comparison between responses with different units. For example, a small percentage increase in shell thickness may produce a large percentage change in minimum Niyama value but only a small percentage change in total solidification time. Such a result would suggest that porosity risk is more sensitive to shell thickness than global cooling time.

Shell thickness also has mechanical consequences. Thicker shells usually increase stiffness and handling strength, but local thickness variations can create stress concentrations and introduce variance in its thermal expansion. Behera et al. [3] reported thermo-mechanical shell analysis in which shell temperature, stress and deformation changed with geometry and local shell conditions. Shell cracking matters because it can cause metal leakage, dimensional error, surface defects or altered heat transfer. Therefore, even if a thinner shell gives favourable thermal results in simulation, it may not be manufacturable. Thus it is useful to separate what might be its thermal and process optimum.

Furthermore there is an issue with the shell's dimensional accuracy. Shell thickness non-uniformity can contribute to local deviations because the ceramic shell is also the mould boundary. Recent turbine-blade work from Dong, R.-Z., et al. [8] connects ceramic shell thickness uniformity with precision in investment casting, reinforcing that shell thickness affects both thermal history and geometry. For a simulation this means that shell thickness should ideally be based on measured or realistic values, not only on nominal coat counts. If only nominal thicknesses are available, the thesis should state this limitation.

3.8 Feeding theory and shrinkage porosity

Shrinkage porosity forms when the volume contraction associated with cooling and solidification cannot be compensated by feeding liquid. In a perfect situation a feeder or reservoir remains liquid longer than the casting and supplies metal to compensate for contraction. During a real casting process feeding may be interrupted by premature freezing, long feeding distances, insufficient temperature gradient, poor

riser design or restricted interdendritic flow in the mushy zone.

It is useful to distinguish macroshrinkage from microporosity. Macroshrinkage is associated with larger cavities or centreline shrinkage in last-to-freeze regions. It is often related to isolated liquid pools and insufficient feeder action. Microporosity is smaller and distributed. It forms in interdendritic regions where liquid feeding becomes difficult near the end of solidification. In aerospace castings, microporosity is important because even small pores may reduce fatigue performance, create leakage paths or fail inspection requirements.

In the mushy zone, liquid feeding can be approximated as flow through a porous medium. Darcy's law can be written as

$$f_l \mathbf{u}_l = -\frac{K}{\mu_l} \nabla p, \quad (3.18)$$

where $\mathbf{u}_l$ is the intrinsic interdendritic liquid velocity, $f_l \mathbf{u}_l$ is the effective Darcy or superficial liquid velocity, K is mush permeability, μ_l is liquid viscosity and p is pressure. This form makes the liquid-fraction dependence explicit. However if $\mathbf{u}_l$ is defined as an effective Darcy velocity, the factor f_l can be put into the velocity definition. Permeability decreases rapidly as the solid fraction increases. The simplified Kozeny–Carman expression illustrates this:

$$K \propto \lambda_2^2 \frac{f_l^3}{(1 - f_l)^2} = \lambda_2^2 \frac{(1 - f_s)^3}{f_s^2}. \quad (3.19)$$

As f_s approaches 1, f_l becomes small and permeability drops. This is the physical reason the late mushy zone is vulnerable to microporosity. The liquid is still present, but it cannot move easily enough to compensate for shrinkage.

The pressure drop needed to feed the mushy zone increases when feeding distance is long, permeability is low, viscosity is high or solidification rate is high. If the pressure in the interdendritic liquid falls below a critical value, pores can nucleate and grow. This explains why a simulation cannot assess porosity only by looking for the last point to freeze. The last point matters, but the feeding path and pressure drop through the mush matter as well. Carlson and Beckermann [5] derived their dimensionless Niyama approach from this Darcy-flow picture and used it to combine thermal conditions, material properties and pore-volume fraction.

For shell-thickness studies such as this one. Shrinkage porosity should be interpreted through the solidification sequence. A shell condition that creates directional solidification toward a feeder can reduce porosity even if the casting cools more slowly. A shell condition that creates isolated hot spots or prematurely freezes feed paths can increase porosity even if the average cooling rate seems acceptable. It is therefore useful to plot and compare the location of predicted pores relative to feeders, gates, thick sections and section transitions.

3.9 Niyama criterion

The Niyama criterion is a common thermal indicator for shrinkage-related porosity. It is defined as

$$Ny = \frac{G}{\sqrt{\dot{T}_c}}, \quad (3.20)$$

where G is the local temperature gradient and $\dot{T}_c$ is the local cooling rate, usually evaluated near the end of solidification. The physical interpretation is that a high gradient supports directional feeding, while a high cooling rate increases the demand for feeding. Low Ny values therefore indicate regions where feeding-related shrinkage is more likely [18, 6].

The Niyama criterion is used since it can combine two separate variables into one. A hot-spot plot which identifies where metal remains hot for a long time, but it does not indicate whether the surrounding thermal gradient is capable of feeding. The Niyama criterion adds this information. Additionally it can be used to compare shell-thickness cases since their thickness can change both G and $\dot{T}_c$. A thicker shell could reduce cooling rate, which can increase Ny . However it might simultaneously reduce the thermal gradient. Which in turn can decrease Ny . Thus the final result depends on which of these effects is the largest.

Since the Niyama criterion is not a universal law it shouldn't be taken as gospel. Carlson and Beckermann [6] mentions that there are thresholds for Niyama values which are dependent on both the alloy and conditions. As well as that the original criterion does not directly provide the amount of porosity. Neither does it correctly predict gas porosity or all forms of isolated-liquid shrinkage. Therefore Niyama is most useful as a comparative risk indicator for shrinkage. It can identify where shell-thickness changes improve or worsen the thermal feeding condition. However it should be combined with fraction-solid plots, solidification sequence and any available experimental evidence.

A practical way to use the criterion is to report the minimum value in critical regions,

$$Ny_{\min,\Omega} = \min_{\mathbf{x} \in \Omega} Ny(\mathbf{x}), \quad (3.21)$$

where Ω is a selected casting region. This avoids relying only on full-field images. It could be useful to compare $Ny_{\min,\Omega}$ across shell thicknesses and materials, while also showing contour plots to verify whether the minimum occurs in a physically meaningful region or only at a numerical boundary.

3.10 Niyama star microporosity

The Niyama star microporosity indicator, written here as Ny^* , is a dimensionless extension of the original Niyama criterion. It was developed to overcome two limitations of the conventional criterion: the need for empirical alloy-specific thresholds

and the inability of Ny alone to directly estimate shrinkage pore volume fraction. The dimensionless form includes local thermal conditions as well as alloy properties and solidification characteristics [5]. This is why it is more appropriate to call it a Niyama star microporosity indicator rather than only a Niyama plot.

One common form is

$$Ny^* = \frac{G\lambda_2\sqrt{\Delta P_{cr}}}{\sqrt{\mu_l\beta\Delta T_f\dot{T}_c}} = Ny\lambda_2\sqrt{\frac{\Delta P_{cr}}{\mu_l\beta\Delta T_f}}, \quad (3.22)$$

where λ_2 is SDAS, ΔP_{cr} is the critical pressure drop for pore formation, μ_l is liquid viscosity, β is solidification shrinkage and ΔT_f is the freezing range. The expression shows why Ny^* is not only a thermal criterion. It includes the material's ability to feed through the mush and the pressure condition for pore formation.

If SDAS is approximated by Equation (3.15) with $n = 1/3$, the thermal dependence of Ny^* can be written as

$$Ny^* = C_\lambda \frac{G}{\dot{T}_c^{5/6}} \sqrt{\frac{\Delta P_{cr}}{\mu_l\beta\Delta T_f}}. \quad (3.23)$$

This is important for shell-thickness sensitivity. The original Niyama criterion scales with $G/\dot{T}_c^{1/2}$, while this simplified Ny^* form scales with $G/\dot{T}_c^{5/6}$. The stronger cooling-rate dependence comes from the effect of cooling rate on SDAS and therefore on permeability. A shell-thickness change that modifies cooling rate can therefore have a greater influence on Niyama star microporosity predictions compared to the regular Niyama values.

Carlson and Beckermann [5] relate Ny^* to the critical liquid fraction $g_{l,cr}$ at which feeding is no longer able to hinder porosity. Here β is the total solidification shrinkage, defined as $\beta = (\rho_s - \rho_l)/\rho_l$. The final shrinkage pore volume fraction can be approximated by

$$g_p = \beta' g_{l,cr}, \quad \beta' = \frac{\beta}{\beta + 1} = \frac{\rho_s - \rho_l}{\rho_s}, \quad (3.24)$$

where g_p is pore volume fraction, ρ_s is solid density and ρ_l is liquid density. In commercial software these might be implemented relationships internally and report a microporosity field rather than Ny^* itself. It is however still useful to explain the theory behind it since it explains why Niyama star microporosity is influenced by both thermal and material data.

Because the model is derived from Darcy feeding through the mushy zone, it should be viewed as a feeding-related shrinkage or microporosity indicator, not as a gas-porosity model or a complete predictor of isolated-liquid shrinkage. This is important because commercial software outputs can appear more definitive than the underlying criterion actually is.

Guo et al. [11] applied a dimensionless Niyama approach with ProCAST to Ni-based

superalloy castings and reported that the method could predict both location and average volume fraction of microporosity in validation cases. This means that the use of Niyama star microporosity as a relevant output for a nickel-superalloy investment casting. Provided that the software settings and material data are documented.

3.11 Relation between shell thickness and Niyama star microporosity

The link between shell thickness and Niyama star microporosity follows from Equation (3.23). Shell thickness influences G and $\dot{T}_c$ through the thermal boundary condition. The alloy properties in the square-root term may remain constant between shell cases, but the local thermal field changes. Therefore, different shell thicknesses can produce different Ny^* values even for the same alloy and geometry.

The effect can go in either direction. If a shell-thickness increase reduces cooling rate more than it reduces the temperature gradient. Then Ny^* may increase and predicted microporosity may decrease. If the same thickness increase instead flattens the thermal gradient or creates a larger isolated hot region. Then Ny^* may decrease and microporosity may increase. This explains why published studies can report different trends depending on geometry, alloy, insulation and process conditions. In one Ni-based superalloy plate study, thermal insulation changed cooling rate and Niyama values, reducing shrinkage-porosity area fraction in the investigated configuration [24]. In turbine-blade work, shell thickness, insulation and casting temperature are treated as interacting variables rather than independent single-factor causes [21].

The shell thickness should be compared using a predefined set of linked inputs such as:

$$t_s \quad \text{solidification time,} \quad (3.25)$$

$$f_s(t, \mathbf{x}) \quad \text{fraction-solid evolution,} \quad (3.26)$$

$$G(\mathbf{x}) \quad \text{thermal gradient,} \quad (3.27)$$

$$\dot{T}_c(\mathbf{x}) \quad \text{cooling rate,} \quad (3.28)$$

$$Ny(\mathbf{x}) \quad \text{standard Niyama value,} \quad (3.29)$$

$$Ny^*(\mathbf{x}) \text{ or } g_p(\mathbf{x}) \quad \text{Niyama star microporosity / pore volume fraction.} \quad (3.30)$$

Reporting only one of these might not give the correct picture. For example a case may show lower cooling rate, but if the thermal gradient also collapses then feeding may not improve. Another case may show better standard Niyama values but still show Niyama star microporosity because the cooling-rate effect on permeability changes the pore-volume prediction.

3.12 Simulation assumptions and credibility checks

Since casting simulations are not a direct measurement of reality, but instead a numerical representation based on geometry, mesh, material data, boundary conditions, solver settings and post-processing criteria. The most important assumptions are likely to be:

1. the thermal properties of IN718 and the ceramic shell;
2. the liquidus, solidus and fraction-solid curve of IN718;
3. the metal-shell interfacial heat-transfer coefficient;
4. the shell outer-surface convection and radiation conditions;
5. the shell thickness distribution and whether it is uniform or locally varied;
6. the feeding and porosity model settings;
7. the fraction-solid threshold used for late-mushy-zone interpretation;
8. the mesh resolution in thin sections and near gates, feeders and hot spots.

Each of these can change predicted hot spots and porosity indicators. The theory chapter should therefore support a transparent methodology chapter. If a value is taken from a software database, it should be stated. If a value is taken from literature or a company standard, that should be stated. If a value is assumed because no data are available, the assumption should be treated as a limitation.

Model credibility should primarily be viewed as a sensitivity analysis and seeing if its at all plausible rather than full experimental validation. Thermal calibration and porosity prediction studies for nickel-superalloy investment castings show that matching thermal histories is important before making strong defect-prediction claims [26, 27].

A useful check is to start from simple expectations. Thin sections should cool off before thick ones unless local insulation or radiation prevents it. Feeders should remain liquid longer than the regions they feed. Low Niyama or Niyama star values should appear in the last-to-feed or low-gradient regions rather than randomly at certain mesh-artefacts. If these checks fail, the model setup should be questioned before shell-thickness conclusions are drawn.

3.13 Implications for the thesis simulation plan

The theory above leads to several practical implications. First, shell thickness should be studied as part of a thermal-feeding system. It should not be discussed only as a mould-building parameter. The simulation matrix should therefore include variables that capture temperature, fraction solid, feeding and porosity.

Secondly, the IN718 material database is to be treated as a central input assumption. Liquidus, solidus, latent heat and fraction-solid curve influence the timing of the mushy zone and the interpretation of the 0.7 solid-fraction cutoff. Since IN718 has

significant segregation and late interdendritic solidification. The last part of the freezing range is important for microporosity. The source of the IN718 data used in the simulations should therefore be documented.

Third, the 0.7 solid-fraction cutoff should be explained as a post-processing threshold for late-stage mushy solidification. It is useful for identifying when feeding changes from open liquid flow to restricted interdendritic flow. It should however not be presented as an absolute.

Fourth, Niyama and Niyama star microporosity should be treated separately. Niyama is a thermal criterion based on $G/\sqrt{\dot{T}_c}$. Niyama star microporosity is a dimensionless microporosity framework that takes the material into account more and is able to be connected to the shrinkage porosity volume fraction.

Finally, results should be compared to each other. The goal is to determine how shell thickness affects the predicted quality for the component and the process conditions. The conclusion should identify whether shell thickness mainly changes total solidification time, feeding direction, Niyama values, Niyama star microporosity or the location of the last-to-freeze region. The purpose of the theory chapter is to make those mechanisms explicit.

3.14 Traceability of formulas and symbols

The formulas in this chapter are included as explanations for simulation results. Table 3.1 summarises where the main formulas come from and what the separate terms mean. This is to prevent them from appearing as unsupported post-processing choices.

Table 3.1: Traceability of the main formulas used in the theory chapter.

Formula or equation	Origin	Meaning of the parts
$\Delta T_f = T_L - T_S$ and $f_l = 1 - f_s$	Standard solidification definitions used in casting-solidification theory [23].	T_L is liquidus temperature, T_S is solidus temperature, ΔT_f is freezing range, f_s is solid fraction and f_l is liquid fraction.
$f_{s,\text{crit}} \approx 0.7$ and $f_s \geq f_{s,\text{crit}}$	Practical late-mushy-zone cutoff. It is a post-processing threshold supported by shrinkage/fraction-solid studies that treat high-solid-fraction regions as restricted-feeding mush [29, 16].	$f_{s,\text{crit}}$ is the selected solid-fraction threshold. Regions above it are mostly solid and likely to require interdendritic rather than open-liquid feeding.
$\rho \partial H / \partial t = \nabla \cdot (k \nabla T) + \dot{q}$ and $H = H(T_0) + \int c_p dT + L f_l$	Energy conservation written in enthalpy form for phase-change heat transfer, commonly used to describe solidification modelling [23, 19].	ρ is density, H is specific enthalpy, k is thermal conductivity, T is temperature, $\dot{q}$ is a source term, c_p is heat capacity, L is latent heat and f_l determines how latent heat is released through the freezing range.
$R''_{\text{tot}} = 1/h_{m-s} + L_s/k_s + 1/h_{s-\infty}$ and $q'' = (T_m - T_\infty)/R''_{\text{tot}}$	Thermal-circuit form of conduction and convection resistances in series. Behera et al. [3] use the same series-resistance logic for investment-casting shells and overall heat transfer	h_{m-s} is metal-shell interface heat-transfer coefficient, L_s/k_s is shell conductive resistance per unit area, $h_{s-\infty}$ is the outside heat-transfer coefficient, T_m is metal temperature and T_∞ is ambient/surrounding temperature.

Continued on next page

Formula or equation	Origin	Meaning of the parts
$k_{s,eq} = L_{tot} / \sum_i (L_i / k_i)$	Composite-wall expression for a multilayer shell. Included because real investment shells are built from several ceramic coats [3, 13].	L_i and k_i are the thickness and conductivity of layer i ; L_{tot} is total thickness; $k_{s,eq}$ is the effective through-thickness conductivity used if the shell is homogenised.
$q''_{rad} = \varepsilon \sigma F (T_s^4 - T_{sur}^4)$ and $h_{rad} = \varepsilon \sigma F (T_s + T_{sur})(T_s^2 + T_{sur}^2)$	Stefan-Boltzmann radiation written with a view factor, then linearised into an equivalent radiation coefficient. Upadhyaya et al. [28] shows why view factors matter in investment-casting defect prediction	ε is emissivity, σ is the Stefan-Boltzmann constant, F is the view factor, T_s is shell surface temperature and T_{sur} is surrounding surface/ambient temperature.
$\dot{T}_c = \partial T / \partial t $, $G = \nabla T $ and $t_s = t(T = T_S) - t(T = T_L)$	Standard thermal post-processing quantities. Niyama-based shrinkage work evaluates G and cooling rate near the end of solidification [18, 24].	$\dot{T}_c$ is cooling rate, G is thermal gradient and t_s is local solidification time. They should be evaluated consistently across all shell-thickness cases.
$\lambda_2 = C_\lambda \dot{T}_c^{-n}$	Empirical SDAS-cooling-rate relation. Szeliga [24] reports an experimentally fitted relation for a Ni-based superalloy casting, while the chapter keeps the general form because constants are alloy/process dependent.	λ_2 is secondary dendrite arm spacing, C_λ is a fitted alloy/process constant, $\dot{T}_c$ is cooling rate and n is the cooling-rate exponent.
$t_s = B(V/A)^n$	Chvorinov-type modulus relation [7], useful for explaining why large-modulus regions solidify late. Upadhyaya et al. [28] discuss investment-casting modifications where mould/radiation conditions change the simple sand-casting behaviour.	B is a mould/process constant, V/A is casting modulus and n is an exponent. It is used qualitatively here, not as the thesis solver.

Continued on next page

Formula or equation	Origin	Meaning of the parts
$S_{y,L_s} = (\Delta y/y_0)/(\Delta L_s/L_{s,0})$	Normalised finite-difference sensitivity measure. This is a mathematical comparison tool rather than a casting law.	y is any output such as $Ny_{\min}$ or t_s , L_s is shell thickness and subscript 0 denotes the baseline case. It expresses percent output change per percent thickness change.
$f_l \mathbf{u}_l = -(K/\mu_l) \nabla p$ and $K \propto \lambda_2^2 f_l^3 / (1 - f_l)^2$	Darcy-flow picture of interdendritic feeding through a porous mush, with Kozeny–Carman-type permeability scaling. Carlson and Beckermann [5] use this feeding concept in their dimensionless Niyama/microporosity framework.	$\mathbf{u}_l$ is intrinsic interdendritic liquid velocity, $f_l \mathbf{u}_l$ is effective Darcy/superficial velocity, K is mush permeability, μ_l is liquid viscosity, p is pressure, λ_2 is SDAS and f_l is liquid fraction.
$Ny = G/\sqrt{\dot{T}_c}$	Original Niyama shrinkage-porosity criterion [18], later used and calibrated in different casting alloys. Szeliga [24] uses the same form for Ni-based superalloy shrinkage-porosity interpretation.	G rewards directional heat flow/feeding, while $\sqrt{\dot{T}_c}$ represents cooling-rate demand. Low Ny indicates higher shrinkage-feeding risk, but thresholds are alloy and process dependent.
$Ny^* = Ny \lambda_2 / \sqrt{\Delta P_{cr}/(\mu_l \beta \Delta T_f)}$	Dimensionless Niyama criterion proposed by Carlson and Beckermann [5] to connect thermal conditions, feeding properties and pore volume fraction; Guo et al. [11] apply the approach to Ni-based superalloy castings.	Ny^* is the Niyama star microporosity indicator, ΔP_{cr} is critical pressure drop for pore formation, μ_l is liquid viscosity, β is solidification shrinkage and ΔT_f is freezing range. It should be treated as a feeding-related shrinkage/microporosity indicator, not as a gas-porosity model.

Continued on next page

3. Theory

Formula or equation	Origin	Meaning of the parts
$g_p = \beta' g_{l,cr}$ and $\beta' = (\rho_s - \rho_l)/\rho_s$	Pore-volume relation associated with the dimensionless Niyama/microporosity model [5].	g_p is shrinkage pore volume fraction, $g_{l,cr}$ is the critical liquid fraction at pore formation, β' is an effective shrinkage factor and ρ_s, ρ_l are solid and liquid density.

3.15 Summary of key theoretical quantities

Table 3.2 Summary of theoretical quantities which while they may not be used directly have an influence on the results and should be used consistently.

Table 3.2: Key theoretical quantities for interpreting shell-thickness and IN718 casting simulations.

Quantity	Expression	Interpretation
Freezing range	$\Delta T_f = T_L - T_S$	Temperature interval over which liquid and solid coexist.
Solid fraction	$f_s = 1 - f_l$	Progress of solidification; used to identify mushy and late-mushy regions.
Late mushy cutoff	$f_{s,\text{crit}} \approx 0.7$	Practical threshold for restricted feeding and shrinkage-risk interpretation.
Cooling rate	$\dot{T}_c = \partial T / \partial t $	Rate of local heat extraction; affects SDAS and feeding pressure demand.
Thermal gradient	$G = \nabla T $	Directional feeding indicator; higher gradients generally support feeding.
Shell resistance	$R'' = 1/h_{m-s} + L_s/k_s + 1/h_{s-\infty}$	Shows how shell thickness enters the heat-flow path.
Layered shell conductivity	$k_{s,\text{eq}} = L_{\text{tot}} / \sum_i L_i / k_i$	Converts a layered ceramic shell into an effective through-thickness conductivity.
SDAS	$\lambda_2 = C_\lambda \dot{T}_c^{-n}$	Microstructural spacing; affects mush permeability.
Niyama	$Ny = G / \sqrt{\dot{T}_c}$	Thermal shrinkage-porosity risk indicator.
Niyama star micro-porosity	$Ny^* = Ny \lambda_2 \sqrt{\frac{\Delta P_{cr}}{\mu_l \beta \Delta T_f}}$	Dimensionless material-aware microporosity indicator.
Pore volume fraction	$g_p = \beta' g_{l,\text{cr}}$	Approximate link between critical liquid fraction and shrinkage pore fraction.

4

Methods

4.1 Approach

The study was structured as a comparative simulation study. A provided CAD model was used as the basis for the simulation setup, after which different shell representations were generated and evaluated. The work first focused on establishing two main shell models: a constant-thickness shell and a non-uniform shell based on physical shell cut-up measurements. These models were then compared with and without insulation wrapping. After this initial comparison, the main sensitivity study was performed using insulation-wrapped models, since this better represented the actual process.

The study was not designed as a material-property investigation. The cast alloy definition used in ProCAST was treated as a fixed input required to perform the simulations. No experimental determination of cast material properties was performed, and no comparison between different cast materials was included. The focus was instead placed on the simulation workflow, shell-thickness representation, process-parameter variation and interpretation of solidification-related defect indicators.

Due to the industrial context of the project, exact component geometry, process parameters, shell dimensions, internal modelling values and company-specific simulation settings are not reported. The methodology therefore describes the modelling procedure, simulation sequence and comparison strategy, while proprietary GAIC information is left out.

4.2 Simulation software

All casting simulations were performed in ProCAST. The software was used to simulate both mould filling and solidification. Including both stages was important because the thermal state after filling provides the starting condition for the solidification calculation. Since shell thickness and process settings can influence the temperature distribution during filling, a solidification-only approach would not have been an accurate representation of the process.

The filling calculation was computationally demanding and accounted for the majority of the simulation time. However it was still included because the filling stage affects

metal temperature, heat loss, local cooling conditions and the later development of solidification indicators. The filling and solidification stages were therefore treated as connected parts of the same process simulation.

4.3 CAD model and casting system

The simulation work was based on a CAD model provided for the project. The CAD geometry was imported into ProCAST and used as the basis for the simulation model. The model included the casting system required for the process simulation. In particular, the gating system was included as part of the model rather than being removed or treated separately.

Including the gating system was necessary because the filling behaviour, metal flow path and temperature distribution are influenced by the complete casting system. The gating system controls how the metal enters the mould cavity and how different regions of the casting are thermally loaded during the early stages of the process. Since the later solidification behaviour depends on the thermal state after filling, the gating system was not changed between the simulations.

Any geometry preparation performed before simulation was limited to making the model suitable for meshing and process simulation. The component and gating layout were treated as industrial input data. Redesigning the component or developing a new gating system was far outside the scope of the study and not something the author was qualified to do.

4.4 Ceramic shell modelling

The ceramic shell was generated within ProCAST. Two main shell representations were used in the study. The first representation was a constant-thickness shell, where the ceramic shell was generated with a uniform nominal thickness around the casting system. This represented the idealised shell setup used as the initial simplified simulation approach.

The second representation was a non-uniform shell. This model was based on improved knowledge of the real shell thickness obtained from physical shell cut-ups. These cut-ups provided measurements of the actual shell-thickness distribution and showed that the real shell could not be represented fully by a single constant thickness. Based on this information, an attempt was made to reproduce the measured shell-thickness variation in ProCAST as closely as practically possible.

The non-uniform shell model should be understood as an approximation rather than an exact digital copy of the real ceramic shell. Reproducing the measured shell geometry precisely was difficult due to the complexity of the casting geometry, the limited number of measurement locations and practical limitations in creating the shell representation inside the simulation environment. The aim was not to model

every local shell-thickness detail exactly, but to create a shell representation that was closer to the measured physical shell than the idealised constant-thickness model.

The comparison between the constant and non-uniform shell representations formed the central part of the study. The constant shell model represented a simplified modelling approach, while the non-uniform shell model represented a more realistic approximation of the measured shell condition.

4.5 Insulation wrapping

Insulation wrapping was included as a part of the simulation setup because the mould is insulated in the actual process. The unwrapped simulations were therefore not intended to represent the real process condition directly. Instead, they were included to show the influence of the insulation wrapping and to provide a reference for how strongly the thermal behaviour changes when the wrapping is added.

Both ceramic shell representations were simulated with and without insulation wrapping. This resulted in four initial model variants: a constant shell without wrapping, a constant shell with wrapping, a non-uniform shell without wrapping and a non-uniform shell with wrapping. The constant shell represented the idealised nominal shell-thickness approach, while the non-uniform shell represented an approximation of the measured shell-thickness variation.

As expected, the comparison showed a large difference between the wrapped and unwrapped simulations. The insulation heavily affected the thermal response of the mould and casting system by changing the heat exchange with the surrounding system. This combined with the actual casted parts utilising insulation wrapping means that these were the relevant models for the continued study.

The difference between the constant and non-uniform shell representations was much smaller when both models included insulation wrapping. This does not mean that the shell representation had no impact, since local differences could still appear in temperature history, heat flux, solidification behaviour and defect indicators. However, the initial comparison showed that insulation wrapping had a much larger influence on the overall thermal behaviour than the difference between the two wrapped representations.

For this reason, the later sensitivity study was performed using insulation-wrapped models. The non-uniform wrapped shell was used for the main process-parameter variations, while the constant wrapped shell was used for corresponding comparison cases.

4.6 Simulation case structure

The simulation cases were organised around two stages. The first stage consisted of the four initial model variants described above: constant shell without wrapping, constant shell with wrapping, non-uniform shell without wrapping and non-uniform shell with wrapping. This stage was used to confirm the influence of insulation wrapping and to compare the two shell representations under process-relevant wrapped conditions. Since the actual process uses insulation wrapping, the unwrapped models were not carried forward into the main sensitivity study.

The second stage consisted of the process-parameter study. This was first performed using the non-uniform shell model with insulation wrapping. The varied parameters were:

- shell preheat temperature,
- mass flow rate during filling,
- enclosure emissivity.

The mass flow rate was compared between the original simulation setting and a lower value based on measured process averages. The shell preheat temperature was evaluated at a lower, standard and higher level. The enclosure emissivity was evaluated using the standard value and one lower value. Exact numerical values are omitted due to confidentiality.

The enclosure emissivity variation had no practical influence on the investigated results. A likely reason is that the mould was exposed to the enclosure radiation condition for a limited part of the relevant process time compared with the stronger effects from filling, shell preheat and metal–mould heat transfer. For this reason the later comparisons focused mainly on shell representation, mass flow rate and shell preheat temperature.

After the non-uniform shell simulations had been evaluated corresponding simulations were performed using the constant-thickness shell model with insulation wrapping. These simulations allowed the constant and non-uniform shell representations to be compared under similar filling and thermal conditions. Since the emissivity variation had shown to have no real influence in the non-uniform shell study, it was not present in the later constant-shell comparison.

The study was not a full factorial design of experiments. Instead, it was a focused comparative simulation study intended to evaluate how shell representation and selected process parameters affected the simulated filling behaviour, solidification sequence and predicted defect indicators.

4.7 Filling and solidification simulation

Each simulation included both filling and solidification. During the filling stage molten metal was introduced through the gating system according to the defined

filling condition. ProCAST calculated the progression of the metal front, the development of the temperature field and the heat transfer between the metal, ceramic shell, insulation and surrounding environment.

After filling, the simulation continued through solidification. The solidification stage was used to evaluate how the different shell representations and process settings influenced cooling behaviour, fraction-solid development, hot-spot formation and predicted defect indicators. Since the thermal field at the end of filling affects the entire solidification sequence, the filling and solidification stages were evaluated as one connected process.

The simulations were compared through differences between cases. The main interest was not only the absolute value of a single output in one simulation, but how the response changed when the shell representation or selected process parameter was changed. The same approach was therefore used across the cases to make the comparisons consistent.

4.8 Post-processing and evaluated outputs

The simulation results were evaluated using both field outputs and time-dependent data extracted from selected locations. The post-processing focused on outputs that describe filling behaviour, solidification development, heat transfer and defect tendency.

The evaluated outputs were:

- fraction-solid development at selected time steps,
- animations of the solidification sequence,
- Niyama microporosity over the full model,
- shrinkage porosity predictions,
- mould-interface behaviour over the process,
- time-temperature curves at selected locations,
- time-heat-flux curves at selected locations,
- temperature-heat-flux plots at selected locations.

Fraction-solid plots and animations were used to study how solidification progressed through the casting. These outputs made it possible to identify which regions solidified early, which regions remained liquid for longer and where last-to-freeze areas developed. Selected time steps were used to compare the different cases at similar stages of the process.

Niyama microporosity and shrinkage porosity outputs were used as defect-related indicators. These outputs were not treated as direct proof of actual defects, but as simulation-based indications of regions where the thermal and feeding conditions may be unfavourable. The defect indicators were therefore interpreted together with the solidification sequence rather than as isolated results.

The mould-interface animation was used to examine the temperature development between the casting and mould throughout the process. This was included because the metal-mould interface is important for heat transfer and can influence the local thermal response of the casting.

Temperature and heat-flux data were extracted at a number of selected locations in the model. These monitoring locations were used to compare local thermal histories between cases. Temperature-time plots showed how quickly different regions heated, cooled and passed through the solidification interval. Heat-flux-time plots showed how local heat extraction changed during filling and solidification. Temperature-heat-flux plots were used to relate local heat extraction to the thermal state of the casting.

4.9 Comparison between simulation cases

The simulation cases were compared by examining both global field behaviour and local thermal response. Global field outputs, such as fraction solid, Niyama microporosity and shrinkage porosity, were used to identify changes in the overall casting response. Local plots from selected monitoring positions were then used to understand the thermal mechanisms behind these changes.

The comparison between the constant shell and non-uniform shell was central to the work. The constant shell case provided an idealised reference case, while the non-uniform shell case provided a more realistic approximation based on physical shell measurements. Differences between the two cases were evaluated in terms of filling behaviour, solidification sequence, hot-spot location, predicted porosity indicators and local heat extraction.

The process-parameter variations were interpreted in relation to their expected physical influence. The mass-flow variation was used to evaluate the effect of the filling condition on the thermal state entering solidification. The shell-preheat variation was used to evaluate the influence of the initial thermal state of the ceramic shell. The emissivity variation was included to test the sensitivity of the model to the enclosure radiation setting, but was not treated as a main controlling parameter after showing limited influence in the investigated cases.

4.10 Validation and comparison with process observations

A limited validation or process comparison was performed using available porosity information. The comparison was qualitative rather than a full quantitative validation. Simulation cases using process parameters close to the actual process were

compared with observed porosity indications from the real casting process.

The observed porosity indications were found to occur in broadly similar regions to the simulated defect indicators. This provided some support that the simulations captured relevant trends in the casting behaviour. However, the comparison was limited and should not be Viewed as complete verification of the model.

Because the available validation data were limited, the results are primarily interpreted as comparative simulation trends. The conclusions are therefore related to how the simulation response changes between shell representations and process-parameter cases, rather than to exact prediction of porosity amount or final component acceptability.

4.11 Confidentiality and reporting of parameters

The project was carried out in an industrial setting, and the simulation work involved proprietary GAIC information. Exact component details, process parameters, shell dimensions, internal modelling values and company-specific simulation settings are therefore not reported in full detail. This applies to both the input data and some aspects of the simulation case definitions.

To maintain clarity while protecting confidential information, the methodology is reported in terms of workflow, model structure and comparison strategy. The work describes which variables were studied, how the simulation cases were constructed and how the results were evaluated, without disclosing proprietary numerical values. This means that the public report is not intended to provide a fully reproducible industrial simulation setup. Instead, it documents the method to explain how the conclusions were reached.

5

Results & Discussion

5.1 Presentation of results

The results are presented as relative simulation trends rather than as absolute numerical process values. Due to confidentiality requirements, exact process parameters, component geometry, gating geometry and numerical axis values are not shown. The plots included in this chapter are therefore used to compare curve shape, relative separation, timing and repeated behaviour between simulation cases.

In total 294 thermal-history plots were generated during the work. These included time-temperature, time-heat-flux and temperature-heat-flux plots from selected locations. The full set of plots was used during the analysis to identify trends, but only a select few of them are included in this chapter. Including all generated plots would reduce readability without adding any real value.

The results focus on how the simulated response changed when insulation wrapping, shell, shell preheat temperature, mass flow rate and enclosure emissivity were varied. The evaluated outputs were local temperature histories, local heat-flux histories, temperature-heat-flux relationships, fraction-solid development, Niyama microporosity and shrinkage-porosity.

The porosity outputs are to be viewed as defect indicators rather than proof of defects. Niyama-type criteria and shrinkage-porosity predictions depend on the simulated thermal field, feeding conditions, material data and post-processing assumptions. For these reasons the results should be seen as comparative trends between simulation cases, not as final acceptance criteria for the component [18, 5, 16].

For confidentiality reasons, the monitoring locations are referred to using anonymised labels in the text and figure captions. Examples for these are shown in table 5.1

Table 5.1: Anonymised monitoring locations used for representative thermal-history plots.

Label used in report	Plot file location used in figures
Location A	mid_part_large_concave_y_2
Location B	mid_part_through_extrusion_thickness_y_9

5.2 Overview of simulation cases

The simulation work was organised around three main result groups. First of all four initial model variants were compared: constant shell without insulation wrapping, constant shell with insulation wrapping, non-uniform shell without insulation wrapping and non-uniform shell with insulation wrapping. This comparison was used to evaluate the influence of insulation wrapping and to compare the two shell representations under wrapped and unwrapped conditions.

Second, the non-uniform shell model with insulation wrapping was used for the main process-parameter sensitivity study. This model was selected because it represented the measured shell condition more closely than the constant shell model while also including the insulation wrapping used in the actual process.

Finally corresponding simulations were performed using the constant shell model with insulation wrapping. These cases were used to compare the idealised constant shell representation with the non-uniform shell representation under similar process conditions.

Table 5.2 shows the test matrix for the non-uniform wrapped shell sensitivity study. Test 1 corresponds to the wrapped non-uniform shell reference case from the initial four-model comparison. Tests 2–6 were used for the main shell-preheat and mass-flow comparison at standard enclosure emissivity. Tests 7–12 repeated the same combinations with lower enclosure emissivity.

Table 5.2: Overview of the non-uniform wrapped shell sensitivity cases. Exact numerical values are omitted due to confidentiality.

Test	Shell preheat temperature	Mass flow rate	Enclosure emissivity
1	Standard	Standard	Standard
2	Lower	Standard	Standard
3	Higher	Standard	Standard
4	Standard	Measured average, lower	Standard
5	Lower	Measured average, lower	Standard
6	Higher	Measured average, lower	Standard
7	Standard	Standard	Lower
8	Lower	Standard	Lower
9	Higher	Standard	Lower
10	Standard	Measured average, lower	Lower
11	Lower	Measured average, lower	Lower
12	Higher	Measured average, lower	Lower

The enclosure emissivity variation had basically no influence on the results. For this reason, Tests 7–12 are included in the case matrix, but the results discussion focuses

on Tests 1–6 and the corresponding constant-shell comparison cases.

5.3 Overall trends from the simulation cases

The investigated simulation cases showed that shell preheat temperature had the strongest influence on the predicted defect indicators. The effect of shell representation was also present, but smaller in comparison. The constant-thickness shell generally gave slightly more favourable results than the non-uniform shell representation, although the difference was not large enough to suggest that shell controlled the overall result. The enclosure emissivity variation had very limited influence, while the lower mass-flow condition produced a smaller negative effect, especially when combined with the lower shell preheat temperature.

The clearest trend was observed for shell preheat temperature. The higher shell-temperature cases consistently produced considerably lower porosity indications than the corresponding standard and lower shell-temperature cases. The standard shell-temperature cases gave intermediate results, where the predicted porosity indication was not as severe as in the lower-temperature cases, but also not as favourable as in the higher-temperature cases. The lower shell-temperature cases gave the least favourable results across the evaluated simulations and showed strong porosity indications compared with the other shell-temperature levels.

The mass-flow variation had a slight influence on the results, but significantly lower than the shell preheat temperature. The lower mass-flow condition generally produced a somewhat less favourable response. This effect was most obvious when the shell preheat temperature was also low. This suggests that the lower mass-flow condition and the lower shell temperature acted in the same unfavourable direction: both contributed to a colder or less favourable thermal state during filling and early solidification. When the shell temperature was higher, the negative effect of the lower mass-flow condition was barely legible.

The enclosure emissivity variation produced almost no visible change in the evaluated results. The temperature-history plots, heat-flux plots and defect indicators were very similar between the standard-emissivity and lower-emissivity cases. This indicates that enclosure emissivity was not a controlling parameter for the investigated simulations. A likely explanation is that the mould was exposed to the enclosure radiation condition for a limited part of the relevant process time, while the thermal response was dominated by filling conditions, shell preheat and heat transfer between the metal and its surroundings.

The comparison between the constant and non-uniform shell representations showed that the constant shell gave slightly better overall results on average. However, the difference between the two wrapped shell representations was small compared with the effect of shell preheat temperature. The non-uniform shell representation can still be relevant because it may affect local heat extraction and defect indicators, but the results suggest that the investigated process was more sensitive to shell preheat

than to the difference between the constant and non-uniform wrapped shell models.

Table 5.3: Qualitative summary of the main simulation trends.

Investigated change	Observed effect	Interpretation
Constant shell compared with non-uniform shell	Slightly more favourable overall result for the constant shell	The effect was present but relatively small compared with the influence of shell preheat temperature. Local differences may still be important.
Higher shell preheat temperature	Clearly more favourable porosity indication	Reduced initial thermal difference between metal and shell likely improved the thermal conditions for feeding and solidification.
Standard shell preheat temperature	Intermediate result	The result was acceptable relative to the lower-temperature case, but less favourable than the higher-temperature case.
Lower shell preheat temperature	Clearly worse porosity indication	Increased heat extraction during filling and early solidification likely promoted unfavourable feeding and stronger defect indications.
Lower mass flow rate	Slightly negative effect, especially with lower shell preheat	The lower mass-flow condition affected the filling thermal history and was most unfavourable when combined with a colder shell.
Lower enclosure emissivity	Negligible effect	The enclosure radiation setting did not appear to control the thermal response in the investigated cases, likely due to limited effective exposure time.

5.4 Influence of insulation wrapping

The initial model comparison showed a large difference between wrapped and unwrapped simulations. This was expected, since insulation wrapping changes the external heat exchange of the mould system. The unwrapped simulations were therefore not representative of the actual process condition, but were useful as comparison cases to show the importance of the insulation wrapping.

The difference between wrapped and unwrapped models was larger than the difference between the constant and non-uniform shell representations when both shell models included wrapping. This is as expected and indicates that insulation wrapping was a dominant part of the overall thermal boundary condition. Since the actual process uses insulation wrapping, the later sensitivity study was performed using wrapped models.

The temperature histories in Figure 5.1 show the difference between the initial wrapped and unwrapped model variants at Location A. The purpose of the figure is not to report exact temperatures, but to show the relative change in thermal behaviour when insulation wrapping is included.




Figure 5.1: Temperature histories for the initial wrapped and unwrapped model comparison at Location A. The wrapped cases show that insulation wrapping strongly affects the thermal response of the mould and casting system. Numerical axis values are omitted due to confidentiality.

The same general behaviour was visible in the heat-flux histories. The wrapped and unwrapped cases did not only differ by a small local shift; the insulation changed the way heat was removed from the system. This supports the decision to use wrapped models for the main sensitivity study. Also a note on the very visible initial spike on the heat-flux in figure 5.2. That comes from when the metal first reaches the plotting point which, since it was located within the part and not the shell, is set to a significantly lower temperature by the software.




Figure 5.2: Heat-flux histories for the initial wrapped and unwrapped model comparison at Location A. The difference between wrapped and unwrapped cases is larger than the difference between the two wrapped shell representations. Numerical axis values are omitted due to confidentiality.

5.5 Influence of shell preheat temperature

The shell preheat temperature had the strongest influence on the predicted porosity indicators. The higher shell-temperature cases gave clearly more favourable results, while the lower shell-temperature cases gave the least favourable results. The standard shell-temperature cases gave an intermediate response.

This trend is consistent with the role of shell preheat in the thermal development of the casting. A colder shell increases the temperature difference between the metal and the mould during filling and early solidification. This increases heat extraction from the melt and can promote earlier local freezing, less favourable feeding conditions and stronger defect indications. A higher preheat temperature reduces the thermal difference between metal and shell, which can keep regions open for feeding for

a longer period and reduce the predicted porosity indicators. This interpretation is consistent with the general investment-casting literature, where shell thermal behaviour, mould preheat, insulation and boundary conditions are treated as coupled factors influencing solidification and defect formation [19, 21, 24].

The influence of shell preheat temperature is shown in Figure 5.3. Tests 1–3 compare the standard, lower and higher shell preheat temperatures at the standard mass-flow condition. Tests 4–6 show the corresponding shell-preheat variation at the lower measured-average mass-flow condition.




Figure 5.3: Temperature histories at Location A for the non-uniform wrapped shell cases. The comparison shows the influence of shell preheat temperature at both standard and lower measured-average mass-flow conditions. Numerical axis values are omitted due to confidentiality.

The corresponding heat-flux histories are shown in Figure 5.4. The lower shell-temperature cases show a less favourable thermal response than the higher shell-temperature cases. This supports the interpretation that the colder shell condition increased early heat extraction and produced less favourable solidification conditions.




Figure 5.4: Heat-flux histories at Location A for the non-uniform wrapped shell cases. The heat-flux response supports the interpretation that shell preheat temperature was the dominant investigated process parameter. Numerical axis values are omitted due to confidentiality.

5.6 Influence of mass flow rate

The lower mass-flow condition produced a less favourable response than the standard mass-flow condition, although the effect was significantly smaller than the influence of shell preheat temperature. The negative influence of the lower mass flow was most visible when it was combined with the lower shell preheat temperature.

This behaviour can be explained by the connection between filling and solidification. Since the simulations included both filling and solidification, the thermal state after filling became the starting condition for the solidification stage. A lower mass flow changes the filling time scale and can increase the opportunity for heat loss before the casting reaches the later solidification stage. This becomes more harmful when the shell is also colder, because the thermal driving force for heat extraction is already larger.

The lower mass-flow condition and the lower shell preheat temperature therefore acted in the same unfavourable direction. Both contributing to a colder state during filling and early solidification. This helps explain why Test 5, which combined lower

shell temperature and lower mass flow, gave a clearly worse response compared with the corresponding higher-temperature cases.

The mass-flow effect is most clearly interpreted by comparing Test 2 with Test 5 and Test 3 with Test 6. Tests 2 and 5 both use the lower shell preheat temperature, but Test 5 uses the lower measured-average mass flow. Tests 3 and 6 both use the higher shell preheat temperature, but Test 6 uses the lower measured-average mass flow. The comparison indicates that mass flow had an effect, but that the severity of this effect depended strongly on the shell-temperature condition.

Figure 5.5 shows a second monitoring location used to support this trend. The repeated behaviour at multiple locations supports the conclusion that shell preheat temperature was the stronger parameter, while mass flow had a much smaller influence.




Figure 5.5: Temperature histories at Location B. The comparison between Tests 2 and 5 shows the influence of reducing mass flow at lower shell preheat temperature, while the comparison between Tests 3 and 6 shows the corresponding effect at higher shell preheat temperature. Numerical axis values are omitted due to confidentiality.

5.7 Influence of enclosure emissivity

The enclosure emissivity variation had practically no influence on the evaluated results. The temperature-history plots, heat-flux-history plots and defect-indicator outputs were very similar between the standard-emissivity cases and the lower-emissivity cases. Compared with shell preheat temperature and mass flow rate, enclosure emissivity was therefore not a controlling parameter in the investigated simulation setup.

This result does not mean that radiation is unimportant in investment casting. Radiative heat transfer can be large at high temperatures and depends on emissivity,

surface temperature, surrounding temperature and view factor [28, 19]. However, in the present simulations, the emissivity boundary condition appears to have a negligible influence compared with the dominant effects from metal-shell heat transfer, shell preheat temperature, filling history and insulation wrapping.

A likely explanation is that the mould was effectively exposed to the enclosure radiation condition for a limited part of the relevant process time. The most important thermal response was instead governed by the early interaction between molten metal, ceramic shell and insulation. Because of this, Tests 7–12 were not treated as a part of later comparisons.

5.8 Influence of shell representation

The constant shell representation gave slightly more favourable overall results than the non-uniform shell representation. However, the difference between the two wrapped shell representations was small compared with the effect of shell preheat temperature. This suggests that shell representation affected the thermal response, but that it was not dominant.

The constant shell model gives a more uniform thermal resistance around the casting system. The non-uniform shell model introduces local differences in shell thickness and can therefore induce differences in heat extraction. Depending on where these shell-thickness variations occur, they may either increase local cooling intensity or retain heat locally. Either effect can become unfavourable if it alters the intended solidification sequence or creates isolated regions that are difficult to feed.

The slightly better result for the constant shell should not be interpreted as meaning that a constant shell is more physically realistic. The constant shell is an idealised representation. The non-uniform shell is closer to the measured physical shell, but it is still an approximation of the real shell geometry. The result instead indicates that the measured shell-thickness variation, as represented in the simulation, introduced local thermal variations that were slightly less favourable for the evaluated defect indicators. This is consistent with the literature, where shell-thickness effects are described as dependent on geometry, shell build, mould preheat, insulation, interface behaviour and process conditions rather than following a universal rule [21, 8, 25].

Figure 5.6 shows representative constant-shell and non-uniform-shell comparisons at Location A. The comparison illustrates that the difference between shell representations was visible, but less dominant than the difference caused by shell preheat temperature.




Figure 5.6: Comparison between constant and non-uniform wrapped shell representations at Location A. The shell representation affects the local thermal response, but the difference is smaller than the influence of shell preheat temperature. Numerical axis values are omitted due to confidentiality.

5.9 Temperature–heat-flux relationship

The temperature–heat-flux plots were used to compare heat extraction at similar thermal states rather than only at equal simulation times. This was useful because some cases developed different temperature histories during filling and early solidification. Comparing heat flux directly against temperature therefore helped identify whether the difference between cases was only a timing effect or whether the local heat-extraction behaviour also changed.

The temperature–heat-flux plots supported the same general trends as the time-based plots. The lower shell-temperature cases showed more erratic thermal behaviour, while the higher shell-temperature cases gave a smoother response. The lower mass-flow condition had a smaller secondary effect, especially when combined with lower

shell preheat temperature. The emissivity variation did not produce a meaningful change in this comparison.

A representative temperature–heat-flux comparison is shown in Figure 5.7. The figure is included to show the development of the heat-flux response relative to local temperature.




Figure 5.7: Representative temperature–heat-flux plots for selected non-uniform wrapped shell cases at Location A. These plots compare local heat extraction at similar thermal states and support the trends observed in the time-based plots. Numerical axis values are omitted due to confidentiality.

5.10 Defect-indicator results

The porosity-related outputs followed the same general trends as the thermal-history plots. The shell preheat temperature had the clearest influence on the predicted

defect indicators. The higher shell-temperature cases gave the lowest and most favourable porosity indications, while the lower shell-temperature cases gave the highest and least favourable indications. The standard shell-temperature cases were between these two conditions, but more towards the acceptable range.

The figures in this section show Niyama star microporosity indicator results for the selected simulation cases. The numerical colour scale has been removed due to confidentiality. The figures should therefore be read as relative comparisons of indication pattern, affected area and repeated behaviour between simulation cases, not as quantitative contour plots.

This result is physically consistent with the heat-transfer interpretation as has been described previously. A colder shell extracts more heat during filling and early solidification. This promotes premature freezing in some regions and reduces the ability of liquid metal to feed solidification shrinkage. The result is a worse solidification sequence and stronger porosity indication. A hotter shell reduces the early thermal driving force and can support a more favourable feeding condition during the critical solidification interval.

The initial model comparison is shown in Figure 5.8. The unwrapped cases show stronger and more extensive defect-indicator regions than the wrapped cases. This is consistent with the earlier temperature-history and heat-flux comparisons, where the insulation wrapping clearly changed the thermal response of the mould and casting system.






Figure 5.8: Niyama star microporosity indicator for the initial wrapped and unwrapped model comparison. The wrapped cases give a more favourable indication pattern than the corresponding unwrapped cases. The numerical colour scale has been removed due to confidentiality.

The shell-preheat, mass-flow and shell-representation comparisons are shown in Figures 5.9 and 5.10. In these figures, each row places the constant-shell and non-uniform-shell result for the same test condition next to each other. This makes the effect of shell representation easier to compare directly, while still showing that the shell-temperature effect was stronger than the difference between the two wrapped shell representations.


Tests 2 and 3 compare lower and higher shell preheat at the standard mass-flow condition. Tests 5 and 6 show the corresponding comparison at the lower measured-average mass-flow condition. The lower mass-flow effect was much smaller than the shell-temperature effect, but it was still somewhat visible in combination with lower shell preheat. Since filling was simulated together with solidification, the filling condition affected the thermal field for the solidification stage. Thus the combination of lower shell temperature and lower mass flow produced the most adverse conditions.




Constant shell, Test 2: lower shell temperature, standard mass flow



Non-uniform shell, Test 2: lower shell temperature, standard mass flow



Constant shell, Test 3: higher shell temperature, standard mass flow



Non-uniform shell, Test 3: higher shell temperature, standard mass flow

Figure 5.9: Niyama star microporosity indicator for Tests 2 and 3. Each row compares the constant-shell and non-uniform-shell representations for the same process condition. The lower shell-temperature case shows stronger indications than the higher shell-temperature case. The numerical colour scale has been removed due to confidentiality.




Figure 5.10: Niyama star microporosity indicator for Tests 5 and 6. Each row compares the constant-shell and non-uniform-shell representations for the same process condition at the lower measured-average mass-flow condition. The numerical colour scale has been removed due to confidentiality.

The constant-shell comparison cases were slightly more favourable on average than the non-uniform shell representation. However, the difference was relatively small compared with the influence of shell preheat temperature. This indicates that shell representation can affect local defect tendency, but that the initial thermal state of the shell was more important for the investigated cases.

The enclosure emissivity variation did not produce a meaningful change in the defect indicators. This is consistent with the local temperature and heat-flux plots, where the standard-emissivity and lower-emissivity cases were very similar. The emissivity variation was therefore judged to have limited influence for this specific simulation setup.

5.11 Comparison with observed porosity indications

A limited comparison with available process observations was performed using porosity indications from process-related cases. The comparison was qualitative rather than a complete validation. The purpose was to check whether the simulated defect indicators appeared in broadly similar regions to observed porosity indications from the real casting process.

The comparison showed that the simulated defect-indicator regions were somewhat consistent with the observed porosity indications. This provides some support that the simulations captured relevant trends in the casting behaviour. However, the comparison was limited and should not be interpreted as a verification of the model. Previous investment-casting simulation studies also distinguish between sensitivity analysis and validation, especially when available process observations are limited [11, 26].

The results should therefore be read primarily as simulation trends. The strongest conclusions concern the relative effect of shell preheat temperature, mass flow rate, enclosure emissivity and shell representation. The results do not define a final accepted production setup or a validated quantitative porosity prediction.

5.12 Summary of main result trends

The main result trends can be summarised as follows:

- Insulation wrapping had a large influence on the thermal response. The wrapped models were therefore used as the relevant basis for the main sensitivity study.
- The difference between constant and non-uniform wrapped shell representations was present, but smaller than the effect of shell preheat temperature.
- The shell preheat temperature was the most influential investigated parameter. Higher shell preheat gave preferable porosity indications, while lower shell preheat gave clearly worse results.
- The standard shell-temperature cases gave intermediate results.
- The lower mass-flow condition had a smaller negative effect, especially when combined with lower shell preheat temperature.
- The enclosure emissivity variation had negligible influence on the evaluated outputs and was not treated as a controlling parameter in the later analysis.
- The constant shell representation gave slightly more favourable overall results than the non-uniform shell representation, but this should not be interpreted as meaning that the constant shell is more realistic.

Overall, the results show that the casting response was governed by the combined thermal history of filling, shell preheat, shell representation, insulation and heat transfer between metal, shell and surroundings. Among the investigated parameters, shell preheat temperature had the strongest influence on the predicted porosity indicators, while enclosure emissivity had the weakest influence.

6

Conclusion

This thesis investigated how ceramic shell representation and selected process parameters influence the simulated solidification behaviour and predicted defect tendency of an investment-cast aerospace component. The work was based on casting simulation of an IN718 component, where the results were interpreted through temperature history, heat-flux response, shell representation, filling and solidification behaviour, and porosity-related defect indicators. The purpose was not to define an approved production setup, but to improve understanding of how the simulated casting response changes when the thermal conditions around the shell and mould system are varied.

The main conclusion is that the casting response was governed by the combined thermal history of the metal, ceramic shell, insulation and surrounding enclosure. The ceramic shell should therefore not be interpreted only as a geometric boundary. It acts as a part of the thermal system and influences how heat is extracted from the melt, how solidification progresses and how favourable the feeding conditions become during the late stages of solidification. For IN718, this is important because the alloy solidifies over a temperature interval and remains sensitive to feeding conditions, last-to-freeze regions and shrinkage-related defect formation.

Insulation wrapping had a strong influence on the thermal response of the model. The wrapped and unwrapped cases differed more than the two wrapped shell representations did. This means that insulation wrapping was a controlling part of the simulation setup and that the wrapped cases were the most relevant basis for the main sensitivity study. The results also show that simplified shell modelling without the surrounding thermal boundary conditions may give an incomplete interpretation of the casting behaviour.

Among the investigated process parameters, shell preheat temperature had the clearest and strongest influence on the evaluated outputs. The higher shell-preheat cases produced the most favourable porosity indications, while the lower shell-preheat cases produced the least favourable indications. The standard shell-temperature cases gave an intermediate response. This trend is physically reasonable. A colder shell increases the temperature difference between the molten metal and the mould during filling and early solidification, which increases heat extraction from the melt. This can promote premature local freezing, reduce the ability of the remaining liquid to feed solidification shrinkage and produce less favourable porosity indications. A hotter shell reduces the initial thermal driving force and can therefore support a

more favourable feeding condition during the critical solidification interval.

The mass-flow-rate variation had a smaller influence than shell preheat temperature. The lower mass-flow condition was generally slightly worse, especially when combined with the lower shell-preheat condition. This indicates that mass flow affected the thermal state established during filling, but it did not dominate the later solidification response in the same way as the shell preheat temperature. The most adverse investigated condition was therefore the combination of low shell preheat temperature and lower measured-average mass flow.

The enclosure emissivity variation had negligible influence on the evaluated results in this simulation setup. Temperature histories, heat-flux histories and defect-indicator outputs were very similar between the standard-emissivity and lower-emissivity cases. This does not mean that radiation is generally unimportant in investment casting. Instead, it means that, for the investigated model and parameter range, emissivity was not a controlling variable compared with shell preheat temperature, insulation wrapping, mass flow and shell representation.

The comparison between constant and varying shells showed that shell representation can affect the thermal response and the predicted defect tendency. The constant shell representation gave slightly more favourable defect-indicator results on average than the non-uniform representation. However, the difference was smaller than the effect of shell preheat temperature and it should not be interpreted as meaning that the constant shell is more realistic. The non-uniform shell representation remains important because real ceramic shells are not perfectly uniform. Local shell-thickness variation can influence heat extraction and should be included when the objective is to represent the industrial casting process as realistically as possible.

The temperature–heat-flux analysis supported the trends observed in the time-based plots and defect indicators. Cases with less favourable porosity indications also showed thermal behaviour consistent with stronger or earlier heat extraction during critical parts of filling and solidification. This supports the interpretation that the defect indicators were linked to the thermal development of the casting rather than being isolated numerical outputs. The results should therefore be read as a connected thermal and feeding problem.

A limited qualitative comparison with available process observations showed that the simulated defect-indicator regions were somewhat consistent with observed porosity indications. This provides support that the simulations captured relevant trends in the casting behaviour. However, the comparison was not extensive enough to validate the model quantitatively. The results should therefore be treated primarily as simulation-based sensitivity trends, not as verified predictions of exact porosity size, amount or acceptance.

The conclusions from this study can be summarised as follows:

- The thermal boundary conditions around the shell, especially insulation wrapping, had a major influence on the simulated casting response.
- Shell preheat temperature was the most influential investigated process parameter.
- Higher shell preheat temperature gave more favourable porosity-related indications, while lower shell preheat temperature gave clearly worse indications.
- Mass flow rate had a smaller but still compared to emissivity still somewhat visible influence, particularly when combined with lower shell preheat temperature.
- Constant and non-uniform shell representations produced different local responses, but the difference was smaller than the influence of shell preheat temperature.
- The simulated defect indicators were broadly consistent with the thermal-history interpretation and showed some qualitative agreement with available process observations.
- The results support the use of casting simulation as a comparative tool for process understanding, but not as a stand-alone validation of production quality.

The work also shows that shell-thickness effects should be interpreted together with material behaviour. For IN718, changes in heat extraction influence the solidification sequence, feeding conditions and the risk of shrinkage-related defects. The shell representation is therefore not only a mould-modelling choice but also a material-quality issue. A shell condition that appears acceptable from one output, such as faster cooling, may still not be if it freezes feeding paths too early or increases the risk of isolated last-to-freeze regions.

For process development, the results suggest that control and modelling of shell preheat temperature should be prioritised. The shell thermal state should be represented carefully in future simulations, and measured or process-based shell-thickness information should be used whenever available. Since the emissivity variation showed little influence in this study, further effort should first be directed toward parameters with stronger demonstrated effects, such as shell preheat, shell representation, filling conditions and interface-related assumptions.

The main limitation of the conclusions is that they are based on a specific component geometry, simulation setup, material dataset and selected process-parameter range. The findings should therefore not be transferred directly to other investment castings without additional analysis. The model also remains limited by assumptions regarding material properties, shell properties, boundary conditions, interface heat transfer, shell representation and porosity post-processing. Since the comparison with observed porosity indications was qualitative, the study cannot define a final approved production setup or a quantitative acceptance criterion.

6.1 Recommendations for future work

Several areas are recommended for future work. The first recommendation is to improve the quality and completeness of the industrial input data used in the simulation model. Several process conditions had to be represented using available or assumed values. Future work should therefore focus on obtaining better information about the actual thermal conditions in the vacuum furnace, the temperature history of the ceramic shell before and during casting, and the thermal properties of the shell system. More detailed studies of the actual mass-flow conditions during filling could also be valuable depending on how strongly they affect the case at hand.

A second recommendation is to characterise the real shell geometry in more detail. The comparison between constant and non-uniform shell representations showed that shell geometry can influence the local thermal response, even if its effect was smaller than the influence of shell preheat temperature in the investigated cases. Future studies should therefore include scanning or measurement of manufactured shells, for example using GOM-scanned data or sectioning, to obtain a more accurate description of local shell-thickness variation. These measurements could then be used to study how real shell geometry affects heat extraction, solidification behaviour, predicted defect locations and ultimately the manufactured cast component.

A third recommendation is to expand the experimental validation. In this work, the comparison with process observations was limited and mainly based on available porosity indications. Future validation should include more types of experimental data, not only additional porosity inspection. Useful validation data would include measured temperature histories, shell-temperature measurements, filling observations where possible, dimensional inspection, defect locations from non-destructive testing, metallographic sections and repeated casting trials. This would make it possible to evaluate whether the simulation captures both the qualitative trends and the quantitative behaviour of the real casting process.

Further work should also include a broader sensitivity study of both process parameters and simulation parameters. The present work focused on selected process-related variables such as shell representation, shell preheat temperature, mass flow rate and enclosure emissivity. Future studies should also investigate how numerical and modelling choices influence the results. Examples include mesh resolution, time-step settings, interface heat-transfer coefficients, radiation assumptions, solver tolerances, material-property datasets and porosity-model settings. This would help separate physical process effects from numerical modelling effects. It would also make it possible to evaluate the balance between simulation accuracy, robustness and computational time.

Finally, a longer-term development step would be to investigate whether the simulation workflow can be extended towards predicting material properties in the finished cast component. This is likely to require substantially more data and a stronger understanding of the specific industrial process than was available in the present

work. Such a study would need reliable thermal histories, validated material models, experimental measurements of microstructure and mechanical properties, and a clear connection between simulated solidification behaviour and final material performance. Although this is further away than the other recommendations, it would be a valuable direction because it would move the simulation work from predicting defect tendency towards predicting the properties of the manufactured part itself.

6.2 Final words

Overall, the thesis demonstrates that the simulated quality response of the investigated IN718 investment casting is controlled by the interaction between shell representation, shell thermal state, filling conditions and solidification behaviour. The most important practical outcome is that shell preheat temperature and the surrounding thermal boundary conditions had a larger influence than the difference between the constant and non-uniform shell representations in the investigated cases. The work therefore contributes a simulation-based understanding of which parameters should be prioritised when developing and interpreting investment-casting procedures for the studied component.

Bibliography

- [1] Altair Engineering. “Result Types: Niyama Criterion and Micro Porosity.” *Altair Inspire Cast 2022.3 Help*, 2022, https://2022.help.altair.com/2022.3/cast/en_us/topics/cast/analysis/result_types_r.htm. Accessed 17 March 2026.
- [2] Alvandi, A., et al. “Effect of Ceramic Shell Thickness on Investment Casting Defects Using Computer Simulation.” *World Academics Journal of Engineering Sciences*, vol. 12, no. 2, 2025, pp. 10–15, https://www.isroset.org/journal/WAJES/full_paper_view.php?paper_id=3881.
- [3] Behera, M. M., et al. (2019) Thermo-Mechanical Analysis of Investment Casting Ceramic Shell: A Case Study. *Measurement*, 147, 106805. <https://doi.org/10.1016/j.measurement.2019.07.033>.
- [4] Campbell, J. (2015) *Complete Casting Handbook: Metal Casting Processes, Metallurgy, Techniques and Design*. 2nd ed. Elsevier.
- [5] Carlson, K. D. and Beckermann, C. (2009) Prediction of Shrinkage Pore Volume Fraction Using a Dimensionless Niyama Criterion. *Metallurgical and Materials Transactions A*, 40, pp. 163–175. <https://doi.org/10.1007/s11661-008-9715-y>.
- [6] Carlson, Kent D., and Christoph Beckermann. “Use of the Niyama Criterion to Predict Shrinkage-Related Leaks in High-Nickel Steel and Nickel-Based Alloy Castings.” *Proceedings of the 62nd SFSA Technical and Operating Conference*, paper no. 5.6, Steel Founders’ Society of America, 2008.
- [7] Chvorinov, N. (1940) Theorie der Erstarrung von Gussstuecken. *Giesserei*, 27, pp. 201–208.
- [8] Dong, R.-Z., et al. (2024) An Investigation of Ceramic Shell Thickness Uniformity and Its Impact on Precision in Turbine Blade Investment Casting. *Journal of Manufacturing Processes*, 131, pp. 507–522. <https://doi.org/10.1016/j.jmapro.2024.09.006>.
- [9] Dong, Yiwei, et al. “Determination of Interfacial Heat-Transfer Coefficient during Investment-Casting Process of Single-Crystal Blades.” *Journal of Materials Processing Technology*, vol. 211, no. 12, 2011, pp. 2123–2131, <https://doi.org/10.1016/j.jmatprotec.2011.07.012>.
- [10] Flow Science. “Prediction of Shrinkage Defects during Investment Casting Process.” *FLOW-3D*, 2016, <https://www.flow3d.com/prediction-shrinkage-defects-investment-casting-process/>. Accessed 12 March 2026.

- [11] Guo, J., et al. (2015) Microporosity Prediction and Validation for Ni-Based Superalloy Castings. *IOP Conference Series: Materials Science and Engineering*, 84(1), 012003. <https://doi.org/10.1088/1757-899X/84/1/012003>.
- [12] The Investment Casting institute (2017). The investment casting process. https://www.investmentcasting.org/uploads/8/1/9/8/81988734/the_investment_casting_process_-_website.pdf
- [13] Jones, S. and Yuan, C. (2003) Advances in Shell Moulding for Investment Casting. *Journal of Materials Processing Technology*, 135(2–3), pp. 258–265. [https://doi.org/10.1016/S0924-0136\(02\)00907-X](https://doi.org/10.1016/S0924-0136(02)00907-X).
- [14] Keysight Technologies. “Casting (ProCAST).” *Computer-Aided Engineering Software*, <https://www.keysight.com/zz/en/products/design-engineering-software/computer-aided-engineering-software/casting.html>. Accessed 30 August 2026.
- [15] Kanyo, J. E., et al. (2020) An Overview of Ceramic Molds for Investment Casting of Nickel Superalloys. *Journal of the European Ceramic Society*, 40(15), pp. 4955–4973. <https://doi.org/10.1016/j.jeurceramsoc.2020.07.013>.
- [16] Manjabacas, M. C. and Miguel, V. (2023) A Methodology to Define the Niyama Criterion Reinforced with the Solid Fraction Analysis: Application to Sand Casting of Steel Bars. *Metals*, 13(10), 1777. <https://doi.org/10.3390/met13101777>.
- [17] Mitchell, A. and Wang, T. (2001) Solidification and Precipitation in IN718. In: Loria, E. A. (ed.) *Superalloys 718, 625, 706 and Various Derivatives*. The Minerals, Metals & Materials Society, pp. 81–90. https://www.tms.org/Superalloys/10.7449/2001/Superalloys_2001_81_90.pdf.
- [18] Niyama, E., Uchida, T., Morikawa, M. and Saito, S. (1982) A Method of Shrinkage Prediction and Its Application to Steel Casting Practice. *AFS International Cast Metals Journal*, 7(3), pp. 52–63.
- [19] Rafique, M. M. A. and Iqbal, J. (2009) Modeling and Simulation of Heat Transfer Phenomena during Investment Casting. *International Journal of Heat and Mass Transfer*, 52(7–8), pp. 2132–2139. <https://doi.org/10.1016/j.ijheatmasstransfer.2008.11.007>.
- [20] Rakoczy, L., and R. Cygan. “Analysis of Temperature Distribution in Shell Mould during Thin-Wall Superalloy Casting and Its Effect on the Resultant Microstructure.” *Archives of Civil and Mechanical Engineering*, vol. 18, 2018, pp. 1441–1450, <https://doi.org/10.1016/j.acme.2018.05.008>.
- [21] Raza, M., et al. (2015) The Effect of Shell Thickness, Insulation and Casting Temperature on Defects Formation during Investment Casting of Ni-Base Turbine Blades. *Archives of Foundry Engineering*, 15(4), pp. 115–123. <https://doi.org/10.1515/afe-2015-0090>.
- [22] Shi, X., et al. (2018) Solidification and Segregation Behaviors of Superalloy IN718 at a Slow Cooling Rate. *Materials*, 11(12), 2398. <https://doi.org/10.3390/ma11122398>.
- [23] Stefanescu, D. M. (2009) *Science and Engineering of Casting Solidification*. 2nd ed. Springer.
- [24] Szeliga, D. (2020) Manufacturing of Thin-Walled Ni-Based Superalloy Castings Using Alternative Thermal Insulating Module to Control Solidification. *Journal*

- of Materials Processing Technology*, 278, 116503. <https://doi.org/10.1016/j.jmatprotec.2019.116503>.
- [25] Szeliga, D., et al. (2017) Investigation of Casting-Ceramic Shell Mold Interface Thermal Resistance during Solidification Process of Nickel Based Superalloy. *Experimental Thermal and Fluid Science*, 87, pp. 149–160. <https://doi.org/10.1016/j.expthermflusci.2017.04.024>.
- [26] Torroba, A. J., et al. (2014) Investment Casting of Nozzle Guide Vanes from Nickel-Based Superalloys: Part I—Thermal Calibration and Porosity Prediction. *Integrating Materials and Manufacturing Innovation*, 3, pp. 344–368. <https://doi.org/10.1186/s40192-014-0025-5>.
- [27] Torroba, A. J., et al. (2014) Investment Casting of Nozzle Guide Vanes from Nickel-Based Superalloys: Part II—Grain Structure Prediction. *Integrating Materials and Manufacturing Innovation*, 3, pp. 369–384. <https://doi.org/10.1186/s40192-014-0026-4>.
- [28] Upadhy, G. K., et al. (1995) Modelling the Investment Casting Process: A Novel Approach for View Factor Calculations and Defect Prediction. *Applied Mathematical Modelling*, 19(6), pp. 354–362. [https://doi.org/10.1016/0307-904X\(95\)90001-0](https://doi.org/10.1016/0307-904X(95)90001-0).
- [29] Yeh, C.-P., Hwang, W.-S. and Lin, C.-H. (2008) Process Design Optimization through Numerical Experimentation for a Brake Disc Casting. *Materials Transactions*, 49(6), pp. 1372–1379. <https://doi.org/10.2320/matertrans.MRA2007233>.



CHALMERS
UNIVERSITY OF TECHNOLOGY