Department of Precision and Microsystems Engineering

Accurate Modelling of Conjugate Heat Transfer for Density-Based Topology Optimisation

Sam Cornelius Adrianus Pijnenburg

Report no : 2025.077
Coaches : Ir. M. Slebioda, Ir. R. J. P. van der Nolle
Professor : Prof. dr. ir. M. Langelaar
Specialisation : Computational Design and Mechanics (CDM)
Type of report : MSc Thesis
Date : November 18, 2025

Accurate Modelling of Conjugate Heat Transfer for Density-Based Topology Optimisation

by

Sam C. A. Pijnenburg

to obtain the degree of Master of Science
at the Delft University of Technology faculty of Mechanical Engineering,
to be defended publicly on Thursday 4 December 2025 at 14:00.

Student number:	4962516	
Project duration:	9 September 2024 — 4 December 2025	
Thesis committee:	Prof. dr. ir. M. Langelaar,	TU Delft, supervisor
	Dr. ir. L. F. P. Noel,	TU Delft
	Ir. M. Slebioda,	TU Delft, PhD supervisor
	Ir. R. J. P. van der Nolle,	ASMPT ALSI B.V.
Submitted on:	18 November 2025	
Document version:	1.2.5 – Committee rectification	

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TU Delft ASMPT

Preface

In front of you is the master thesis Accurate Modelling of Conjugate Heat Transfer for Density-Based Topology Optimisation: TopOpt for Thermofluid design problems, *the culmination of twelve months of work in studying the basics of topology optimisation, multiphysics topology optimisation, in particular, implementing various techniques from literature into my own simulations and optimisation models and trial and error. In the end, this thesis project has proven to be a great challenge in pursuit of satisfactory results. It also has turned out to be a true test of my creativity, patience and sanity.*

A special thanks goes out to prof. dr. ir. M. Langelaar and ir. M. Slebioda for their continued efforts to help me throughout this past year to the best of their abilities, despite my struggling to explain whatever new techniques I was trying to devise and implement.

I also extend my gratitude to dr. ir. M. J. B. Theulings for enabling me to read, consider and implement his at-the-time unpublished research. Any of the accomplishments in this thesis, frankly, would not have been achievable without his work as a foundation and as a source of inspiration.

Furthermore, I would like to thank Rens and everyone else at ASMPT's CoC in Beuningen, The Netherlands, for their support throughout the project and for serving as rubber duckies over the past twelve months. In particular, as promised, a 'thank you' goes out to Jasper for his contributions to and feedback on the fill patterns of the x-bar data plots featured in chapter 4.

Sam C. A. Pijnenburg
Leusden, November 2025

Abstract

In the high-tech engineering sector, industry is always looking for the competitive, technological edge. Through optimisation, in particular the optimisation of component designs, performance gains can often be realised compared to conventional design geometries. Even within existing implementations, this so-called topology optimisation allows for improvements in component performance, simply by exchanging its existing conventional design, and ultimately to higher-level machine assemblies and modules. This low barrier to entry makes topology optimisation a field of engineering that has been gaining traction over the past decades within various fields of high-tech engineering and, of course, research. While much research has been done into the field of topology optimisation, including research into multi-physics optimisation problems, such as conjugate heat transfer problems, more work is still to be done to improve models, improve optimisation schemes and approaches, reduce computational time, increasing results accuracy and much more.

This thesis, in particular, is focussed on devising a strategy to derive and implement a thermofluidic model specifically for density-based topology optimisation applications. In doing so, emphasis is placed on the accuracy of the optimisation model when compared to numerical results found in regular thermofluidic analyses of systems. In this thesis, a new technique is introduced to refine existing fluidic topology optimisation models for density-based methods: dubbed IGPP, Implicit Gradient-Parallel Penalisation for fluid velocity fields is devised, implemented and evaluated numerically. Furthermore, a parameterisation is created for the material properties relevant to conjugate heat transfer problems. Finally, the model's performance is evaluated.

Keywords: density-based topology optimisation, thermofluidic optimisation, conjugate heat transfer, modelling accuracy, IGPP.

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

Abbreviation	Definition
CFD	Computational Fluid Dynamics
CHT	Conjugate Heat Transfer
D(-penalisation)	Darcy-only fluid permeation penalisation
DF(-penalisat.)	Darcy-with-Forchheimer fluid permeation penalisation
DFF(-penalisat.)	Darcy-with-filtered-Forchheimer fluid permeation penalisation
GCMMA	Globally convergent version of MMA
IGPP	Implicit Gradient-Parallel (flow) Penalisation
KC	Kays-Crawford (model)
LSM	Level-Set Method
MMA	Method of Moving Asymptotes
NS (equations)	Navier-Stokes momentum and mass continuity equations for fluids
OC	Optimality Criteria method
PAMP	Porous Anisotropic Material with Penalisation approach
RAMP	Rational Approximation of Material Properties approach
SIMP	Solid Isotropic Material with Penalisation approach
TF (problem)	Thermofluidic (optimisation problem)
TO	Topology Optimisation

Symbols

Symbol	Definition
D	Diameter
$\mathcal{D}$, $\mathcal{D}^{\mathrm{glob}}$	Truly optimal design, truly globally optimal design
$\mathcal{F}$	Objective function
$\mathcal{G}$	Constraint function
T, T_{ref}	Temperature field, reference temperature (293.15 K)
V_0	Maximum (solid) volume
Da	Darcy number
Gz	Graetz number
Ma	Mach number
Nu	Nusselt number
Pr, $\mathrm{Pr_T}$	Prandtl number, turbulent Prandtl number
Re, $\mathrm{Re_L}$	Reynolds number, Reynolds number for characteristic length-scale L
c_p	Material (linear) heat capacity at constant pressure
i, j, k	Generic (tensor) indices
p	Fluid pressure field

Symbol	Definition
t	Time
u, v, w	Fluid velocity field x-, y-, z-components, respectively
x, y, z	Spatial coordinates
$\dot{\mathbf{q}}$	Heat flux field
$\mathbf{u}$	Fluid velocity field
Γ	Boundary
Γ^{sf}	Solid-fluid interface boundary
Ω	Design domain, simulation domain
Ω^{f}	Fluid material-filled sub-domain of design domain
Ω^{intf}	Solid-fluid interface / intermediate pseudo-density-filled sub-domain
Ω^{s}	Solid material-filled sub-domain of design domain
γ, $\gamma(\mathbf{x})$, γ_{f}, γ_{s}	Element density, density distribution, fluid and solid element density
ζ	Fluid dynamic bulk viscosity
κ	Thermal conductivity
μ	Fluid dynamic shear viscosity
ρ	Physical material density
$\rho^{\star}$	Non-physical material density for flow field solution
$\boldsymbol{\gamma}$	Design variable vector
$\boldsymbol{\psi}$	(Generic) state field vector
$\tilde{\square}$	Filtered variable / field
$\hat{\square}$	Projected variable / field
$\mathfrak{Z}$	Bourbaki caution symbol for complex arguments – found in page margin

1

Introduction

In pursuit of ever faster and more accurate machinery to satisfy the demand for ever higher-quality products, high-tech industries, such as the automotive, aerospace and electronics and semiconductor sectors, seek to research new technologies that enable the designing, manufacturing and operation of better-performing machines, subsystems and parts.

Knowing that topology optimisation might be a valuable tool in trying to further increase the performance of machinery, research into its potential for the optimisation of complex integrated parts within their specific applications, such as those for motion stage design, thermal management systems or drive systems, is desirable.

In many applications, the topology optimisation of parts that are subject to (a subset of) thermal, fluidic and even structural physics are of interest. Such design problems may include an actively-cooled sleeve for an electric motor to minimise thermal resistance of the motor internals to ambient temperature or heat exchanger plates for electric motor DC-AC inverters, allowing it to run cool while increasing mechanical power output. Alternatively, heat insulation may be of interest, meaning structures are to be optimised to maximise thermal resistance in some regions in the design domain, and minimised in others by the introduction of well-placed cooling channels.

As outlined by Alexandersen and Andreasen (2020), in the domain of thermofluidic topology optimisation using density-based methods, gains are still to be made in terms of the accuracy of the optimisation models. Moreover, this is the case even in the field of purely fluidic topology optimisation. Considering their recommendations, this is what this thesis seeks to achieve: a fairly robust description of the thermofluidic problem for a density-based topology optimisation problem. Through a better understanding of the accuracy of the model used, ideally, a better understanding of the accuracy and optimality of generated designs can be created and better-performing designs may be realised.

This research will focus on making improvements to topology optimisation for thermofluid-related implementations.

To further illustrate the underlying problem, in figure 1.1, a schematic representation of the design variable field $\hat{\gamma}$ is given. The core problem facing modelling accuracy in thermofluidic, density-based topology optimisation lies in the various steps required to ensure a generated design is realisable. In order to achieve this, a minimum feature size requirement should be imposed, to avoid structures that cannot be produced with existing manufacturing techniques. In doing so, however, a filtering step is applied to the design field, that blurs the regions between solid and fluid, causing intermediate-pseudodensity regions as seen in the figure.

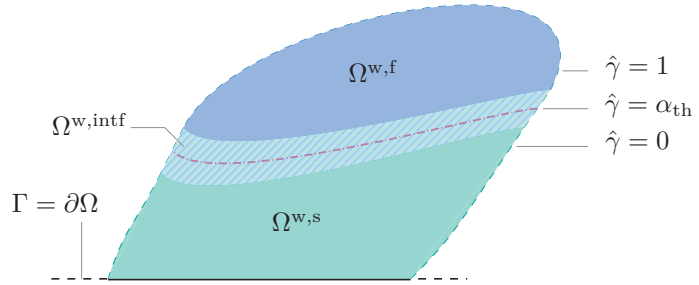

Figure 1.1: *Schematic illustration of a selected subdomain Ω^{w} in Ω utilising a density-based approach. Independently represented are a fluid region ($\Omega^{\mathrm{w,f}}$), a solid region ($\Omega^{\mathrm{w,s}}$) and an intermediate-density interface region ($\Omega^{\mathrm{w,intf}}$). $\hat{\gamma}$ represents the projected pseudodensity, or design variable field. Note that in simulation, the domain Ω is discretised.*

It is through this blurry interface that not only any information on the 'real' position of the designed interface is lost, but also any material properties take on non-physical intermediate values between solid and fluid. Considering that a good representation of fluid-solid interfaces is critical in accurate numerical modelling of even regular heat transfer problems, this blurry interface becomes exceptionally problematic as the relative effects of heat conduction and advection become hard to model accurately with the intermediate material property values found in topology optimisation. Moreover, also meshing the domain to refine the regions where boundary layers are formed at the solid walls, is problematic for density-based optimisation methods, further increasing inaccuracy of the simulation.

After a design has been generated, it will need to be discretised into a solid section and a fluid section. This discretisation is realised by splitting the entire design domain Ω at the threshold value α_{th} of $\hat{\gamma}$. This process of discretising the design using the threshold contour surfaces, is hereafter referred to as "thresholding". The design generated after thresholding a continuous-density design will be referred to as the discrete-equivalent design.

1.1. Research scope & focus

For this master thesis, some limitations on its scope are set. These limitations are imposed on both the considered effects and physical phenomena that are considered as well as some simplifications and working assumptions about the physics involved. With thermal management as one of several key design considerations and challenges in high-tech equipment, topology optimisation applied to design elements subject to thermofluidic physics, forms the focus of this thesis. This technology could significantly improve cooling performance and ultimately the overall performance of machinery and products and is therefore of interest. As a result of this focus, this research considers heat transfer, in the form of conduction and thermofluidics, and fluid dynamics in topology optimisation problems.

Disregarded physics include structural mechanics, including thermomechanics and buckling mechanics, as well as thermal radiation, unsteady and quasi-turbulent flows from the RANS-equations, anisotropic or temperature-dependent material properties (thereby also excluding natural convection effects) and material phase transitions. Furthermore, some other key working assumptions are that the considered fluid is Newtonian and that solid material permeability in reality is zero. The commercially available COMSOL Multiphysics® 6.2 software will be used for the majority of simulation and optimisation work presented in this thesis.

Within the outlined scope, the following main research question of this thesis is distilled:

"How can thermal and fluidic requirements be combined into a sufficiently accurate multidisciplinary topology optimisation method?"

Listed here, are the most relevant research sub-questions to this thesis, devised to support the main research question:

- *What multiphysics topology optimisation approaches exist for this scope and how do they, generally, work?*
- *What issues are found in the research gap between the existing work and the desired outcome of this thesis?*
- *How can the individual identified issues be solved, mitigated or even worked around to generate satisfactory results?*
- *How does the proposed approach perform on reference problems and optimised designs?*

Very relevant to this thesis' research question is the topic of how accurate "sufficiently accurate" is, though it is also quite arbitrary. For this research, an error of 10% across the various simulations would be deemed a successful improvement, and a consistent 5% would exceed expectations.

For context, currently the most accurate density-based optimisation models for conjugate heat transfer problems see somewhere near 15% errors between the optimisation result and its post-processed and realisable equivalent. This means that, considering the non-linearity of fluid dynamics problems in general, designs may well be optimised to far exceed 15% error in performance when compared to the truly optimal design. As such, a more accurate model not only brings the performance of generated designs closer to performance upon their realisation into a physical component or assembly, but also reduces the likelihood of generating topologies that perform inferiorly to those generated with a more accurate model, as it leads to better gradient information during the optimisation process.

Considering also that models used in the numerical analysis of conjugate heat transfer problems typically sees errors up to the order of 1%, a consistent 5% error between the optimisation and the numerical analysis models can be considered a very good result, especially in the light that the optimisation problem also realistically eliminates the use of boundary layer mesh elements as an option to enhance the accurate capture of heat transfer across solid-fluid interfaces.

1.2. Thesis structure

Following this introduction of the thesis and its scope, the literature review of Chapter 2 covers the state-of-the-art of topology optimisation research, with specific focus on thermofluidic system applications the outlined scope. Chapter 3 regards the methodology of implementation, used techniques and specific reference problems utilised in the evaluation of different optimiser and solver configurations as well as the specific design cases considered. Any identified problems from the methodology and the found research gap are covered in Chapter 4. This chapter includes the use of aforementioned methods to evaluate the performance of the newly devised implicit gradient-parallel penalisation scheme (IGPP) for fluids in permeability-based topology optimisation. It also concerns the ultimate construction and parameterisation of the thermofluidic model used in optimisation.

Chapter 5 deals with results of model validation and verification reference simulations and optimisations as well as the results of the specific design cases considered. Any observations based off of these results are discussed in Chapter 6, before the final conclusions on the posed research questions and remaining remarks in Chapter 7.

2

Literature review

2.1. Existing relevant work in topology optimisation

2.1.1. General techniques in topology optimisation

This section will cover the existing research on implementation of topology optimisation utilising density-based methods. It pertains to work specific to the density-based method, such as material properties interpolation and filtering methods, as well as suitable approaches in design iteration and optimisation algorithms.

Techniques for density-based methods

Two-material density-based methods for topology optimisation rely on the discretisation of the design domain Ω into (volume) elements and assigning to each a continuous pseudodensity between 0 and 1. This pseudodensity resembles the designed structure, in that 0 and 1 refer to two different materials or material phases. In structural problems these often represent 'void' and 'solid' material elements and for thermofluidic optimisation problems, they generally represent 'solid' and 'fluid' material elements, respectively. This thesis will follow this convention.

Density-based methods, since their inception by Bendsøe and Kikuchi (1988) in the pioneering paper on topology optimisation, have seen wide-spread application. Despite some complications in early years, such as mesh-dependencies, checkerboarding and grayscaling, this method has evolved to generate reliable results for a broad range of different applications. As discussed in papers akin to Zhou, Shyy, and Thomas (2001), these issues have been addressed through several techniques, including filtering and projection methods applied to the pseudodensity variable field. While early continuous density methods utilised rather more rudimentary approaches to solve these problems, such as gradient control and sensitivity filtering, more recent work relies on somewhat more involved methods like three-field methods that include Helmholtz PDE filtering and projection steps to also realise effective minimum feature size control. Specifically works like Guest, Prévost, and Belytschko (2004), Guest (2009) and Guest, Asadpoure, and Ha (2011) have contributed significantly to the development towards effective use of these methods.

This thesis will make use of Helmholtz PDE filtering and subsequent hyperbolic tangent projection of the design variable field.

Techniques for design iterations and sensitivity analysis

For continuous pseudodensity methods, the most commonly used optimisation algorithm is the Method of Moving Asymptotes, or MMA. Svanberg (1987) devised this method and it was later revised to ensure global convergence, the method named GCMMA, by Zillober (1993).

Since this method relies on gradient information, a sensitivity analysis of the problem has to be performed. Since the governing equations consist of the non-linear Navier-Stokes partial differential equations for fluid flow and the diffusion-advection partial differential equation for heat transfer, analytic sensitivities are hard to come by. In this thesis, then, the adjoint state method will be used to find the numerical values of the sensitivities of the objective and constraints to the design variable field.

Material properties interpolation

Since the projected design variable field forms the definition of the design and uses the projected pseudo-density $\hat{\gamma}(\mathbf{x}) \in [0, 1]$, the material properties of an element should be specified as a function of this variable. Several common types of interpolation functions exist, such as the Solid Isotropic Material with Penalisation (SIMP) interpolation functions and the Rational Approximation of Material Properties (RAMP) interpolation functions.

For this research, any interpolation functions for material properties will be of the RAMP-type unless explicitly differently defined in Chapter 4. RAMP, after its introduction by Stolpe and Svanberg (2001), has become the interpolation scheme of choice in research for multidisciplinary topology optimisation, on account of the non-zero gradients at both 0- and 1-values of the projected pseudodensity variable.

2.1.2. Current techniques for topology optimisation of thermofluidic problems

Existing fluidic topology optimisation research

On laminar flow conditions in fluidic topology optimisation – akin to those that will be considered in this thesis – Gersborg-Hansen, Sigmund, and Haber (2005) published the first work to consider the full steady-state, incompressible Navier-Stokes equations as governing equations. Following their paper on two-dimensional optimisation problems, fluidic topology optimisation was researched heavily, specifically with level-set methods, as opposed to density-based methods. Notable papers include the progress made by Duan and others in 2008: Duan, Ma, and R. Zhang (2008a), Duan, Ma, and R. Zhang (2008c), Duan, Ma, and R. Zhang (2008b).

However, since this thesis will make use of and looks to make improvements to the continuous density method, improvements to the Darcy penalisation scheme for density-based methods by Theulings, Langelaar, et al. (2023) are of particular interest. Moreover, the paper by Theulings, Noël, et al. (2025) forms the basis upon which this research is built. It introduced two new penalisation methods with a new term to realise more effective fluid permeation penalisation. The introduced methods enable enhanced control over in-solid fluid velocities and, consequently, fluid leakage throughout the optimisation domain to yield more realistic flow field results in the fluid solver. Since their work is principal for this thesis, it is further covered throughout section 4.1.

Existing thermal topology optimisation research

In the realm of thermal physics-only topology optimisation, research is limited to almost only conduction-based problems. Lohan, Dede, and Allison (2019) performed a range of different heat conduction topology optimisation studies, evaluating various types of objective functions and design cases for these problems. More complex problems have also been researched, including transient problems by Wu, Y. Zhang, and Liu (2019), with Li et al. (2022) even introducing multi-material optimisation into their transient problem.

Since conjugate heat transfer (thermofluidic) problems, such as those considered in this thesis, are rather sensitive to the interpolation of the thermal conductivity in generating meaningful (realistic) results, the specific interpolation functions presented in these papers are likely not suitable. This perspective is supported by the remarks of Alexandersen and Andreasen (2020) on the need for further research into modelling accuracy for thermofluidic topology optimisation, in particular, as well as the fact that, since their publication, little emphasis has been placed on this in literature.

Existing thermofluidic topology optimisation research

Work on conjugate heat transfer in topology optimisation was picked up near 2007, arguably started in the paper by Bruns (2007). However, in their work, the flow field was not calculated. Instead using a fixed coefficient for heat advection. Yoon (2010), instead, included a flow field solver in their optimisation. They used a similar staggered, or segregated, solution method to the one that will be employed in this thesis, where first the flow fields are solved, and the temperature field is calculated after, since all (including fluid simulation-related) material properties are temperature-independent.

Matsumori et al. (2013) improved on this by also successfully modelling the solid material domain temperature field under some constant input power, enabling more relevant constraints to be placed on the problem. However, their implementations lacked analysis of the heat transfer accuracy and they clearly demonstrated that results were mesh-dependent and that grayscaling occurred. Conversely, Yaji et al. (2015) achieved crisp results, albeit using a level-set method, for three-dimensional problems.

Later, Haertel and Nellis (2017) provided an improvement to density-based conjugate heat transfer topology optimisation, allowing for arbitrary, though non-turbulent (Re smaller than 2300), flow conditions under steady-state conditions. While they provided some comparison between numerical simulation and known analytical solutions, they provided no comparison between a numerical, discrete-density result and an equivalent continuous-density result to quantify the effects of the continuous design variable field on the accuracy of the optimisation problem.

At this point, the use of turbulent flow models in topology optimisation models also gained in popularity, shifting research focus. For instance, Dilgen et al. (2018) implemented a density-based approach for a turbulent conjugate heat transfer problem with around 5 million degrees of freedom, using the k-ω-turbulence model. Then, in 2022, L. Noël and Maute (2022) also considered conjugate heat transfer topology optimisation, utilising a level-set method. Meanwhile, Holka et al. (2022) used a density-based method with RAMP-interpolation instead and used it to optimise a surface cooler design. From this point forward, most research has been focussed on the extension of the possible physics and more exotic applications, both of which are well beyond the scope of this thesis. Nevertheless, plenty of research can still be conducted into improved methods and implementations for even relatively simple pure fluid and thermal topology optimisation problems.

In this research, the interpolation functions and further implementation of the thermofluidic model as described by Dilgen et al. (2018) will be used as a benchmark for the 'state-of-the-art' in optimisation model accuracy.

2.2. Identified research gap

The research gap between the desired outcome of this thesis and literature has been reduced to several key issues and considerations. The need for a focus on accurate modelling in topology optimisation problems follows largely from Alexandersen and Andreasen (2020) and the following reasoning.

Optimisation of any kind requires a sufficiently true-to-reality optimiser model for the results to be meaningful in reality. An optimisation algorithm may successfully optimise for some model with a large discrepancy to reality, but will result in an output that is in turn non-optimal in actuality. Consequently, also in thermofluidic topology optimisation, there is a need to ensure the optimiser model's realism. To achieve this, interpolation functions and penalisation schemes should be implemented in such a way as to result in an optimiser model that is sufficiently similar to reality or, at least, to known analytical solutions or to numerical solutions of reference problems. This, in turn, also requires the design of adequate reference problems to be devised to verify and validate the model. Furthermore, the resulting model should be verified to be suitable for optimisation.

A research gap is also found in the independent validation and implementation of the novel work by Theulings, Noël, et al. (2025), that shows promise in the effective control over flow penalisation in solid regions.

2.3. Update to research questions

Based off the conducted literature review, the originally posed research sub-questions may be in part answered and revised:

- *What multiphysics topology optimisation approaches exist for this scope and how do they, generally, work?*

Based off of the literature review, this question may be considered answered: a selection of important works to the state of general topology optimisation research and techniques as well as work specifically focussed on fluid, thermal or thermofluidic optimisation have been high-lighted.

- *What issues are found in the research gap between the existing work and the desired outcome of this thesis?*

This question also can be considered answered by the previous section.

- *How can the individual identified issues be solved, mitigated or even worked around to generate satisfactory results?*

Some key problems have come to light throughout the review that would affect the modelling accuracy for any interpolated thermofluidic model in optimisation. Specifically, these issues include the effective penalisation of non-zero fluid velocities in regions that are designed solid and the construction of adequate interpolations to the various material properties that are relevant to heat transfer modelling, such as the thermal capacity and conductivity. To tackle the first issue, the approach of Theulings, Noël, et al. (2025) could provide a good baseline, while the second would require some further analysis of the governing equations at play.

From this third question, two new research questions arise, as well as the final sub-question that, so far, remains unanswered:

- *How may a sufficiently accurate fluid velocity field solution be achieved? Is the work of Theulings, Noël, et al. (2025) adequate, or are additional steps required?*
- *How can the various material properties be interpolated for density-based topology optimisation to achieve accurate temperature field results?*
- *How does the proposed approach perform on reference problems and optimised designs?*

3
Methodology and techniques

3.1. Topology optimisation description

Several problems, both two-dimensional and three-dimensional, will be considered and outlined here, including a description of the topology optimisation formulation as well as a description of the relevant physics.

3.1.1. Physics and governing equations

In every considered problem, optimisation requires the simulation of multiple state fields in the design space in order for the objective and constraints to be evaluated. These fields are the fluid pressure field p, the fluid velocity fields $\mathbf{u}$ and the temperature field T. Throughout the optimisation process, these fields will be solved numerically using COMSOL, following the governing equations and boundary conditions described in this section. The numerical solutions are all executed on mapped square element meshes following the finite element method.

Design variable fields

This thesis will make use of a three-field approach to define the 'designed' structure in the (simulation) domain Ω. In such a three-field approach, three scalar pseudodensity fields are used (with lower and upper limits 0 and 1) to define the design. First the control pseudodensity field $\gamma(\mathbf{x})$ will represent the state of the optimisation variables that capture the material properties assigned to every element in the simulation domain.

To avoid single-element-sized features to be designed and simulated and instead realise a minimum feature size in the problem, this pseudodensity field will be filtered using the Helmholtz PDE filter approach:

$$-R_\gamma^2 \nabla^2 \tilde{\gamma} + \tilde{\gamma} = \gamma. \tag{3.1}$$

In this formulation, $\tilde{\gamma}$ represents the filtered pseudodensity field, the second of the three fields considered in defining the optimiser-designed geometry. R_γ represents the filtering radius, that is responsible for the minimum feature size control. A larger value leads to a greater minimum feature size, but also to larger intermediate-pseudodensity regions in the domain. Finally, since the filtered field will have a significantly increased area of intermediate-pseudodensity regions in Ω when compared to γ, a projection step is applied to the filtered field to reduce the total size of intermediate density areas by forcing elements to take on values closer to 0 or 1. To this end, the projected pseudodensity field $\hat{\gamma}$ is defined:

$$\hat{\gamma} = \frac{\tanh\left(\beta(\tilde{\gamma} - a)\right) + \tanh\left(\beta a\right)}{\tanh\left(\beta(1 - a)\right) + \tanh\left(\beta a\right)}. \tag{3.2}$$

Throughout this thesis, a will be kept constant at 0.5. β_γ is a measure of the 'strength' with which the projection occurs, specifically for elements that have filtered pseudodensities close to 0.5. All simulations are performed on this projected field $\hat{\gamma}$ and the chain rule is used to define the design changes that should be made to γ during optimisation.

Fluid state fields governing equations

The fields p, resembling fluid pressure, and $\mathbf{u}$, resembling fluid velocity, are governed by the Navier-Stokes equations:

$$\rho^\star \left(\mathbf{u} \cdot \nabla\right) \mathbf{u} + \nabla p - \mu \nabla^2 \mathbf{u} - \left(\frac{1}{3}\mu + \zeta\right) \nabla \left(\nabla \cdot \mathbf{u}\right) - \mathbf{F} = \mathbf{0}, \tag{3.3a}$$

$$\nabla \cdot \left(\rho^\star \mathbf{u}\right) = 0. \tag{3.3b}$$

These specific formulations are the momentum and continuity equations, respectively, for compressible flow with invariant shear (μ) and bulk (ζ) viscosities. It will be assumed that the Stokes hypothesis holds: $\zeta = 0$. Furthermore, $\rho^\star$ denotes a non-physical material density interpolation function and $\mathbf{F}$ the fluid permeation penalisation.

While flow conditions, in reality and the scope of this thesis, are taken to be incompressible on account of the Mach number Ma satisfying $\mathrm{Ma} < 0.3$, the compressible flow description of the Navier-Stokes equation will prove necessary.

Temperature field governing equation

The temperature state field T is governed by the steady-state diffusion-advection equation:

$$\nabla \cdot \left(\kappa \nabla T\right) - \rho c_p \mathbf{u} \cdot \nabla T + \dot{q}_{\mathrm{gen}} = 0. \tag{3.4}$$

In this governing equation, κ denotes the material thermal conductivity field, ρ the material physical density field and c_p the material heat capacity field, all of which are interpolated from the projected pseudodensity field. $\dot{q}_{\mathrm{gen}}$ resembles an out-of-plane heat flux that may be enabled for two-dimensional problems.

Materials simulated and their real properties

The material that are sought to be simulated are aluminium alloy EN-AW 6082-T6, at $\hat{\gamma} = 0$, and water, for which $\hat{\gamma} = 1$. Table 3.1 contains the relevant properties of these materials pertaining to thermofluidic simulation. Unless specified otherwise, these properties will be used throughout this thesis for simulation and optimisation as the 'solid' and 'fluid' material properties associated with 0 and 1, respectively, in the projected pseudodensity interpolations.

Property		Solid γ^{s}	Fluid γ^{f}	Unit
Heat capacity	$c_p(\hat{\gamma})$	900	4184	$\mathrm{J\ kg^{-1}K^{-1}}$
Dynamic viscosity[1]	$\mu(\hat{\gamma})$	$\to \infty$	0.9	$10^{-3}\ \mathrm{Pa\ s}$
Thermal conductivity	$\kappa(\hat{\gamma})$	237	0.6	$\mathrm{W\ m^{-1}\ K^{-1}}$
Physical density	$\rho(\hat{\gamma})$	2700	1000	$\mathrm{kg\ m^{-3}}$

Table 3.1: *Overview of selected temperature-independent material properties.*

[1] Let it be understood that, in this thesis, μ will be kept constant at its 'fluid' value.

3.1.2. Optimisation problem formulations

Objective definition

For the considered thermofluidic problems, the objective is to minimise the effective thermal resistance of the design to the heat flux. For two-dimensional problems, heat will be 'generated' inside the domain to simulate influx into the design space from the out-of-plane direction. Assuming a steady-state solution and that the effective thermal resistance is measure to a control temperature, that may be a local temperature, an averaged temperature or a global maximum temperature:

$$\mathcal{F} = \mathrm{R}_{\mathrm{th,eff}} = \frac{T^\star - T_{\mathrm{ref}}}{\dot{Q}_{\mathrm{tot}}}; \quad \text{where} \quad \dot{Q}_{\mathrm{tot}} = \iiint_\Omega \dot{q}_{\mathrm{gen}} \,\mathrm{d}\Omega \quad \text{in 2D}. \tag{3.5}$$

Here, $T^\star$ represents the control temperature. The objective $\mathcal{F}$ depends – on the lowest level – only on the design variable field and the heat influx distribution into the domain Ω.

This definition for the objective is chosen exactly for that reason: it is dependent only on the design geometry and the thermal load distribution, which includes the boundary conditions. This means that any optimised design would be maximally effective under all total heat influx values or under all temperature constraints, so long as the heat influx is distributed the same over the domain and only scaled. The argument to support that the effective resistance is only dependent on this thermal load distribution and the design itself, is provided in appendix A.1.

Constraints definitions

For specific optimisation cases, additional constraints will apply. Such constraints might include inequality constraints $\mathcal{G}_i$, for which $\mathcal{G}_i \leqslant 0$:

1. A maximum fluid pressure drop constraint between all fluid outlets and inlets – (3.6a);
2. A maximum temperature constraint on the temperature-controlled section of the design domain – (3.6b);
3. A maximum average temperature constraint on the temperature-controlled section of the design domain or the domain boundary – (3.6c).

$$\mathcal{G}_1 = \frac{1}{|\Gamma^{\mathrm{in}}|} \iint_{\Gamma^{\mathrm{in}}} p \,\mathrm{d}\Gamma - \frac{1}{|\Gamma^{\mathrm{out}}|} \iint_{\Gamma^{\mathrm{out}}} p \,\mathrm{d}\Gamma - \Delta p_{\max}, \tag{3.6a}$$

$$\mathcal{G}_2 = \max\left(T\right) - T_{\max}, \tag{3.6b}$$

$$\mathcal{G}_3 = \frac{1}{|\Omega^{\mathrm{t.c.}}|} \iiint_{\Omega^{\mathrm{t.c.}}} T \,\mathrm{d}\Omega - T_{\mathrm{avg,max}} \quad \text{or} \quad \mathcal{G}_3 = \frac{1}{|\Gamma^{\mathrm{t.c.}}|} \iint_{\Gamma^{\mathrm{t.c.}}} T \,\mathrm{d}\Gamma - T_{\mathrm{avg,max}}. \tag{3.6c}$$

In these constraint functions, the maximum allowable pressure drop between the averaged in- and outlet boundaries is represented $\Delta p_{\max}$. The maximum allowable temperature is $T_{\max}$ and the maximum allowable domain- or boundary-averaged temperature is $T_{\mathrm{avg,max}}$. $\Omega^{\mathrm{t.c.}}$ and $\Gamma^{\mathrm{t.c.}}$ represent the temperature-controlled domains – $\mathbb{R}^2$ or $\mathbb{R}^3$ – and boundaries – $\mathbb{R}^1$ or $\mathbb{R}^2$ –, respectively.

Note that any selected objective or constraint function should be fully differentiable to enable optimisation. As such, the $\max$-function should be approximated using a p-norm approximation applied to, in this case, the temperature field T.

Optimiser overview and applied techniques

The process for a thermofluidic topology optimisation problem is outlined in figure 3.1.

Design iterations are performed through the gradient-based globally convergent method of moving asymptotes (GCMMA) as devised by Svanberg (2002). The gradients (sensitivities) of the objective and constraints are acquired through the adjoint state method. The construction of the specific interpolations used in the final model, will be covered in chapter 4.

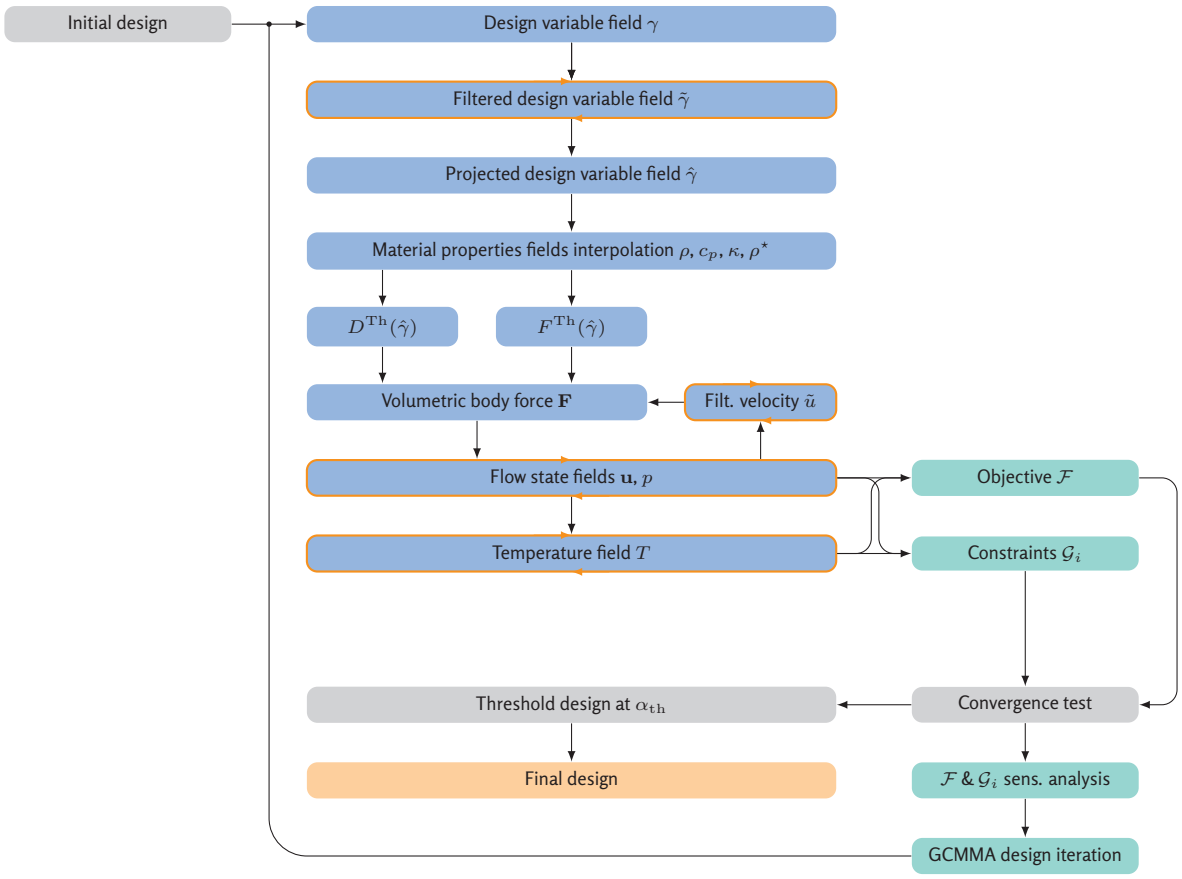

Figure 3.1: *Outline of the topology optimisation process for the single-design variable problem, using a Darcy-with-filtered-Forchheimer flow penalisation approach. Steps depicted within orange bounds are (internally) iterative processes.*

'Solid' material in the flow fields solver step will be modelled as a finite-impermeability medium. To this end, the momentum equations of the Navier-Stokes governing equations are modified to include carefully constructed penalisation terms for flow in 'solid' material elements in the volumetric body force term $\mathbf{F}$. The implementation of the three penalisation implementations outlined by Theulings, Noël, et al. (2025) will be considered and discussed in section 4.1. An additional implicit penalisation on gradient-parallel flow, abbreviated IGPP, will be introduced in section 4.2, in an effort to mitigate flow in solid material elements with the aim of improving the thermofluidic model accuracy in the optimisation loop.

The construction of a sufficiently accurate thermal and thermofluidic model for topology optimisation – especially considering that the problem involves a heat flux objective, rather than a more commonly used temperature-based objective – is key to obtaining meaningful optimisation results. The main identified issues that will be discussed and for which mitigation strategies will be proposed and implemented include the design for non-negligible advective heat transfer in solid-pseudodensity regions by the optimiser and the approach to deriving interpolation functions for the various relevant material properties and selection of penalisation parameters for accurate modelling.

3.2. Optimiser model construction

This section will outline the process employed to construct a more accurate (thermo)fluidic model for density-based topology optimisation.

The set-up of the optimisation model is undertaken in several steps, to ease the process of verification and validation of the model. This stepwise tuning and validation can be done without loss of maximum

attainable accuracy due to the sequential nature of the solver steps: see figure 3.1. The exception is the iterative solver process between $\mathbf{F}$ and flow fields p and $\mathbf{u}$. Consequently, this process will be regarded as a single step in the model set-up.

3.2.1. Approach

Firstly, various implementations and parameterisations of the of fluid permeation penalisation will be tested, as they are outlined in the paper by Theulings, Noël, et al. (2025). These tests will be performed on reference problems and equivalent discrete-density models will also be used to compare performance of the continuous-density descriptions to 'reality'. This work is done in section 4.1.

After, the gradient-parallel flow penalisation (IGPP) will be set up in section 4.2, though the concept will be introduced in section 3.3, where an interpolation function $\rho^\star(\hat{\gamma})$ is introduced for use in the fluid dynamics governing equations, such that it effectively reduces flow along projected-pseudodensity gradients.

To evaluate the constructed fluidic model, consisting of the fluid permeability penalisation and IGPP, any considered implementation will be tested in a two-dimensional 'island formation' simulation, in which a circular region of variable but uniform pseudodensity is placed in a uniform velocity field. In the simulation, the variable pseudodensity is decreased from 1 (fluid) to 0 (solid) gradually, to simulate a solid-material island forming within the flow field, so that the sensitivities of the constraints and objective may be analysed in the process of such an island forming. Reference simulations with a discrete solid-fluid interface will also be performed to compare results.

Then, the interpolation functions for κ, c_p and ρ will be constructed to build the eventual thermofluidic model, with a focus on accurate heat flux results through the solid-fluid interface. Evaluation of these functions will also be undertaken with various simulations, both single- and two-dimensional. This work will be outlined in section 4.3.

3.2.2. Performance indicators

Since model validation simulations for the optimiser are performed in an effort to quantify the performance of different implementations and parameterisations on model accuracy, several global performance indicators are used:

$$U_{\max} = \max_{\Omega} |\mathbf{u}| \qquad \text{Maximum absolute fluid velocity,} \tag{3.7a}$$

$$\Delta P = \frac{1}{|\Gamma^{\text{in}}|} \iint_{\Gamma^{\text{in}}} p \,\mathrm{d}\Gamma - \frac{1}{|\Gamma^{\text{out}}|} \iint_{\Gamma^{\text{out}}} p \,\mathrm{d}\Gamma \qquad \text{Pressure drop,} \tag{3.7b}$$

$$\dot{M} = \frac{1}{2} \oiint_{\Gamma} \rho^\star \frac{\mathbf{u} \cdot \nabla\hat{\gamma}}{|\mathbf{u} \cdot \nabla\hat{\gamma}|} \mathbf{u} \cdot \mathrm{d}\mathbf{\Gamma} \qquad \text{Leakage mass flux,} \tag{3.7c}$$

$$\partial_n U = \left| \nabla\mathbf{u} \cdot \frac{\nabla\hat{\gamma}}{|\nabla\hat{\gamma}|} \right|_{\hat{\gamma}=\alpha_{\text{th}} \in \Gamma} \qquad \text{Normal at-threshold abs. velocity gradient,} \tag{3.7d}$$

$$T_{\max} = \max_{\Omega} (T) \qquad \text{Maximum temperature,} \tag{3.7e}$$

$$\mathcal{Q} = -\iint_{\Gamma^{\text{in}}} \rho c_p T \mathbf{u} \cdot \mathrm{d}\mathbf{\Gamma} - \iint_{\Gamma^{\text{out}}} \rho c_p T \mathbf{u} \cdot \mathrm{d}\mathbf{\Gamma} \qquad \text{Net domain heat flux.} \tag{3.7f}$$

As the accuracy of the thermofluidic model is highly dependent on the accuracy with which the fluid velocity field can be simulated, the first three quantities listed are purely measures of fluid model accuracy. Depending on inlet conditions and the intended purpose of the reference problem in a fluid simulation either $U_{\max}$ or ΔP is used, though the 'normal-to-wall-flow' reference problem includes both. Meanwhile, quantity $\dot{M}$ is used in simulations to evaluate the efficacy of both selected fluid permeation penalisation schemes as well as IGPP implementations.

$\partial_n U$, in turn, is a measure of exclusively the fluid velocity field, but is only of importance due to its substantial effect on heat transfer across a solid-to-fluid interface, provided the leakage mass flux is sufficiently small.

Then, both $T_{\max}$ and $\mathcal{Q}$ are measures of the thermal and thermofluidic performance in a simulation and may be used as an optimisation constraint and optimisation objective interchangeably.

3.2.3. Model verification and validation reference simulations

Listed in the following sections are detailed explanations of the various simulations used to evaluate (thermo)fluidic model accuracy for the continuous-design variable problem that the optimiser will be solving for. Results will also include those found for an equivalent discrete-density design. Figure 3.2 further specifies the application of the various reference simulations during the construction of the density-based thermofluidic optimisation model.

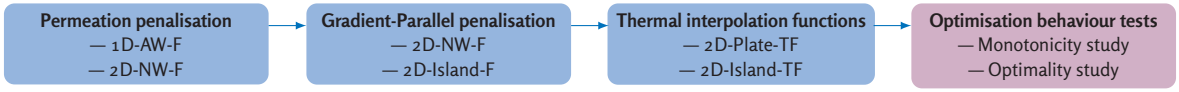

Figure 3.2: *Outline of the approach to thermofluidic model construction and verification, using Darcy penalisation, Implicit Gradient-Parallel Penalisation and minimal-complexity thermally-relevant material properties interpolations. Construction is followed by some simulations to test expected optimisation behaviour of the model.*

The dimensions of the 2D-NW-F, 2D-Island-F/TF and 2D-Plate-TF problem geometries and the approach to obtaining the discretised equivalent design are provided in section B.1.

Single-dimensional flow-along-wall fluidic simulation (1D-AW-F simulation)

This simulation is a reduced two-dimensional simulation, where the single spatial y-coordinate is taken to be normal to the fully-developed fluid velocity $u(y)$. This simulation is used to evaluate near-wall flow for gradient walls in comparison to the velocity profile of the equivalent discrete design.

The simulation can be performed for Darcy (D), Darcy-with-Forchheimer (DF) and Darcy-with-filtered-Forchheimer (DFF) penalisation schemes for in-material impermeability. The continuous designs simulated are step functions in the control pseudodensity γ and then filtered and projected using a Helmholtz PDE filter and hyperbolic tangent projection, respectively. The penalisation parameters are kept variable to enable the analysis of their effect on the along-wall flow profile. The IGPP implementation will be omitted when the fluid permeation penalisation is tested.

Results are generated using COMSOL's coefficient form PDE modules: one of which solves for the filtered pseudodensity field $\tilde{\gamma}$ with the Helmholtz PDE filter at 'radius' R_γ, one solves the single-dimensional simplified flow momentum conservation equation (equation (3.8)) for velocity profile $u(y)$ and – exclusively for the DFF approach – a final PDE module to solve for the filtered velocity field $\tilde{u}(y)$ using another Helmholtz filter with 'radius' R_u, applied to the absolute value of the velocity field. The equation that follows is a single-dimensional Navier-Stokes momentum equation, where the flow permeability penalisation is explicitly implemented in accordance with Theulings, Noël, et al. (2025).

$$\frac{\partial p}{\partial x} = \mu \frac{\partial^2 u}{\partial y^2} + f_1\left(\hat{\gamma}(y)\right) \frac{\mu}{h^2} u + f_2\left(\hat{\gamma}(y)\right) \frac{\rho}{h} U u \tag{3.8}$$

In equation (3.8), f_2 is 0 for a D-penalisation scheme, $U = |u(y)|$ for a DF-penalisation scheme and $U = \tilde{u}(y)$ for a DFF-fluid permeation penalisation scheme. For the DFF-approach $f_1(\hat{\gamma})$ and $f_2(\hat{\gamma})$ are identical, whereas f_2 is modified to double the exponents q and $\hat{q}$ for Darcy-with-Forchheimer, in accordance with Theulings, Noël, et al. (2025). For D-penalisation, f_1 is also scaled with the elemental Reynolds number $\mathrm{Re_h}$. Note that $\frac{\partial p}{\partial x}$ is assumed constant and resembles the pressure gradient in the flow direction. Finally h

resembles the characteristic mesh size, μ and ρ are material properties, where ρ is kept independent of y at ρ^{f} in this simulation.

The results generated with this reference problem can be found in section 4.1.1.

Two-dimensional normal-to-wall fluidic simulation (2D-NW-F simulation)
This simulation is two-dimensional in x- and y-spatial coordinates. An influx profile is specified on a boundary section, and a zero-control-pseudodensity wall is specified close behind this fluid inlet. For illustrative purposes, figure 3.3, depicts the problem geometry and illustrative fluid streamlines for both the (continuous-density) reference problem the discrete-density equivalent.

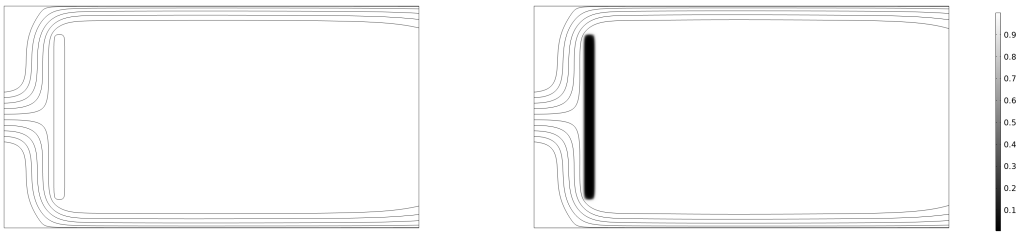

(a) *Discrete-density design simulation domain.* ***(b)*** *Continuous-density design simulation domain.*

Figure 3.3: *Domains for two-dimensional normal-to-wall fluidic simulations. Streamlines are added to indicate the inlet and outlet of the geometry.*

The aim of this simulation is to evaluate the through-wall flux, pressure drop and maximum absolute velocity for various in-material flow and IGPP penalisation schemes and parameterisations. Specifically, the simulation is performed to find the fluidic penalisation model implementation that will yield the most accurate velocity vector field and pressure field results compared to the discrete equivalent design, as well as to evaluate the efficacy of a gradient-parallel penalisation implementation. To this end, it uses COMSOL's laminar flow solver module and a coefficient form PDE solver module to find the filtered pseudodensity field and filtered velocity field. The in-material flow penalisation formulation is implemented as a volumetric body force term in the model, specified as $\mathbf{F}$ in governing equations (3.3).

The results generated with this reference problem can be found in section 4.1.2 and section 4.2.2.

Two-dimensional island fluidic simulation (2D-Island-F)
Also this simulation is two-dimensional in x- and y-spatial coordinates. A uniform inflow velocity is specified on the inlet edge and a 0-control-pseudodensity island is specified behind this fluid inlet. Figure 3.4, shows the problem geometry and illustrative fluid streamlines for both the (continuous-density) reference problem the discrete-density equivalent.

(a) *Discrete-density design simulation domain.* ***(b)*** *Continuous-density design simulation domain.*

Figure 3.4: *Domains for selected two-dimensional island problem fluidic simulations. Streamlines are added to indicate the inlet and outlet of the geometry.*

The aim of this simulation is to evaluate the through-wall flux, maximum absolute velocity and velocity gradient at-threshold for the three fluid permeation penalisation schemes, for the various parameterisations. The implementation of this penalisation as well as the solver is identical to the 2D-NW-F reference problem.

The results generated with this reference problem can be found in section 4.1.3 and section 4.2.2.

Two-dimensional flat plate flow thermofluidic simulation (2D-Plate-TF simulation)
In figure 3.5, the problem geometry and illustrative fluid streamlines are depicted. The problem consists of a simple geometry, where fluid flows over a flat surface. Heat is generated only in the solid elements using a RAMP interpolation function $\dot{q}_{\mathrm{gen}}$. The velocity field is fully developed at the inlet and the total mass flux through the domain equalised between the discrete and continuous problems. This is done explicitly, since an average-velocity-based inlet condition would act on a relatively longer domain boundary for the continuous problem.

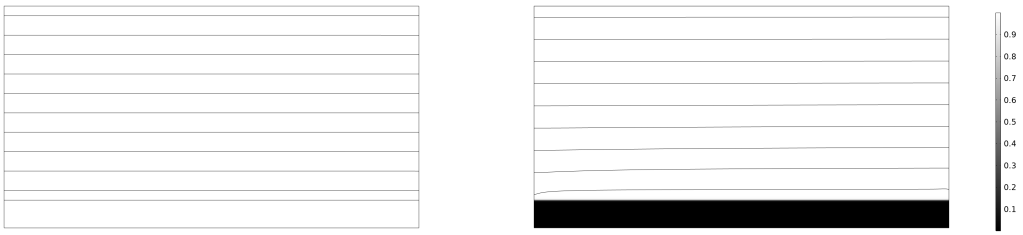

(a) *Discrete-density design simulation.* ***(b)*** *Continuous-density design simulation domain.*

Figure 3.5: *Domains for selected two-dimensional flat plate flow problem thermofluidic simulations.*

This simulation is intended to provide insight into the performance of various relevant material properties to the thermal problem, κ, c_p, ρ, $\dot{q}_{\mathrm{gen}}$, etc. Ultimately, the results will be used to decide on a well-performing optimisation model formulation. The global variables to the problem, such as pseudodensity projection function parameters and Helmholtz filtering radius, are carried over from the fluid reference simulations.

With the influx of heat into the domain constructed to be equal between the continuous- and discrete-density problems, the maximum temperature observed in the field becomes the best measure of the accuracy of any formulation or parameterisation of the continuous problem. This results from the maximum temperature being both a constraint that will be applied to the problem, as well as a measure of the domain's total thermal resistance under the same heat influx and distribution and fluid velocity field.

The results generated with this reference problem can be found in section 4.3.2.

Two-dimensional island thermofluidic simulation (2D-Island-TF simulation)
The geometry and definition of this problem are identical to those found in the two-dimensional fluidic island simulation (2D-Island-F). However, added to it, are the conduction-advection equation to solve the temperature field over the simulation domain and the necessary material properties interpolation functions, as well as a heat source interpolation. In figure 3.6, the problem geometry and illustrative fluid streamlines are presented, for both the discrete-density and a continuous-density construction.

(a) Discrete-density design simulation domain. *(b) Continuous-density design simulation domain.*

Figure 3.6: *Domains for selected two-dimensional island problem thermofluidic simulations.*

The objective of these simulations is to evaluate the accuracy of the continuous-pseudodensity formulation to their discrete-equivalent. This is done for various interpolations of material properties relevant to the thermal problem should aid in constructing a good optimisation model, as is the case with the 2D-Plate-TF reference problem.

The results generated with this reference problem can be found in section 4.3.2.

3.2.4. Model optimisation behaviour simulations

Beyond the described reference problems, several checks will need to be performed to evaluate the performance of the selected formulations and parameterisations of the model when it is applied in optimisation. Such checks include: a check for monotonicity of the objective function with the generation of a solid island in the domain for which desired behaviour is known and a check for optimality of the objective to occur with the narrowest achievable solid-fluid pseudodensity interface.

The first monotonicity check is necessary to ensure that the optimiser is able to generate any fully-solid regions into the domain. If the the behaviour of the problem is not monotonous, it may well either design intermediate-density regions or not design any features at all. It is currently unclear whether or not strict monotonicity can be achieved and if it is a necessary condition. The tests involve a check on the monotonicity of both the heat flux and the maximum temperature found in the simulation domain, as both could be used interchangeably as objective and constraint in optimisation. To further elucidate the results, the design's effective thermal resistance will be considered as a key performance metric for any continuous-density-based thermofluidic formulation.

Then, a reference problem is devised to verify that the crispest solid-fluid interfaces result in optimality of the objective with the constructed model. If so, it serves as evidence that the optimiser will tend to generate the narrowest intermediate-density regions possible in the projected pseudodensity field, resulting in the highest achievable accuracy in the model's results.

The results generated with this reference problem can be found in section 4.4.

3.3. Implicit Gradient-Parallel flow Penalisation – A conceptual exposition

Implicit gradient-parallel flow penalisation for fluidic density-based topology optimisation, or IGPP, refers to the use of the compressible (steady-state) Navier-Stokes equations, see equations (3.9), for penalisation of fluid flow into and out of regions and elements of pseudodensities that are associated with non-fluid material properties.

IGPP is 'implicit', as it does not add any terms to the governing equations or to the constraints or objective of the problem. Instead, the terms responsible for the realisation of effective penalisation are inherent to the existing (admittedly, now compressible) fluid dynamics governing equations. It is also 'gradient-parallel', as this technique only penalises fluid flow locally, where the gradient of the projected pseudodensity field and the velocity vector field share a non-zero component. In other words, it only acts where projection of the fluid

velocity vector field onto the density gradient field fields is non-zero and proportional to the local alignment of both fields.

As a starting point, consider the compressible Navier-Stokes equations for compressible fluids with uniform viscosities:

$$\boldsymbol{\mathcal{R}}^{\mathrm{m}} = \rho^{\star}\left(\mathbf{u}\cdot\nabla\right)\mathbf{u} + \nabla p - \mu\nabla^{2}\mathbf{u} - \left(\frac{1}{3}\mu + \zeta\right)\nabla\left(\nabla\cdot\mathbf{u}\right) - \mathbf{F} = \mathbf{0}, \tag{3.9a}$$

$$\mathcal{R}^{\mathrm{c}} = \nabla\cdot\left(\rho^{\star}\mathbf{u}\right) = 0. \tag{3.9b}$$

In this description, $\boldsymbol{\mathcal{R}}^{\mathrm{m}}$ denotes the residual vector field of the momentum equation and $\mathcal{R}^{\mathrm{c}}$ the residual scalar field of the continuity equation. In the real solution to the Navier-Stokes equations, all residuals evaluate to zero. In simulation, the equations are solved numerically to a specified tolerance.

Through the use of the compressible fluid governing equations, it becomes possible to leverage the continuity equation (3.9b) by specifying the material density in the fluid simulation to equate to the non-physical density field $\rho^{\star}(\hat{\gamma})$. The idea is to specify $\rho^{\star}$ to be exclusively dependent on the projected pseudodensity variable field $\hat{\gamma}$. In this way, the compressibility effects in the fluid are modelled to be purely design-dependent, instead of pressure-dependent.

Through the design variable-dependent description of this non-physical density, one may derive, from the continuity equation:

$$\mathcal{R}^{\mathrm{c}} = \nabla\cdot\left(\rho^{\star}\mathbf{u}\right) = 0 \tag{3.10}$$

$$\mathcal{R}^{\mathrm{c}} = \rho^{\star}\nabla\cdot\mathbf{u} + \nabla\rho^{\star}\cdot\mathbf{u} = 0 \tag{3.11}$$

$$\mathcal{R}^{\mathrm{c}} = \rho^{\star}\nabla\cdot\mathbf{u} + \rho^{\star\prime}\nabla\hat{\gamma}\cdot\mathbf{u} = 0, \tag{3.12}$$

$$\tag{3.13}$$

where $\rho^{\star\prime}$ denotes the first derivative of $\rho^{\star}$ with respect ot the projected pseudodensity variable. Note, from the last expression, that the residual is made up of two terms, of which the first is also found in the incompressible fluid formulation of the continuity equation, wherein the density – or $\rho^{\star}$, in this case – is fully invariant. The second term, however, includes $\nabla\hat{\gamma}\cdot\mathbf{u}$, a measure of fluid flow parallel to the design gradient, or, equivalently, normal to the projected pseudodensity field level contours. It, therefore, provides a means of control over fluid flow perpendicular to the 'walls' of a design. The specific interpolation of $\rho^{\star}$ with respect to the projected pseudodensity variable will be addressed in section 4.2.

3.4. Filtered pseudodensity fields and curvature-induced shrinkage

One complexity to the proper analysis of the various reference problems considered in this thesis, in particular the 2D-Island problem specification (2D-Island-F and 2D-Island-TF), concerns the geometry of the control pseudodensity field γ defined onto its simulation domain Ω. The issue is that the PDE filter does not only blur the interfaces, causing larger intermediate-density regions, but also causes the filtered and projected pseudodensity contours to shift in the domain. In essence, this means that the filtering step alters the generated design compared to the control pseudodensity variable γ.

This effect is, at least in part, outlined in the paper by Trillet, Duysinx, and Fernández (2021) and should be accounted for when setting up the discrete reference simulations: if the discrete simulation, that is not filtered, specifies a larger island diameter than is actually found in the filtered and projected field of the equivalent continuous-pseudodensity reference problem, this could significantly reduce the accuracy and the consequent usefulness of the reference simulations.

Now, its filtered pseudodensity field $\tilde{\gamma}$, will be solved for to illustrate the effect of curvature of the contours of $\gamma(\mathbf{x})$ on the position of the contours of $\tilde{\gamma}(\mathbf{x})$ and $\hat{\gamma}(\mathbf{x})$. Particular focus is placed on how it affects the contour $\hat{\gamma} = \alpha_{\mathrm{th}}$, where α_{th} is defined to be 0.5.

Since γ is circular in nature in the 2D-Island reference problems and may be centred on the origin, the control variable field γ can be defined:

$$\gamma(r) = H\left(r - \frac{D_\gamma}{2}\right). \tag{3.14}$$

with H the Heaviside step function. For the PDE filter:

$$-R_\gamma^2 \nabla^2 \tilde{\gamma} + \tilde{\gamma} = \gamma. \tag{3.15}$$

R_γ, the filtering radius that is also responsible for the minimum feature size control, is typically presented in the form $\frac{N_\gamma h}{2\sqrt{3}}$. In this expression, N_γ is some number of elements and h the characteristic mesh element length scale. This thesis makes use of $N_\gamma = 8$, unless differently specified.

The solution process to arrive at the closed expression for the filtered density field is covered in appendix A.2. However, the result is presented as equation (3.16):

$$\tilde{\gamma}(r) = 1 - \frac{1}{2\pi\, R_\gamma^2} \int_0^{\frac{D_\gamma}{2}} \int_0^{2\pi} K_0\left(\frac{\sqrt{r^2 + r'^2 - 2rr'\cos\theta'}}{R_\gamma}\right) r' \,\mathrm{d}\theta' \,\mathrm{d}r' \tag{3.16}$$

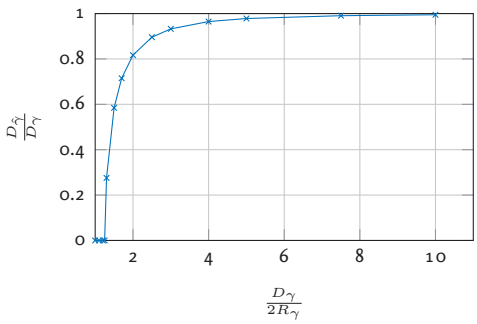

Figure 3.7: *Data on the effective diameter reduction in the projected pseudodensity field as a consequence of the curvature of the wall as designed in the control variable field relative to the pseudodensity filtering radius. Evaluated for $\hat{\gamma}$ at 0.5.*

All data presented throughout this thesis is corrected for this effect. Specifically, the geometries in the discrete-density simulations corresponding to the two-dimensional 'island' reference problems are explicitly corrected for the reduction in island diameter caused by pseudodensity filtering.

4

Optimisation model construction

4.1. Solid-region fluid permeation penalisation

In the search for an accurate continuous-density model for thermofluidic optimisation, an adequate fluid velocity penalisation scheme is necessary in order for solid-region flow to be kept to a minimum. In keeping flow minimal in those regions, the velocity field in the optimiser simulations will more closely resemble the equivalent discrete design that is obtained by thresholding the continuous design at $\hat{\gamma} = \alpha_{\mathrm{th}}$.

For the in-material penalisation schemes, the work of Theulings, Noël, et al. (2025) will be used, wherein three possible penalisation schemes are laid out for laminar flows with Re greater than 1. While it is important to mention that Theulings' work is principally applied to pure flow topology optimisation problems with pressure drop minimisation objectives, these schemes will be analysed and tuned for applicability to the extension into a thermofluidic topology optimisation model.

These three penalisation schemes include a Darcy penalisation scheme (D), a Darcy and Forchheimer penalisation scheme (DF) and a Darcy with filtered Forchheimer penalisation scheme (DFF):

$$\mathrm{D}: \quad \mathbf{F} = -\frac{10^{q-\hat{q}}(1-\hat{\gamma})}{10^{-\hat{q}}+\hat{\gamma}}\mathrm{Re}_h^{\mathrm{f}}\frac{\mu}{h^2}\mathbf{u} \tag{4.1a}$$

$$\mathrm{DF}: \quad \mathbf{F} = -\frac{10^{q-\hat{q}}(1-\hat{\gamma})}{10^{-\hat{q}}+\hat{\gamma}}\frac{\mu}{h^2}\mathbf{u} - \frac{10^{2q-2\hat{q}}(1-\hat{\gamma})}{10^{-2\hat{q}}+\hat{\gamma}}\frac{\rho}{h}|\mathbf{u}|\mathbf{u} \tag{4.1b}$$

$$\mathrm{DFF}: \quad \mathbf{F} = -\frac{10^{q-\hat{q}}(1-\hat{\gamma})}{10^{-\hat{q}}+\hat{\gamma}}\frac{\mu}{h^2}\mathbf{u} - \frac{10^{q-\hat{q}}(1-\hat{\gamma})}{10^{-\hat{q}}+\hat{\gamma}}\frac{\rho}{h}\tilde{U}\mathbf{u}, \quad \text{where} - R_u^2\nabla^2\tilde{U} + \tilde{U} = |\mathbf{u}| \tag{4.1c}$$

In the expressions for the D, DF and DFF fluid permeation penalisation schemes, interpolation is achieved via RAMP interpolation approach. Theulings, Noël, et al. (2025) achieved these general constructions, that provide a high degree of control over the fluid velocities found in solid regions. In short, their proposed implementations allow the use of a single parameter q to achieve a velocity reduction in solid elements of q orders-of-magnitude compared to the maximum velocities found in fluid regions. They also specify $\hat{q}$ as the interpolation parameter in their RAMP-functions.

From the advection-diffusion thermal governing equation (equation (3.4)), it is clear that any heat flux is the result of either diffusion or advection, where, generally, the effect of advection is large even for small fluid velocities. To the end of creating a sufficiently accurate continuous-pseudodensity model to use in optimisation, selecting an in-material flow penalisation that minimises fluid velocity in solid regions effectively

is therefore a necessity. Three tests will be performed to evaluate the penalisation schemes' relative performances based on this and other indicators. The results in this chapter are provided in plots akin to the one presented in figure 4.1, which is given as a clarification.

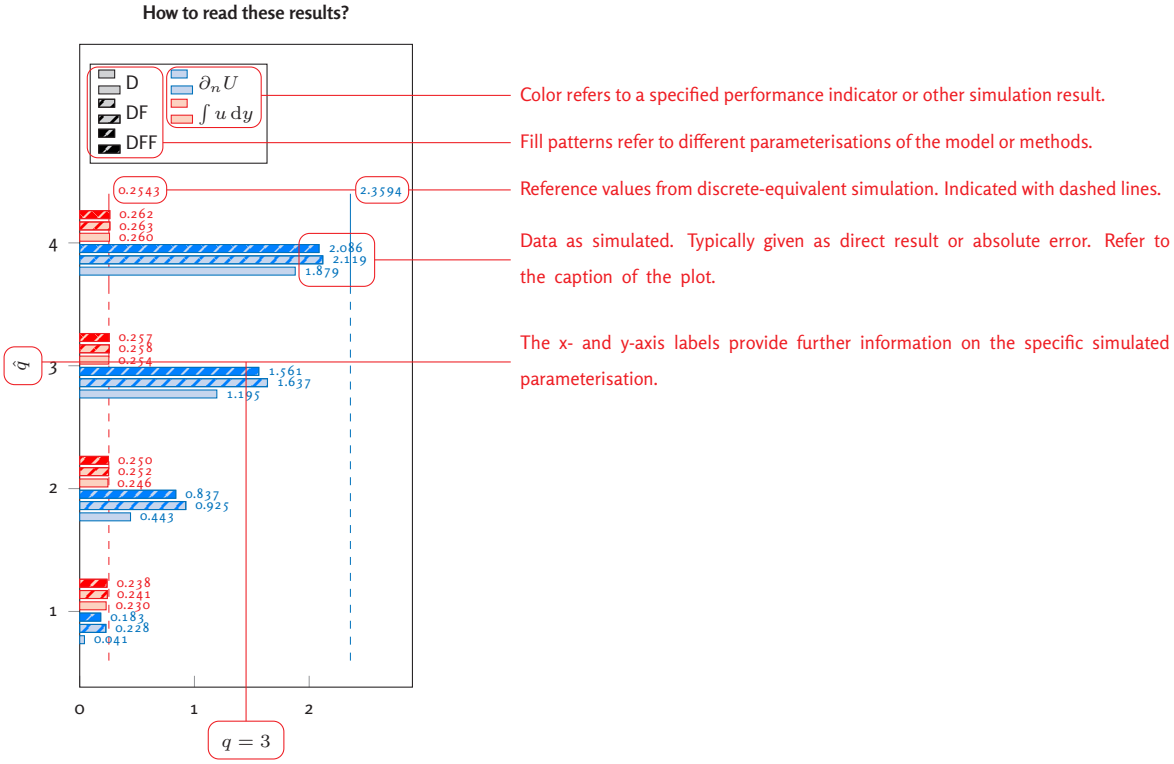

Figure 4.1: *Illustration of results plot as depicted for the various reference simulations in this chapter with comments on its lay-out as clarification.*

4.1.1. Along-wall in-material penalisation analysis

Figure 4.2 presents data on the single-dimensional along-wall fluidic simulation (1D-AW-F) performed using the D-, DF- and DFF-flow penalisation schemes as proposed by Theulings, Noël, et al. (2025) for continuous-design variable simulation. Eight permutations of values for $q \in \{2, 3\}$ and $\hat{q} \in \{1, 2, 3, 4\}$ have been simulated, with the aim of evaluating flow profile accuracy under variation of these penalisation parameters. Results of the discrete equivalent design are included.

The relevant performance indicators for this simulation are the velocity gradient in y-direction at the threshold projected pseudodensity ($\hat{\gamma} = \alpha_{\mathrm{th}}$) in s^{-1} and the total volume flux in $10^{-3}\,\mathrm{m}^2\,\mathrm{s}^{-1}$ under constant pressure gradient.

As a first observation, it is clear that the results appear to be relatively invariant with a parameter $r = \hat{q} - q$. This behaviour is expected, as the variation in the various simulation results is mostly caused by changes in fluid penalisation in pseudodensity gradient regions and only to a lesser degree in solid-designed regions. This, in turn, is subject to penalisation near the interface, $\hat{\gamma} \approx 0.5$, at which penalisation interpolation functions are of the order $10^{(q-\hat{q})}$. The only exception to this is the Forchheimer term in the DF-scheme, where this exponent is doubled.

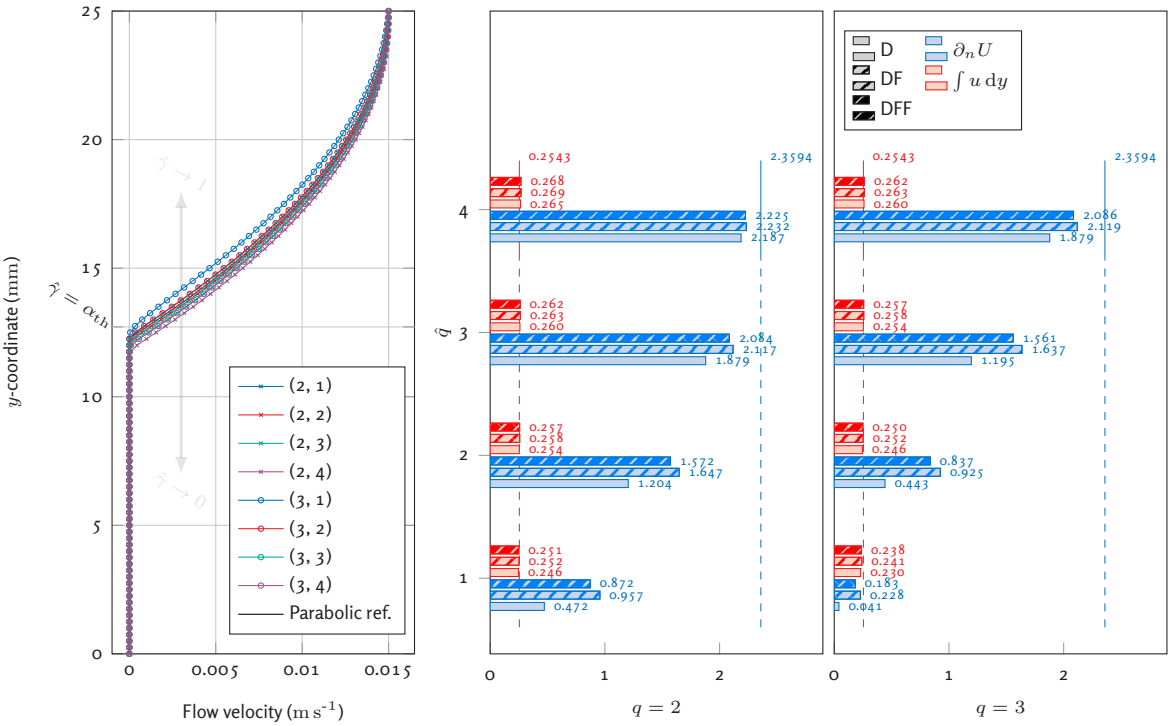

Figure 4.2: *Selected data from single-dimensional flow analyses. In the left-most plot, the velocity profile of the fluid is shown for DFF flow penalisation with $(q, \hat{q})$-parametisation. In the centre and right plots: relevant performance indicator figures $\partial_n U$ in s^{-1} and total volume flux in $10^{-3}\ \mathrm{m}^2\,\mathrm{s}^{-1}$ are shown for various values of $\hat{q}$, using the three seperate penalisation schemes, where $q = 2$. In the right-most plot, the same is shown, where, instead, $q = 3$. Used global parameters are N_γ 8, β 10 and N_u 10. Discrete simulation results are also depicted.*

Furthermore, the results show that increasing r leads to more accurate results on the at-threshold velocity gradient, but exacerbates the overestimation of total volume flux. The most accurate total flux results are found between $r = -1$ and $r = 0$ for all three penalisation schemes. Generally, the Darcy-only approach underperforms noticeably (on velocity gradient) in comparison to the DF- and DFF-approaches, though this relative performance deficit decreases for increased r. No stability or robustness issues are observed.

4.1.2. Through-wall in-material penalisation analysis

In figure 4.3 results of the two-dimensional fluidic simulation of a fluid inlet normal to a wall (2D-NW-F) are presented. Simulations have been performed for D-, DF- and DFF-penalisation schemes, each for the eight permutations of q and $\hat{q}$ that also are used in the single-dimensional along-wall flow simulation.

The first, left-most plot shows the horizontal velocity component in the simulation domain at a vertical cut plane, located at x 20 mm. It, thereby, intersects the wall at the center of its width. In other words, in the geometry from figure 3.3, the section is placed at the wall's vertical mid-plane.

This section plane should provide the most accurate reading of through-wall flow, by not intersecting the rounded-off corners of the continuous wall that result from the filtering and projection steps to the pseudo-density. Still, however, the 'curvature' of the wall at the endpoints will lead to the PDE filter slightly decreasing the projected width of the wall as a consequence of curvature-induced shrinkage. Wall thickness is 4 mm, its length is 60 mm at a 20 mm distance from the inlet, pseudodensity projection slope β is set at 10, and filtering radii to the pseudodensity, N_γ, and to the velocity field for DFF, N_u, are set at 8 and 10, respectively. Inlet velocity profile maximum is set to be 15 $\mathrm{mm\,s}^{-1}$.

The quad mesh size is set to be 0.5 mm ±3 % for all elements in the continuous simulation. All fields are simulated with linear elements, with the exception of the velocity field, that is simulated using Lagrangian

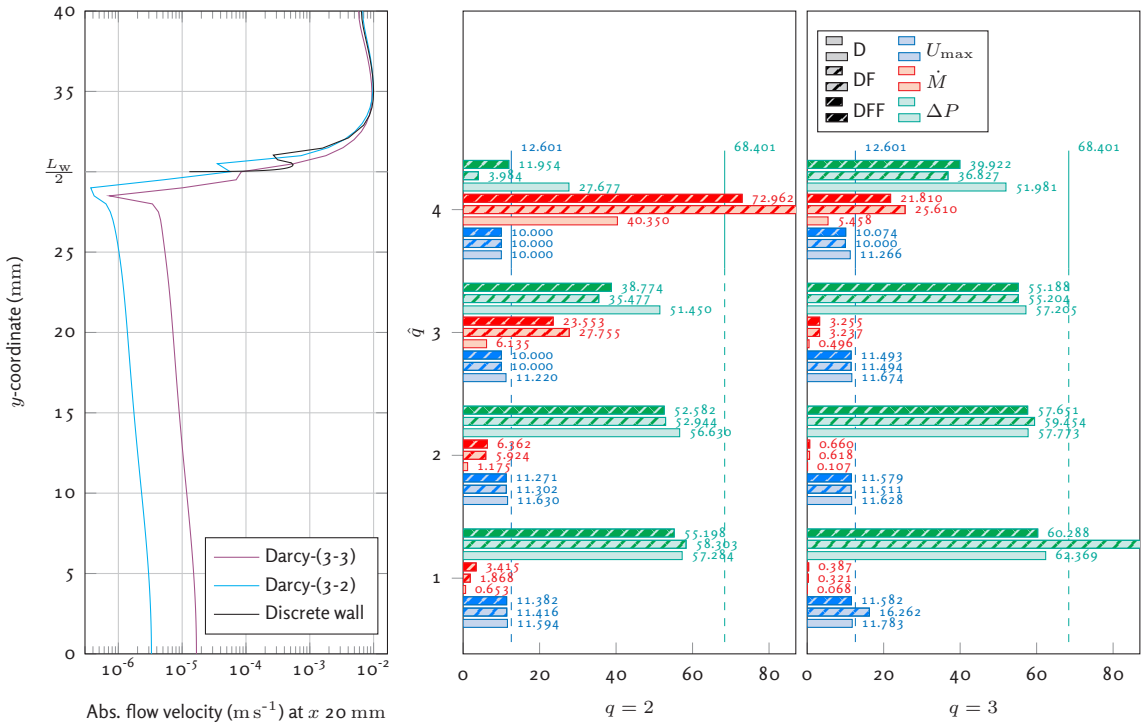

Figure 4.3: *Data from the 2D-NW-F simulation for the D-, DF- and DFF-penalisations of in-solid flow for eight $(q, \hat{q})$-permutations. The left plot shows flow profiles for $x = x_{\text{wall}}$ for illustration. The center and right plot show data on the relevant performance indicators for q 2 and 3, respectively. N_γ 8, β 10 and N_u 10. $U_{\max}$ is shown in $10^{-3}\,\mathrm{m\,s^{-1}}$, $\dot{M}$ in $10^{-3}\,\mathrm{m^2\,s^{-1}}$ and ΔP in* mPa. *Discrete simulation results are also depicted.*

quadratic elements.

The relevant performance indicators from the simulation data are maximum absolute velocity recorded in the simulation, presented in blue in figure 4.3, the through-wall volume flux (where $\hat{\gamma} \leqslant 0.5$), presented in red, and the average pressure drop over the entire simulation domain, presented in green.

From the results, it is clear to see that robustness issues with the DF-penalisation scheme occur for (3, 1)-parameterisation, as the maximum velocity as well as the pressure drop far exceed the values of even the discrete equivalent problem. Moreover, the DF-penalisation also underperforms dramatically in comparison to both D- and DFF-implementations for leakage flux for (2, 4)-parameterisation.

Surprisingly, the leakage flux of the Darcy-only penalisation scheme is consistently between 20 and 35% of both Forchheimer-modified approaches for all parameterisations except (2, 4). A data-driven understanding of the underlying cause of this observed behaviour is missing. Based on the simulation data, the selection of Darcy-(3, 1) for fluid permeation penalisation would be obvious as it yields the most accurate results on these three performance indicators at only a minimal increase in computational effort compared to $\mathbf{F} = \mathbf{0}$.

4.1.3. Island leakage simulation

Figure 4.4 shows the results of the two-dimensional island simulation (2D-Island-F) for the selected fluid permeation penalisation schemes and parameterisations of q and $\hat{q}$.

In this island simulation, parameters are selected in accordance with those from the 2D-NW-F reference simulation, with some exceptions. The control pseudodensity field is designed with a zero-density circular region with its center at x 60 mm and diameter D_γ 10 mm. Outside this region, γ is 1. The uniform normal inlet flow is defined to to have a magnitude of 5 $\mathrm{mm\,s^{-1}}$. The mesh is still a 0.5 mm ±3% mesh with square elements. Quadratic lagrangian elements are used for the velocity fields and linear elements for the pressure

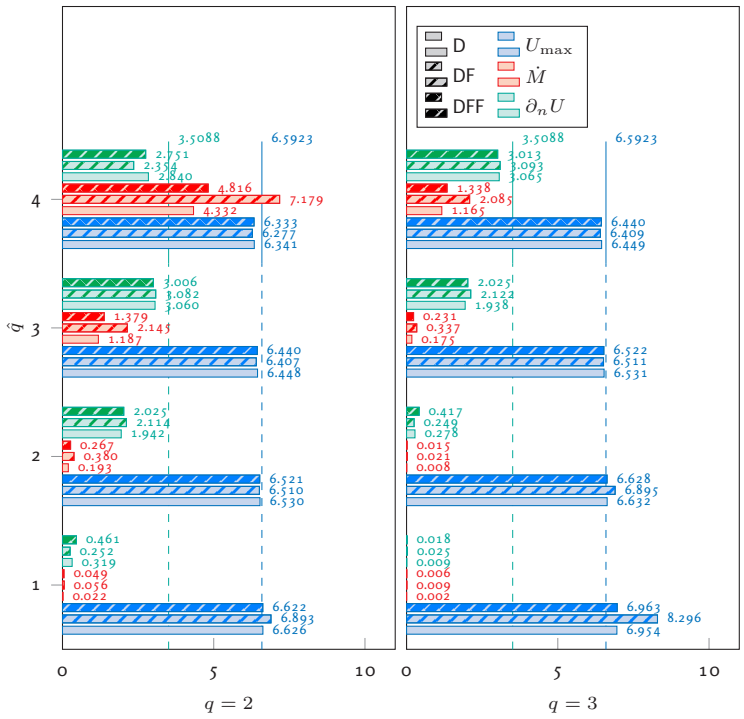

Figure 4.4: *Data from the 2D-Island-F simulation for the D-, DF- and DFF-penalisations of in-solid flow for eight $(q, \hat{q})$-permutations. Both plots show data on the relevant performance indicators for q 2 and 3, respectively. N_γ 8, β 10 and N_u 10. $U_{\max}$ is shown in $10^{-3}\,\mathrm{m\,s^{-1}}$, $\dot{M}$ in $10^{-3}\,\mathrm{m^2\,s^{-1}}$ and $\partial_n U$ in $\mathrm{s^{-1}}$. Discrete simulation results are also depicted.*

field and pseudodensity fields.

The relevant performance indicators found through this simulation data (displayed in figure 4.4) are maximum absolute velocity recorded in the simulation domain Ω, presented in blue, the through-wall volume flux, measured at the vertical measurement line $x = x_{\text{wall}}$ in red and the velocity gradient at the threshold projected pseudodensity on the measurement line in green.

The data shows problems with the Darcy-with-Forchheimer approach similar to those found in the normal-to-wall simulation: high penalisation with low interpolation coefficients, (3, 1), for instance, lead to stability problems, with the solver finding dramatically higher maximum velocities. Also, for the more relaxed penalisation parameterisations such as (2, 3), (2, 4) and (3, 4), leakage flux is significantly higher than equivalent Darcy-only and Darcy-with-filtered-Forchheimer implementations. In this simulation, too, the Darcy-only penalisation outperforms DFF in terms of leakage flux across all simulated parameter permutations, with it suffering from 40 to 85% of the leakage, comparatively, with this disparity increasing for higher penalisation and lower interpolation coefficients.

In terms of the velocity gradients found at vertical section line at $x = x_{\text{wall}}$, at the threshold pseudodensity, it is evident that higher interpolation parameters lead to better results up to a certain point: when increasing $\hat{q}$ causes the effective penalisation on leakage flux to deteriorate dramatically, such as observed between (2, 3) and (2, 4), especially the high-intermediate-density regions ($0.5 \leqslant \hat{\gamma} \lesssim 0.95$) suffer from that increase in fluid velocity. In turn, this causes the resulting velocity gradient at the threshold pseudodensity to reduce. Following this data, selection of $r = -1$ yields the best results on modelled velocity gradient accuracy, with all three penalisation schemes suffering from similar inaccuracy, ranging from -14.5% to -12% to the equivalent discrete simulation.

4.1.4. Scheme and parameter selection

From the data presented in figures 4.2, 4.3 and 4.4, it is clear that the results, in terms of the selected performance indicators are fairly invariant with $q - \hat{q} = r$. Some robustness problems are observed with the Darcy-with-Forchheimer approach that are of the same nature as those described by Theulings, Noël, et al. (2025). While it is much less computationally complex than its filtered equivalent, DFF and, thus, takes substantially less time, robustness problems appear for integer r, if $|r| > 1$, limiting its realistically achievable accuracy across all performed simulations.

The results also show that any parameterisation change, particularly to r, to improve velocity gradient accuracy comes at the cost of increased leakage mass flux. Given the relatively unknown efficacy the proposed IGPP penalisation, selecting two formulations with meaningfully different parameter r and comparing those, will provide insight into the most effective parameterisation of the fluid permeation penalisation in conjunction with their, then individually simulated, IGPP-parameterisations. The better the performance of IGPP is, the more may be compromised on leakage mass flux performance figures in the fluid permeation penalisation to improve velocity gradient figures.

Finally, both the Darcy-only and Darcy-with-filtered-Forchheimer schemes are stable within the simulated ranges and problems. Considering the fairly minimal observed performance difference between the two schemes – Darcy-only even outperforming DFF in some simulations – both will be considered in IGPP performance evaluation and tuning. Moreover, the D-scheme is particularly interesting due to the substantial reduction in computational effort it provides in comparison to DFF.

Four fluid permeation penalisations are selected for further testing in the IGPP reference simulations: r is selected -2 and 1 and applied to, both, D- and DFF- approaches. Simulations will be performed with $q = 3$, so $\hat{q}$ will then be 5 and 2, respectively. This leads to the following selection of implementations and parameterisations: D-(3,2), D-(3,5), DFF-(3,2) and DFF-(3,5). Values for r are selected to be very differing between -2 and 1, since the efficacy of further penalisation is unknown: while the additional implicit penalisation scheme may benefit from permeation penalisation factors as will be argued in section 4.2, it is also possible that IGPP will prove to be significantly less effective at already-high penalisation on the Darcy term.

Type	q	$\hat{q}$	r	Name
D	3	2	1	D-(3,2)
D	3	5	-2	D-(3,5)
DFF	3	2	1	DFF-(3,2)
DFF	3	5	-2	DFF-(3,5)

Table 4.1: *Overview of selected permeation penalty schemes and parameterisations for further investigation.*

For simulations with $\hat{q} = 5$, the efficacy of IGPP in limiting leakage mass flux will need to be very good and should then automatically lead to accurate velocity gradients at the threshold pseudodensity as a result of the favourable parameterisation of the fluid permeation penalisation term(s). Conversely, for $\hat{q} = 2$, the efficacy of IGPP in mitigating the leakage flux need not be as good to achieve desired levels, but leaves the velocity gradient accuracy compromised and in need of a precisely tuned interpolation of 'thermofluidic' material properties to reach the desired accuracy in simulated heat flux, compared to the discrete-equivalent problems.

4.2. Implicit Gradient-Parallel Penalisation (IGPP)

In pursuit of more accurate fluid field solver models for density-based topology optimisation problems, considering specifically thermofluidic modelling accuracy, the results found in section 4.1 are not satisfactory. The substantial remaining inaccuracies inevitably lead to lesser-optimised topologies, as the optimisation

model is of too limited accuracy to reality, creating designs that will almost always underperform when physically realised and implemented.

To improve the continuous-pseudodensity fluidic model, this thesis introduces a new, additional penalisation scheme, dubbed IGPP. Introduced in section 3.3, Implicit Gradient-Parallel Penalisation of fluid flow will now be further elaborated upon and evaluated for performance in this section.

How IGPP should work: a quasi-quantitative estimation of performance

To capture the essence of the manner in which implicit gradient-parallel penalisation on fluid flow achieves in-solid fluid flow reductions is fairly easily explained. This technique – at least in the implementation as it is shown throughout this thesis – relies heavily on effective in-solid, or fluid permeation, penalisation. With the introduction of the reduced-tuning formulations by Theulings, Noël, et al. (2025), adequate control over this penalisation is possible, meaning some very rough estimates can be made in an effort to quantify the effect of IGPP. Using the compressible Navier-Stokes equations and a (non-physical) density interpolation function $\rho^\star(\hat{\gamma})$, the following line of reasoning may be reached.

Consider fluid permeation penalisation in accordance with Theulings, Noël, et al. (2025). In their paper, they find that fluid velocity in solid elements can reasonably be controlled via

$$|u^{\mathrm{s}}| \leqslant 10^{-q}\, u_\infty. \tag{4.2}$$

simultaneously, through integration with respect to spatial coordinate x, we find for the continuity equation, if $u = u_\infty$ and $\rho^\star = \rho^{\mathrm{f}}$ as $x \to -\infty$:

$$\rho^\star u = \rho^{\mathrm{f}} u_\infty - \int_{-\infty}^{x} \rho^\star(\hat{\gamma}(\xi)) \frac{\partial v}{\partial y}\, \mathrm{d}\xi\,. \tag{4.3}$$

Now assume that along x, pseudodensity starts off at 1 (fluid) for sufficiently negative x, 0 for a solid region near $x = 0$ and returning to 1 for large x. As $x \to \infty$, $\rho^\star$ and u return to ρ^{f} and u_∞, respectively. Finally, $\rho^\star$ is constructed such that $\rho^{\star\mathrm{s}} = 10^{-a}\,\rho^{\mathrm{f}}$ for some real, positive a. In this case, the continuity expression can be observed at four values of x: x_1 well 'before' the solid region, x_2 within the solid region, possibly close to its trailing 'edge' (near to the second intermediate-density region), at x_3 just after the second intermediate-density region and, finally, x_4 for which $x \to \infty$. A two-dimensional schematic depiction of a problem of this form is given in figure 4.5. Consider, from this particular schematic depiction, $y = 0$, for which $\frac{\partial \rho^\star}{\partial y} = 0$.

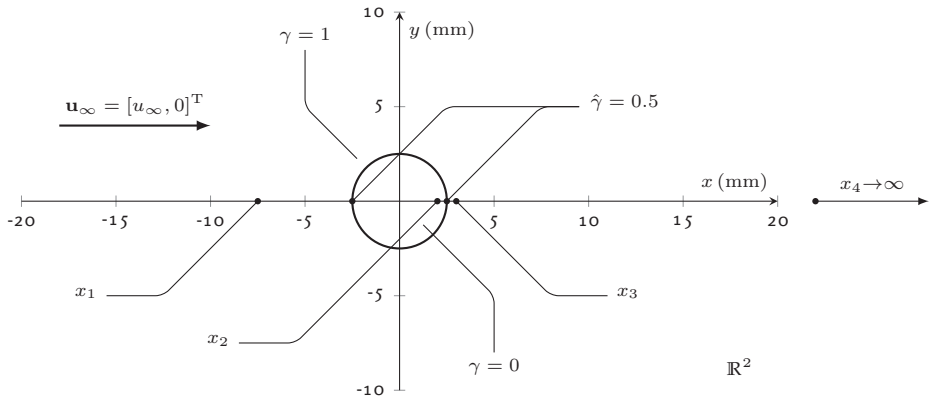

Figure 4.5: *A possible representation of the problem for illustrative purposes. This particular design, characterised by control variable field γ, is axisymmetric and makes a discrete step, to result in the sharpest possible interface. Only symmetry plane $y = 0$ is to be considered.*

First, observe for x_1 the integrated continuity equation. From $\hat{\gamma} = 1$ for fully-fluid elements, we find that

$$\rho^{\mathrm{f}} u_\infty = \rho^{\mathrm{f}} u_\infty - \int_{-\infty}^{x_1} \rho^\star(\hat{\gamma}(\xi)) \frac{\partial v}{\partial y} \,\mathrm{d}\xi \implies \int_{-\infty}^{x_1} \rho^\star(\hat{\gamma}(\xi)) \frac{\partial v}{\partial y} \,\mathrm{d}\xi = 0.$$

This is not particularly surprising: far-field fluid flow is uniform, so negligible 'divergence' of the fluid occurs until quite a short distance before the leading edge of the solid feature.

Then, consider the case for x_2, where fluid is (somewhere) inside the solid region. We assume that the fluid permeation penalisation is adequately controllable following the relation in equation (4.2). Also assuming, for now[1], that $\rho^\star$ may be constructed to yield $\rho^{\star\mathrm{s}} = 10^{-a}\,\rho^{\mathrm{f}}$, one deduces that, in the theoretical worst case:

$$\rho^\star u|_{x_2} = 10^{-a-q}\rho^{\mathrm{f}} u_\infty = \rho^{\mathrm{f}} u_\infty - \int_{-\infty}^{x_2} \rho^\star(\hat{\gamma}(\xi)) \frac{\partial v}{\partial y} \,\mathrm{d}\xi \implies \int_{-\infty}^{x_2} \rho^\star(\hat{\gamma}(\xi)) \frac{\partial v}{\partial y} \,\mathrm{d}\xi = \rho^{\mathrm{f}} u_\infty \left(1 - 10^{-q-a}\right).$$

Under the working assumption that $x_3 - x_2$ is small, meaning that the intermediate-density region at the trailing edge of the solid feature in the simulation domain is 'sufficiently narrow', the assumption

$$\int_{x_2}^{x_3} \rho^\star(\hat{\gamma}(\xi)) \frac{\partial v}{\partial y} \,\mathrm{d}\xi = 0$$

is reasonable, since $\frac{\partial v}{\partial y}$ will be negligibly small sufficiently near to the trailing edge *when flow is not both (nearly) inviscid and globally irrotational*. A case where both conditions are met, would result in potential flow, where $\frac{\partial v}{\partial y}$ would be very significant in size near the trailing edge of the solid feature. However, neither condition is met for the cases considered in this thesis, certainly not both simultaneously. Consequently, one may assume with reasonable accuracy:

$$\rho^\star u|_{x_2} = \rho^\star u|_{x_3} \implies u|_{x_3} = 10^{-a-q}\, u_\infty.$$

Flow order reduction at x_3 is even greater than in the solid domain (x_2) as a consequence of the compression due to the severe compressibility effects introduced with $\rho^\star$. Since fluid velocity along the gradient of the trailing edge of solid regions is responsible for the largest contribution to non-physical heat transfer in thermofluidic modelling with a density-based impermeability model, this additional order reduction is expected to drastically improve thermofluidic modelling accuracy.

Finally, for a sufficiently large x_4, one may assume 'convergence' to have taken place to completion in the wake of the solid region, resulting in $u|_{x_4} = u_i nfty$ to within reasonable accuracy. This yields, by extension, that

$$\int_{x_3}^{x_4} \rho^\star(\hat{\gamma}(\xi)) \frac{\partial v}{\partial y} \,\mathrm{d}\xi = -\left(1 - 10^{-q-a}\right) \rho^{\mathrm{f}} u_\infty.$$

This consequently, as expected via far-field flow conditions, also implies:

$$\int_{-\infty}^{\infty} \rho^\star(\hat{\gamma}(\xi)) \frac{\partial v}{\partial y} \,\mathrm{d}\xi = 0.$$

In conclusion, there are some observations to be made and a warning to be issued about this quasi-quantitative approximation to IGPP's performance in conjunction with the fluid permeation penalisation formulations based on the work of Theulings, Noël, et al. (2025).

For one, *within* the solid region, mass flux is reduced by IGPP if and only if the fluid permeation penalisation is able to achieve the desired order reduction as given in equation (4.2), but the fluid velocity is, theoretically, reduced with order q, not $q + a$. Only after 'compression', due to the increase in $\rho^\star$ through the intermediate density region of the trailing edge (from x_2 towards x_3), is the fluid velocity further reduced to be $a + q$ orders of magnitude smaller than u_∞. The increased reduction in mass flux in the solid region due

[1] The relation $\rho^{\star\mathrm{s}} = 10^{-a}\,\rho^{\mathrm{f}}$ is permissible under the necessary conditions placed on $\rho^\star(\hat{\gamma})$. See section 4.2.1.

to the introduction of IGPP also directly implies that the mass flux inside the fluid region (around the solid region) is closer to that found for a discrete-walled design and, theoretically, a more accurate solution to the flow fields.

As for the warning, it is important to mention that the efficacy of IGPP is directly dependent on the efficacy of the fluid permeation penalisation scheme for any given problem: poorer penalisation of fluid velocity leads to a proportionally poorer penalisation of leakage mass flux and reduced effects of trailing edge velocity reduction as a direct consequence. A specific issue that may come up is that the drastic reduction in the non-physical density inside solid elements could cause the fluid permeation penalisation parametrisation to be unable to achieve the expected order-of-magnitude velocity reduction as it was not derived for such excessively compressible (i.e. variable-density) flow conditions.

4.2.1. The compressible fluid continuity equation in numerical simulations

While the true residual $\mathcal{R}^{\mathrm{c}}$ of the continuity equation might be defined as defined above, a discretised equivalent is utilised in finite element or volume methods such as those employed by the numerical models that are used throughout this thesis. The COMSOL documentation provides the techniques that are available as termination criteria in the iterative (non-linear) solver. The COMSOL documentation (see COMSOL AB (2025)) provides the following non-linear solver termination criteria for stationary problems:

$$\sqrt{\frac{1}{M}}\sqrt{\sum_{j=1}^{M}\frac{1}{N_j}\sum_{i=1}^{N_j}\left(\frac{|E_{i,j}|}{W_{i,j}}\right)^2} \leqslant \varepsilon_{\mathrm{tol,sol}} \tag{4.4}$$

$$\sqrt{\frac{1}{M}}\sqrt{\sum_{j=1}^{M}\frac{1}{N_j\bar{W}_j^2}\sum_{i=1}^{N_j}|F_{i,j}|^2} \leqslant \varepsilon_{\mathrm{tol,res}} \tag{4.5}$$

where $E_{i,j}$ is the error estimate of element i after a solver iteration corresponding to the numerical solution of field(s) j, U_j. The error estimate follows from the (damped) Newton method in the non-linear solver. $W_{i,j}$ is a scaling factor, which, by default, divides the absolute error by the solution value to result in the elementwise relative error of i for field j. By default, $\bar{W}_j$ is the average of all DoF's (all i) residuals for the field j. N_j resembles the number of degrees of freedom in field j and M the number of total fields. $F_{i,j}$ may be found through the discrete-equivalent evaluation of scalar field $\mathcal{R}^{\mathrm{c}}$ for individual elements i in field j.

While for the residual termination criterion of equation (4.5) it is more straight-forward to find an appropriately-sized tolerance value, simulations have shown earlier on, that this termination criterion results to much poorer convergence behaviour. Consequently, the solution-based termination expression in equation (4.4) will be utilised. Also note, that in both expressions, the left-hand side 'error' terms of the inequalities include only the square of the absolute value of the solution error estimate or the residual value.

From hereon, the working out turns problematic. Since the desired outcome of IGPP is for the numerically solved fluid fields to still exhibit incompressible behaviour *and* mitigate fluid flow along projected pseudo-density field gradients simultaneously, it is necessary to devise $\rho^\star$ in a manner such that the terms do not 'trivially' cancel. That is

$$\rho^\star\nabla\cdot\mathbf{u} = -\rho^{\star\prime}\nabla\hat{\gamma}\cdot\mathbf{u}$$

must never hold for $\mathbf{u}\neq\mathbf{0}$. It is possible for this to occur, because the solver, logically, for compressible fluids, should allow these results and only limits the absolute square of the residual, $|\mathcal{R}^{\mathrm{c}}|^2$. Besides this issue, defining a solution tolerance $\varepsilon_{\mathrm{tol,sol}}$ that corresponds to the desired tolerance on the continuity equation residual will require some attention.

Conditions for $\rho^\star$

In total, two conditions exist that the non-physical density function must adhere to:

1. $\rho^\star(1) = \rho^\mathrm{f}$;
2. $\rho^\star \nabla \cdot \mathbf{u} \neq -\rho^{\star\prime} \nabla\hat{\gamma} \cdot \mathbf{u}$.

Considering the second condition, this expression is realised by the stricter-than-necessary condition that

$$(\rho^\star \nabla \cdot \mathbf{u}) \left(\rho^{\star\prime} \nabla\hat{\gamma} \cdot \mathbf{u}\right) > 0.$$

Under the assumption that at the gradient in the projected pseudodensity, an adequate, dominant fluid permeation penalisation gradient is found, it is reasonable to assume that compressive effects (with fluid velocity decreasing) will occur upon fluid entering into solid regions (decreasing $\hat{\gamma}$) and expansive effects upon increasing $\hat{\gamma}$ into fluid regions. Following this reasoning we find that this implies:

$$(\nabla \cdot \mathbf{u}) (\nabla\hat{\gamma} \cdot \mathbf{u}) \geqslant 0.$$

Then, finally, this implies

$$\rho^\star \rho^{\star\prime} > 0 \implies \rho^{\star\prime} > 0, \quad \text{if } \rho^\star > 0$$

as the second necessary condition in the construction of $\rho^\star$ for effective gradient-parallel flow penalisation. With this description, flow velocities will generally be increased in the solid regions by the reduction in $\rho^\star$. However, since the fluid permeation penalisation is known to be effective in fluid velocity order-of-magnitude reduction for solid regions, certainly in fully-solid regions, these effects should combine to further reduce fluid mass flux into solid regions. Note, that the interpolation used for $\rho^\star$ throughout this thesis will be linear in $\hat{\gamma}$.

A brief note on the additional term in the compressible fluid momentum equation
In the case of successful penalisation of along-gradient fluid flow with IGPP, as the result of effective construction of non-physical density interpolation function $\rho^\star(\hat{\gamma})$, results should – up to the allowable tolerances – satisfy the true residual expression, leading to

$$\nabla\hat{\gamma} \cdot \mathbf{u} = 0 \implies \nabla \cdot (\rho^\star \mathbf{u}) = \rho^\star \nabla \cdot \mathbf{u} = 0.$$

It naturally follows, of course assuming $\rho^\star \neq 0$ in the entirety of Ω – as can be ensured through proper construction –, that $\nabla \cdot \mathbf{u}$ turns zero in the entire domain. The necessary condition is for trivial cancellation not to occur within Ω.

With all conditions met, $\nabla \cdot \mathbf{u}$ being equal to zero in all of the simulation domain, will directly imply that the gradient of the velocity field divergence, $\nabla (\nabla \cdot \mathbf{u})$, should also evaluate to the zero vector in all of Ω. Consequently, the momentum equation will reduce to its incompressible form.

4.2.2. Reference simulation results for IGPP

Figures 4.6 and 4.7 show the results generated by the 2D-NW-F and 2D-Island-F reference simulations, respectively, this time including IGPP. The selected fluid permeation penalisation formulations from section 4.1.4 are utilised and subjected to a range of $\Delta\rho^\star$ parameter values, including 0, to demonstrate the effect of IGPP against the case without it.

From the displayed data of the 2D-NW-F simulations in figure 4.6 it is evident that the introduction of IGPP has a significant relative effect on the performance indicators through its effect on the fluid velocity fields for all four simulated fluid permeation penalisation formulations, leakage flux decreases substantially under increasing $\Delta\rho^\star$, as it increases the 'strength' of the gradient-parallel penalisation effect. In the cases of ineffective penalisation near the threshold value $\hat{\gamma} = \alpha_\mathrm{th}$, D-(3,5) and DFF-(3,5), the absolute effect of IGPP on the indicators is greatest: in reducing leakage mass flux dramatically, the fluid velocity field becomes so much more accurate that it also brings the pressure drop measure within less than 10% error compared to

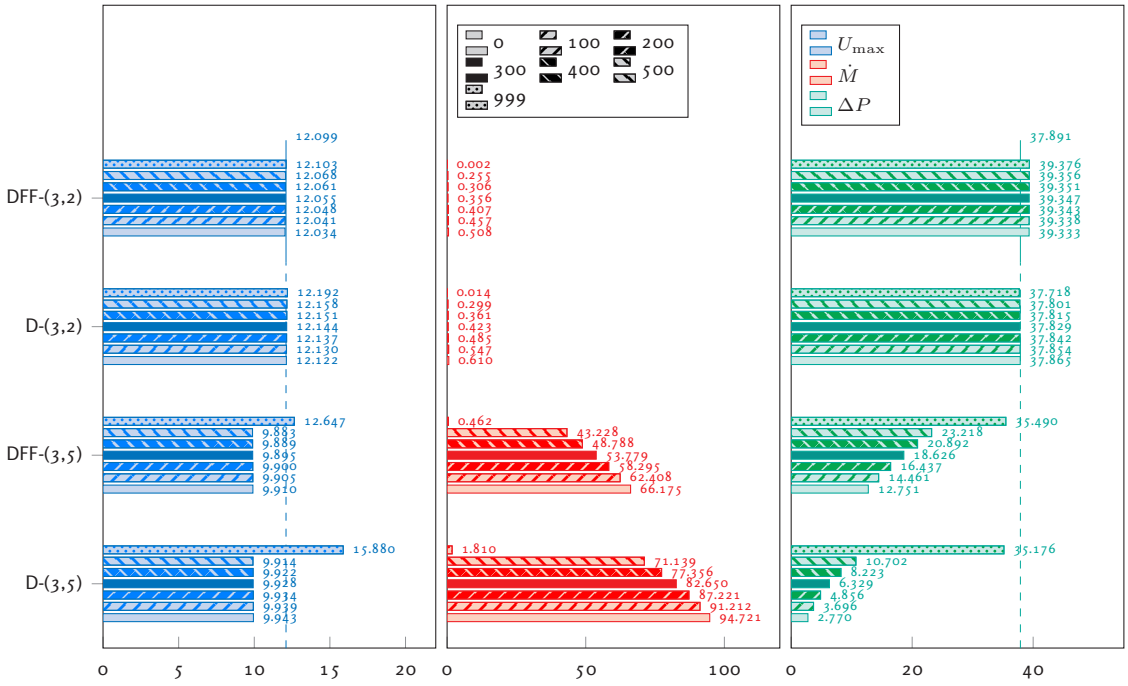

Figure 4.6: *2D-NW-F reference simulation results for selected performance indicators. Data is generated for four impermeability penalisation schemes and with for each seven parameterisations of $\rho^\star(\hat{\gamma})$, with $\Delta\rho^\star$ at 0, 100, 200, 300, 400, 500, 999 $\mathrm{kg\,m^{-3}}$. In the left plot, the maximum velocity in the simulation domain is shown, in the center plot the leakage mass flux and in the right plot the average pressure drop over the domain. Discrete simulation results are also depicted.*

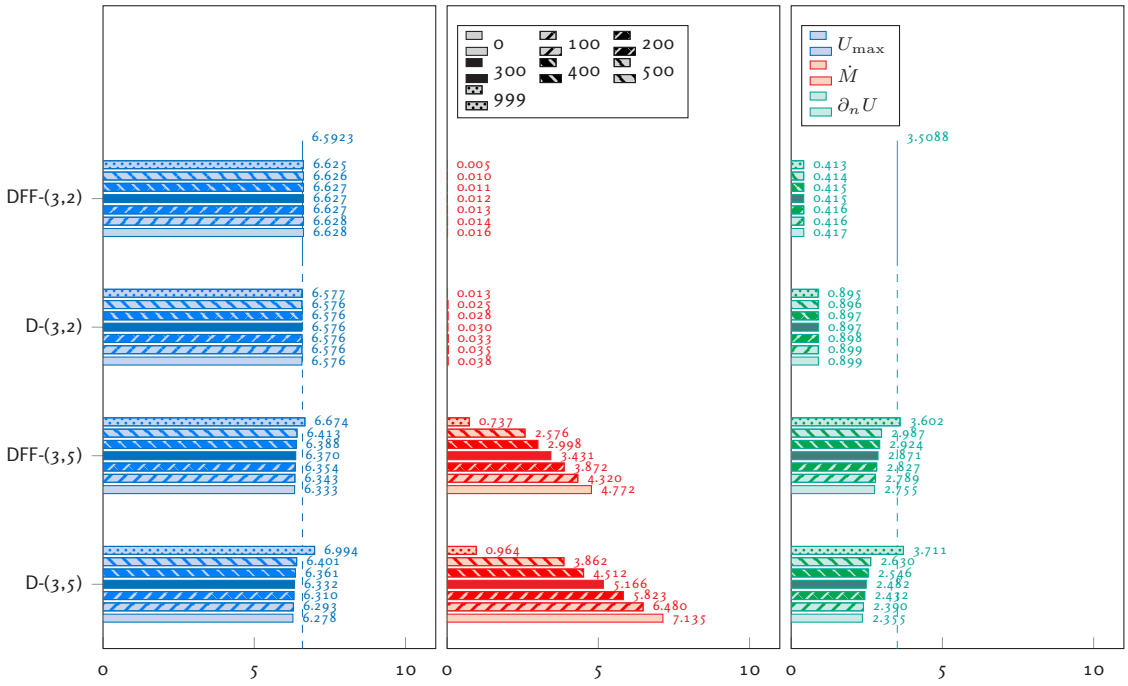

Figure 4.7: *2D-Island-F reference simulation results for selected performance indicators. Data is generated for four impermeability penalisation schemes with for each seven parameterisations of $\rho^\star(\hat{\gamma})$, with $\Delta\rho^\star$ at 0, 100, 200, 300, 400, 500, 999 $\mathrm{kg\,m^{-3}}$. In the left plot, the maximum velocity in the simulation domain is shown, in the center plot the leakage mass flux and in the right plot the at-threshold velocity gradient at the vertical mass flux measurement plane through the center of the island. Discrete simulation results are also depicted.*

the discrete-density reference simulation for $\Delta\rho^\star$ equal to 999, despite the poor fluid permeation penalisation in the intermediate-density region of the problem. While caution is necessary, these results also seem to support the hypothesis that achieving orders-of-magnitude reduction in leakage mass flux through effective reduction of $\rho^\mathrm{f} - \Delta\rho^\star$ is, in fact, possible for at least some circumstances.

Especially as $\Delta\rho^\star$ approaches ρ^f, penalisation is very effective, with 999 reducing leakage flux to only 0.4% - 2.5% compared to the case without IGPP.

Meanwhile, figure 4.7 shows that the effect of IGPP on the island simulation (2D-Island-F) is not as pronounced. The reason for this is not entirely apparent. In spite of this, IGPP, with $\Delta\rho^\star$ at 999, is still able to realise a factor 6 to 7 reduction in $\dot{M}$ for cases with $\hat{q}$ at 5 and a factor of approximately 3 for cases where $\hat{q}$ is set at 3. Also here, it is apparent that IGPP is working – albeit to a seemingly lesser degree than in the 2D-NW-F simulation – to increase the velocity gradient indicator $\partial_n U$ to more realistic levels and reduce leakage simultaneously. The effect of improving the at-threshold velocity gradient is only observed for $\hat{q} = 5$-simulations.

In both of the performed reference simulations, there are indications of possible stability and / or robustness problems occurring at very high $\Delta\rho^\star$: the maximum measured velocity in the domain as well as the velocity gradient actually exceeding the discrete simulation results is indicative of such problems, similar to how these effects are observed for the Darcy-with-Forchheimer fluid permeation penalisation scheme when $q - \hat{q}$ takes on large values. In and effort to construct an effective, yet robust penalisation scheme for the fluid solver, some additional simulations are performed to find a more favourable parameterisation of IGPP.
Table 4.2 includes additional 2D-Island-F simulation results for selected parameterisations of the fluid permeation penalisation, as well as $\Delta\rho^\star$. On the data presented, several interesting observations can be made. Firstly, under constant $q - \hat{q}$, the results are nearly completely invariant to change of the fluid permeation penalisation parametrisation. Then, the effect of increasing $\Delta\rho^\star$ from 990 to 999 has a minimal effect (ca. 1.5% mass flux reduction), despite the expectation that this increase corresponds to $\rho^{\star\mathrm{s}}$ reducing by another order of magnitude and should reduce leakage mass flux by an order of magnitude in the process. This most likely illustrates that, for a currently unknown reason, the fluid permeation penalisation in the 2D-Island-F problem is unable to realise the order reduction $u^\mathrm{s} \leqslant 10^{-q}\, u_\infty$, as the density is guaranteed to drop to the predefined level at $\hat{\gamma} = 0$.

q	$\hat{q}$	$\Delta\rho^\star$	$\dot{M}$ $(10^{-3}\,\mathrm{kg\,s^{-1}\,m^{-1}})$	U_max $(\mathrm{mm\,s^{-1}})$	$\partial_n U$ $(\mathrm{s^{-1}})$
3	4	990	0.4855	6.5488	3.4713
4	4	990	0.1434	6.5246	2.6611
4	5	990	0.4852	6.5487	3.4712
4	5	999	0.4721	6.5511	3.4761
3	3	990	0.1438	6.5246	2.6616
3	3	999	0.1411	6.5249	2.6621
Discrete			0.0000	6.5923	3.5088

Table 4.2: *Results of additional 2D-Island-F simulations using Darcy-only impermeability penalisation and IGPP.*

The effect of changing $\hat{q}$ independent of q, however, is quite astounding: equalising both parameters reduces flux by a factor of 3 to 4 and also exacerbates velocity gradient error from circa -1% to around -25%.

Having evaluated all generated data in this section on the efficacy of various parameterisations of this linear IGPP formulation in conjunction with several selected parameterisations of Darcy-only and Darcy-with-filtered-Forchheimer fluid permeation penalisation descriptions, the Darcy-only formulation is selected for

all simulations from hereon. It is parameterised with q 4 and $\hat{q}$ 5. The IGPP formulation with linear non-physical density interpolation will be utilised, setting $\Delta\rho^\star$ at 990. This formulation is expected to yield the best compromise between the accuracy of the velocity gradient at the threshold pseudodensity (approximately 1% error from discrete simulations), to aid thermofluidic simulation heat transfer accuracy, and sufficiently restrict non-physical fluid mass flux through solid regions, to mitigate non-physical heat transfer mechanics in thermofluidic simulations.

Furthermore, no overestimation occurs on any of the performance indicators in comparison to the discrete-equivalent simulations, meaning there is no reason to expect any model robustness- or stability-related problems throughout further simulation and optimisation. While no performance indicators exceed the levels of the discrete simulation for $\Delta\rho^\star$ equal to 999 either, the performance deficit of choosing 990 is so minimal, that sacrificing it for additional margin on stability is found favourable.

4.3. Thermofluidic material properties interpolation schemes

With a parameterisation and formulation selected for the fluidic optimisation problem, attention turns to the construction of the material properties that are relevant to solving the thermal governing equation.

4.3.1. Theoretical approach to realistic thermal field solutions in optimisation

As the starting point to refining the continuous-density model for the most accurate heat transfer results, we should consider the governing equation for the temperature field:

$$\mathcal{R}^{\mathrm{t}} = \nabla \cdot \{\kappa \nabla T\} - \rho c_p \mathbf{u} \cdot \nabla T + \dot{q}_{\mathrm{gen}} = 0, \tag{4.6}$$

wherein T is the temperature field to solve for and $\dot{q}_{\mathrm{gen}}$ responsible for the 'influx' of heat into the simulation / optimisation domain – for two-dimensional problems, only. In the case of a real, discretised simulation domain, the thermal conductivity κ has a zero-gradient in all of Ω. While the conductivity is constant within the fluid and the solid regions, respectively, it is, however, only piecewise constant in the entirety of Ω: κ is not $\mathcal{C}^0$-continuous in Ω. In fact, the same is true for all material properties fields in any discretised problem. However, for the other material properties, their solid-element values are not relevant when dealing with steady-state problems concerning the diffusion-advection equation.

Consequently, if one were to further derive from the governing equation, they would find:

$$\mathcal{R}^{\mathrm{t}} = \frac{\nabla\kappa}{\kappa} \cdot \nabla T + \nabla^2 T - \frac{\rho c_p \mathbf{u}}{\kappa} \cdot \nabla T + \frac{\dot{q}_{\mathrm{gen}}}{\kappa} = 0.$$

Considering that the maximum value of temperature field T in Ω will be dependent mostly on the efficacy of heat transfer, that is, the thermal resistance, specifically within fluid-density and intermediate-density regions, rather than the thermal resistance of the solid features, to achieve high simulation accuracies, the expression without heat generation term is a decent approximation to the thermal governing equation, to still obtain adequate results.

Furthermore, since flow velocity field u is defined to be near perpendicular to $\nabla\hat{\gamma}$, it is also near perpendicular to $\nabla\kappa = \kappa'\nabla\hat{\gamma}$. Considering that all heat is (eventually) advected, we consider this term for now irrelevant:

$$\text{For non-solid regions:} \quad \cancel{\frac{\nabla\kappa}{\kappa}}\cdot\nabla T + \nabla^2 T - \frac{\rho c_p \mathbf{u}}{\kappa}\cdot\nabla T + \cancel{\frac{\dot{q}_{\text{gen}}}{\kappa}}^{\approx 0} = 0$$

$$\nabla^2 T - \frac{\rho c_p \mathbf{u}}{\kappa}\cdot\nabla T = 0$$

$$\nabla^2 T - \frac{\mathrm{Pe}_{L'}}{L'}\frac{\mathbf{u}}{|\mathbf{u}|}\cdot\nabla T = 0$$

$$\nabla^2 T - \frac{\mathrm{PrRe}_{L'}}{L'}\frac{\mathbf{u}}{|\mathbf{u}|}\cdot\nabla T = 0$$

where $\mathrm{Pe}_{L'}$ is the dimensionless Péclet number for characteristic length scale L', $\mathrm{Re}_{L'}$ is the Reynolds number of the same length scale, and Pr the Prandtl number. Under the premise that the construction of the interpolation functions for material properties c_p, ρ and κ does not allow for any corrections to the temperature field that originate from an erroneous fluid velocity field beyond possibly reducing advective heat transfer resulting from leakage mass flux, the term $\frac{\mathrm{Pe}_{L'}}{L'}$ should be constructed to result in minimal error between any continuous- and discrete-density problems. Regardless of its size, the characteristic length scale L' will be identical in both problems.

Consequently, where $\dot{q}_{\text{gen}} \approx 0$, results are most accurate when $\frac{1}{L'}\mathrm{PrRe}_{L'}$ is interpolated to be as close as possible to the discrete-equivalent throughout Ω.

Since the Prandtl number is exclusively relevant in the fluid- and intermediate-density regions of the simulation domain, κ and c_p need only be accurate for $\hat{\gamma} \geq 0.5$. As such, one can set c_p constant at c_p^{f}, or 4184 $\mathrm{J\,kg^{-1}\,K^{-1}}$. κ may not be set constant, since the heat influx term into solid elements depends on it. Instead, κ will be interpolated using a hyperbolic tangent function akin to that used for projection of the filtered pseudodensity field. This way, the thermal conductivity interpolation function is differentiable and close to the step function that would be expected through the discrete problem.

For the Reynolds number, only one relevant assumption can be made, namely that $\mathbf{u}$ in any continuous problem is solved with sufficiently good accuracy to the discrete simulation results. So long as that holds, also the Reynolds number need only be accurate in the fluid- and intermediate-density regions, meaning ρ may be kept constant at ρ^{f} (1000 $\mathrm{kg\,m^{-3}}$), since there is no need to compensate for in-solid advection. Finally, the heat influx will be interpolated using RAMP.

$$\rho = \rho^{\text{f}} \tag{4.7a}$$

$$c_p = {c_p}^{\text{f}} \tag{4.7b}$$

$$\dot{q}_{\text{gen}}(\hat{\gamma}) = \dot{q}_{\text{gen}}^{\max}\frac{10^{-q_{\text{heatgen}}(1-\hat{\gamma})}}{10^{-q_{\text{heatgen}}}+\hat{\gamma}} \tag{4.7c}$$

$$\kappa(\hat{\gamma}) = \kappa^{\text{s}} + \left(\kappa^{\text{f}} - \kappa^{\text{s}}\right)\frac{\tanh\left(\beta_\kappa(\hat{\gamma}-\gamma_\kappa)\right)+\tanh\left(\beta_\kappa\gamma_\kappa\right)}{\tanh\left(\beta_\kappa(1-\gamma_\kappa)\right)+\tanh\left(\beta_\kappa\gamma_\kappa\right)} \tag{4.7d}$$

This leaves some parameters to be optimised for: β_κ (the 'sharpness' of the interpolation of κ) and q_{heatgen} (the heat generation RAMP interpolation parameter), specifically. Parameter $\gamma\kappa$ will be fixed at a value of 0.5. Besides these remaining parameters, also the effect of IGPP and the fluid permeation penalisation parameterisation will be simulated to provide some insight into the sensitivity of the modelling accuracy to changes to them.

4.3.2. Reference simulation results for the thermofluidic problem

The figures 4.8 and 4.9 display the results of the performed reference simulations 2D-Plate-TF and 2D-Island-TF, respectively.

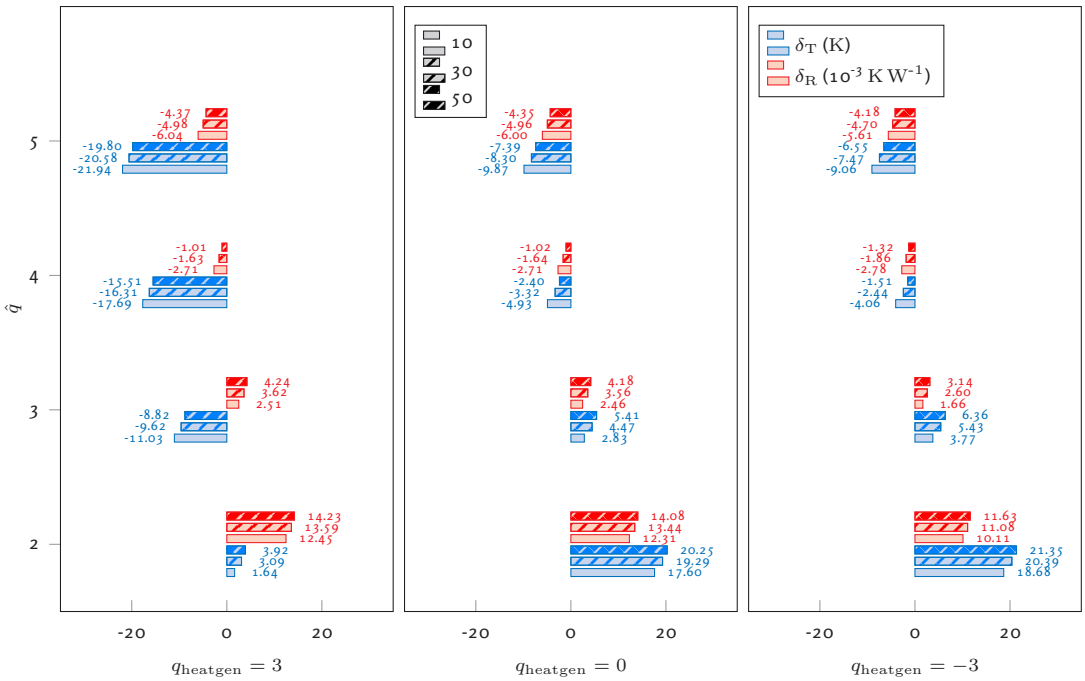

Figure 4.8: *Results of the two-dimensional flat plate flow thermofluidic reference simulation (2D-Plate-TF). Depicted data are the temperature error between the simulated case and the discrete-equivalent (in* K*) in blue and the linearised error in the effective total thermal resistance of the geometry to the heat influx distribution (in* $10^{-3}\,\mathrm{K\,W^{-1}}$*) in red. The data is generated for* $q_{heatgen}$ *equal to 3, 0 and -3 with* $\hat{q}$ *ranging from 2 to 5. For each permutation,* β_κ *is simulated at 10, 30 and 50.* $q = 5$*.*

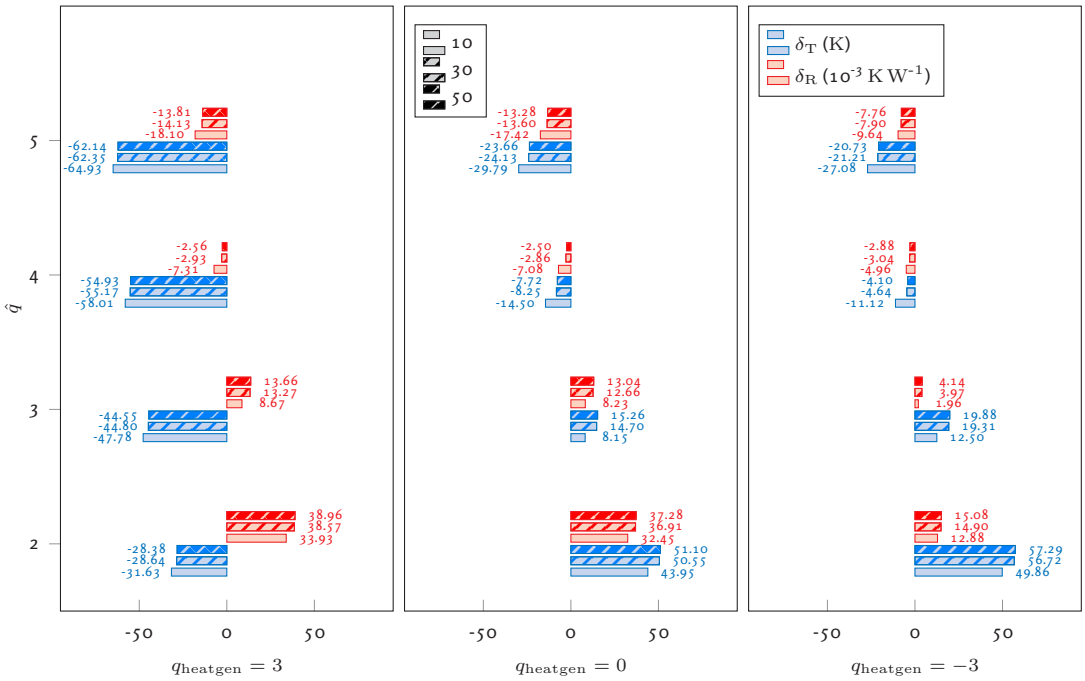

Figure 4.9: *Results of the two-dimensional island flow thermofluidic reference simulation (2D-Island-TF). Depicted data are the temperature error between the simulated case and the discrete-equivalent (in* K*) in blue and the linearised error in the effective total thermal resistance of the geometry to the heat influx distribution (in* $10^{-3}\,\mathrm{K\,W^{-1}}$*) in red. The data is generated for* $q_{heatgen}$ *equal to 3, 0 and -3 with* $\hat{q}$ *ranging from 2 to 5. For each permutation,* β_κ *is simulated at 10, 30 and 50.* $q = 5$*.*

The data presented as ϵ_{R} is a measure of the error of the model's total effective thermal resistance. This total thermal resistance, while it is quite crude and dependent on the specific design under consideration as well as the heat influx distribution, is, at least, compensated for any deficit in the total heat flux as a result of the non-linearity of its interpolation function. Considering this total effective thermal resistance as

$$\mathrm{R}_{\mathrm{th,eff}} = \frac{T_{\max} - T_{\mathrm{ref}}}{\dot{Q}_{\mathrm{tot}}}. \tag{4.8}$$

The corresponding linearised error in the thermal resistance between the discrete-density solution and any other continuous-density simulation result is then found to be:

$$\delta_{\mathrm{R}} = \frac{\delta_{\mathrm{T}}}{\dot{Q}_{\mathrm{tot}}^{\mathrm{disc}}} - \frac{\delta_{\dot{Q}}}{\left(\dot{Q}_{\mathrm{tot}}^{\mathrm{disc}}\right)^2} \left(T_{\max} - T_{\mathrm{ref}}\right)^{\mathrm{disc}}, \tag{4.9}$$

by differentiation and substitution. In this expression δ_{T} is the difference between simulations in maximum recorded temperature in Ω, $\delta_{\dot{Q}}$ the total heat influx deviation to the discrete-equivalent simulation integrated over the simulation domain. For the discrete problems, $T_{\max}$ are 388.12 K and 387.12 K and $\dot{Q}_{\mathrm{tot}}$ 1500.0 $\mathrm{W\,m^{-1}}$ and 1478.6 $\mathrm{W\,m^{-1}}$ for the 2D-Plate-TF and 2D-Island-TF problems, respectively.

From the data, it is evident that the effect of $\hat{q}$, the fluid permeation penalisation interpolation parameter, on the results is substantial. This is a consequence of the – apparently – much higher sensitivity of the accuracy of the temperature field to the leakage mass flux, rather than the velocity gradient at the threshold density. By extension, D-(4,5)-990 is not an adequate formulation to get accurate temperature fields from the continuous-density problem.

Instead, for both problems, the much more accurate results are found for $\hat{q} = 4$ – note that results are generated with $q = 5$ – and q_{heatgen} either 0 or -3. In fact, given the relatively small differences observed between both of these parameterisations, it seems reasonable to assume that the effect of a change in q_{heatgen} in this region is quite limited. However, to increase the likelihood of the optimality of sharp interfaces, choosing q_{heatgen} larger is advantageous, so 0 is selected.

Furthermore, for the 2D-Island-TF problem, limited gain is seen in increasing β_κ from 30 to 50, and so, limited gain is expected in increasing it further. As such, 50 is selected.

All in all, this results in a thermofluidic formulation that provides, at least for these two specific problems, accuracy to within 5% to 8% of the discrete equivalent simulation on the maximum observed temperature and within 1% error in the problems' total effective thermal resistance. We find for this parameterisation that the total heat influx into the simulation domains is 1485.6 $\mathrm{W\,m^{-1}}$ for the flat plate simulation and 1411.7 $\mathrm{W\,m^{-1}}$ for the island geometry, corresponding to approximately a 1% and a 4.5% error, respectively. Considering that specifically the derived effective thermal resistance of the problem, pertaining to the maximum temperature in the domain, is assumed to be the key indicator for the model's accuracy, this circa 1% error greatly exceeds expectations.

Parameter/scheme	Value/selection (SI units)
N_γ	8
β	10
γ_{proj}	0.5
Permeation penalisation	Darcy, RAMP
q	5
$\hat{q}$	4
$\rho^\star$-interpolation	Linear
$\Delta\rho^\star$	990
c_p-interpolation	None, constant

Parameter/scheme	Value/selection (SI units)
c_p	4184
ρ-interpolation	None, constant
ρ	1000
$\dot{q}_{\mathrm{gen}}$-interpolation	RAMP
q_{heatgen}	0
κ-interpolation	Hyperbolic tangent proj.
κ^{f}	0.6
κ^{s}	237
β_κ	50
γ_κ	0.5
μ	$9 \cdot 10^{-4}$

Table 4.3: *Selected model parameters and implementations.*

4.4. Optimisation model behaviour-under-optimisation tests

In order to get some sense of the suitability of any thermofluidic formulation as an optimisation model for density-based topology optimisation, some admittedly fairly rudimentary tests can be performed. Specifically, in this section, the constructed thermofluidic formulation and parameterisation will be tested for several desirable properties: the (strict) monotonicity of the total heat flux and maximum temperature in the simulation domain during the process of island formation in an otherwise fully fluid design; monotonicity of the continuity equation residual during the same island formation process; optimality of the sharpest achievable interfaces in both the 2D-Plate-TF and 2D-Island-TF reference simulations.

While none of these test will guarantee the model to behave as desired during optimisation, they should at least serve to identify and discard any model that would exhibit any obvious undesirable behaviour.

4.4.1. Monotonicity of objective and constraints in island formation

The plots in figure 4.10 show a selection of the data generated through the island formation simulation performed. The simulation is identical to the 2D-Island-TF, though with a parameter sweep performed on the control pseudodensity field's 'island' region density: 0 indicates a fully solid island (in the control density) and 1 corresponds to a fully fluid simulation domain.

Displayed in the upper plot, are the objective, or total heat flux through the domain, and the maximum recorded temperature in Ω for the case where the interpolation parameter q_{heatgen} of the heat influx RAMP-interpolation is set at 0. Note that T_{ref} is set to be $293.15\ \mathrm{K}$.

In the lower plot, total effective thermal resistance data is presented for various values of q_{heatgen} throughout the island formation process. Again, note that the thermal resistance displaced is dependent on both the design and distribution of heat influx into Ω. It is calculated from $T_{\max}$ and $\dot{Q}_{\mathrm{tot}}$ through equation (4.8).

From figure 4.10, it is evident that both the maximum recorded temperature $T_{\max}$ and total heat flux $\dot{Q}_{\mathrm{tot}}$ monotonically increase during the island formation process (equivalent to γ_{island} decreasing from 1 to 0). As a result of the limited numerical accuracy of the model, the derived $\mathrm{R}_{\mathrm{th,eff}}$ is ill-defined as the island control density approaches 1: as the design approaches its fully-fluid form, both the denominator and numerator of the expression in equation (4.8) approach 0. However, we can make the qualitative claim for the behaviour of the problem as it nears $\gamma_{\mathrm{island}} = 1$, that $\mathrm{R}_{\mathrm{th,eff}}$ is non-zero and generally well-behaved. The argument for this can be found in section A.3.

Based on the data presented and the knowledge that $\mathrm{R}_{\mathrm{th,eff}}$ is non-zero as the the design nears its fully-fluid form, the temperature and total heat flux are not only monotonically increasing, they are – on account of the non-zero gradient at $\hat{\gamma} = 0$ in the RAMP interpolation for heat generation into the domain – even strictly

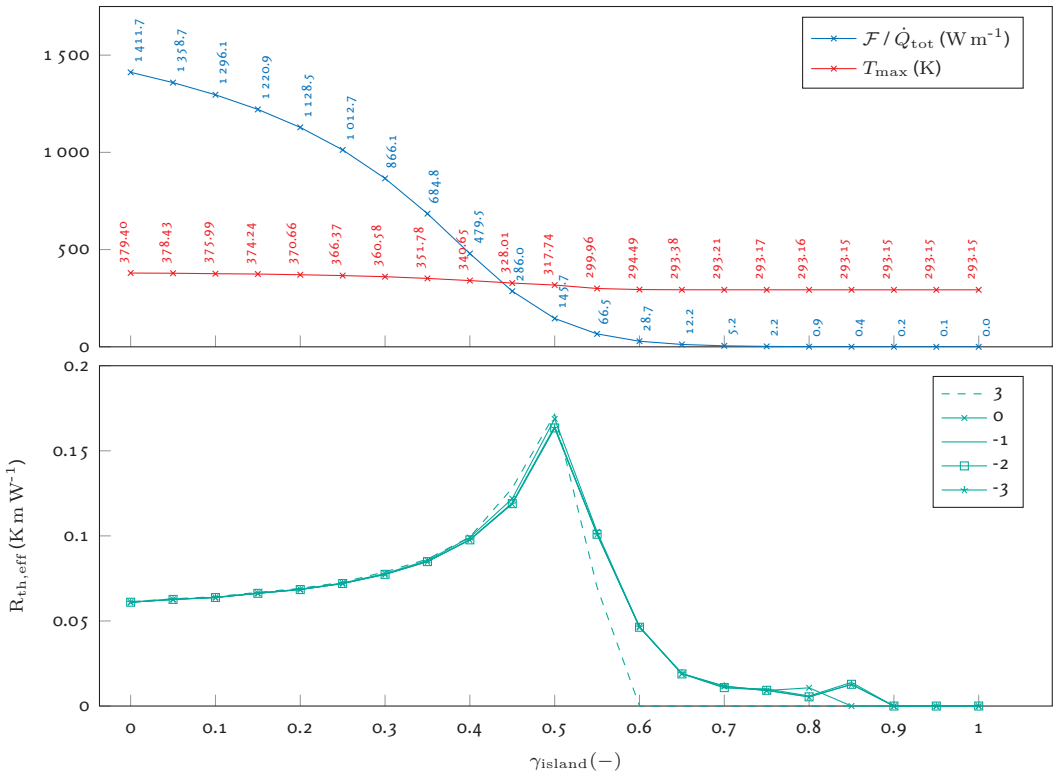

Figure 4.10: *Data generated using the two-dimensional island formation thermofluidic reference simulation. In the upper plot, the total heat flux objective and maximum recorded temperature in the simulation domain are depicted. Results correspond to $q_{\mathrm{heatgen}} = 0$. In the lower plot, the effective total thermal resistance of the geometry is plotted, given the specific heat influx distribution. Results displayed are generated with various values of q_{heatgen}.*

monotonically increasing in for $\gamma_{\mathrm{island}} \in [0, 1]$.

Simultaneously, the fact that 0 and 1 are both local minima in the effective thermal resistance is promising. This, too, may be an indication that the generated designs through optimisation, when using the these model configurations, will tend to 0-1-solutions with minimal intermediate density regions, as those appear to be more 'cost-effective' on the objective function under the complementary constraint.

4.4.2. Optimality of sharpest interfaces

To verify that an objective – be it maximum temperature-based or total heat flux-based – causes the GCMMA iterations to converge onto a design with sharp interfaces is hard. However, in an attempt to gain some insight, both the 2D-Plate-TF- and 2D-Island-TF-problems are simulated with the constructed thermofluidic model, this time for various pseudodensity filter radii in the Helmholtz PDE. This is achieved by variation of N_γ. Note, that N_γ is set at 8 by default. Results are provided in table 4.4.

N_γ	8	10	12	14	16
2D-Plate-TF	0.0624	0.0630	0.0638	0.0648	0.0660
2D-Island-TF	0.0611	0.0651	0.0692	0.0741	0.0803

Table 4.4: *Wall sharpness simulation data. Depicted results are calculated effective total thermal resistance $R_{\mathrm{th,eff}}$ in K m W^{-1} of the problem in the given design configuration with its respective influx distribution based on the heat influx interpolation. Results are generated with decreasing 'crispness' with N_γ increasing from 8 to 16.*

From the results, it is quite clear that increasing the pseudodensity filtering range causes the effective thermal resistance of the design to increase, leading to less favourable objective function evaluations for 'thicker' walls or, equivalently, shallower gradients. Since this effect is observed for both problem geometries, this thermofluidic model will be considered acceptable. The observed behaviour also serves as an indication of the optimality of narrower intermediate-density regions, suggesting that crisper design variable fields should be found favourable to the objective during design optimisation.

5 Results

5.1. Thermofluidic model performance evaluation

This section covers the results generated to quantify the accuracy of the selected thermofluidic optimisation model formulation, the gains in accuracy achieved through the current implementation of IGPP in the reference problems and to quantify the gains in accuracy made when compared to existing schemes. To this end, the formulation and parameterisation of the problem as constructed by Dilgen et al. (2018) is used, but the D-(5,4)-990 fluidic penalisation is applied instead, that is also applied in the constructed optimisation model. This formulation, though likely not the most accurate that exists in literature, is established through its ability to be used in large-scale (ca. 5 million DoFs) optimisation problems. It also appears to have yielded quite well-defined (crisp) designs, devoid of notable grayscaling issues.

In figure 5.1, the numerically solved temperature fields are displayed for all three considered models as well as the solution to the discrete-design problem.

	Discrete	D-(5,4)-990	D-(5,4)-0	Dilgen2018
$T_{\max}$	388.32	385.92	385.59	385.59
$\dot{Q}_{\text{tot}}$	1500.0	1485.6	1485.6	1543.5
$\mathrm{R}_{\text{th,eff}}$	0.0634	0.0624	0.0624	0.0599
Rel. Err. (%)	±00.00	-01.58	-01.58	-05.52

Table 5.1: *2D-Plate-TF reference simulation data for various thermofluidic models. Presented are: the solution to the simulation of the discrete equivalent design, the solution to the constructed model, both with and without IGPP as well as the selected fludic model with a thermal model taken from Dilgen et al. (2018).*

	Discrete	D-(5,4)-990	D-(5,4)-0	Dilgen2018
$T_{\max}$	387.12	379.40	377.08	350.58
$\dot{Q}_{\text{tot}}$	1478.6	1411.7	1411.7	1604.3
$\mathrm{R}_{\text{th,eff}}$	0.0636	0.0611	0.0595	0.0358
Rel. Err. (%)	±00.00	-03.93	-06.45	-43.71

Table 5.2: *2D-Island-TF reference simulation data for various thermofluidic models. Presented are: the solution to the simulation of the discrete equivalent design, the solution to the constructed model, both with and without IGPP as well as the selected fludic model with a thermal model taken from Dilgen et al. (2018).*

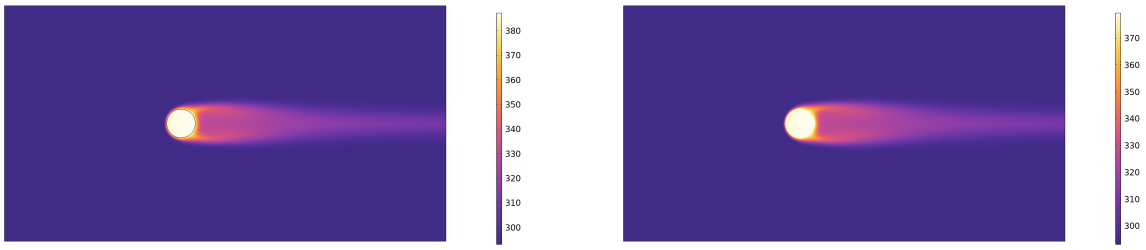

(a) *2D-Island-TF discrete simulation.*

(b) *2D-Island-TF D-(5,4)-990 simulation.*

(c) *2D-Island-TF D-(5,4)-0 simulation.*

(d) *2D-Island-TF Dilgen2018-990 simulation.* $n = 3$.

(e) *2D-Plate-TF discrete simulation.*

(f) *2D-Plate-TF D-(5,4)-990 simulation.*

(g) *2D-Plate-TF D-(5,4)-0 simulation.*

(h) *2D-Plate-TF Dilgen2018-990 simulation.* $n = 5$.

Figure 5.1: *Temperature fields corresponding to the 2D-Island-TF and 2D-Plate-TF simulations performed. Data is displayed for the discrete equivalent simulation, the constructed continuous-pseudodensity model, the constructed model without IGPP and the thermal-problem's material properties interpolations as taken from Dilgen et al. (2018).* $n = 3$ *has been found to yield the most accurate results in the 2D-Island-TF simulation and* $n = 5$ *for the 2D-Plate-TF simulation. In both cases, IGPP is applied and the D-(5,4) flow permeation penalisation scheme is used.*

In figure 5.1, the resulting thermal fields of the simulations performed are presented. From these fields, the only notable observation is that the temperature distribution accuracy of the Dilgen2018-implementation of the thermal model is considerably less accurate than the thermal model that has been constructed in this thesis in the 2D-Island-TF problem. Meanwhile, no notable differences are observed between the four models in the 2D-Plate-TF simulation.

5.2. Flow inverter problem

The flow inverter problem is a problem commonly seen in fluidic topology optimisation, frequently serving as a reference problem to quantify the performance of new techniques in the research field. This section covers a 4:1 aspect ratio flow inverter. Select design iterations are shown in figure 5.2. Inflow occurs at the left boundary of the domain and outflow on the right. A pressure drop constraint of 0.1 Pa is applied for 0.5 $\mathrm{mm\,s^{-1}}$ inflow velocity. The objective is to minimise the horizontal velocity component u at the geometric center of the domain. Figure 5.3 shows the evolution of the objective $\mathcal{F}$ across the simulated design iterations.

The results displayed, specifically the visualised last iteration of the projected pseudodensity field, appear to be quite similar to designs frequently encountered throughout literature. While it is clearly not fully optimised, as can also be deduced from the objective function evaluations plot, some interesting things can be said about these results.

Both the case with and without IGPP converge fast, in part due to the selected initial pseudodensity field. Both models appear to be very well behaved throughout the optimisation, with both objectives converging after a similar number of iterations. It is worthwhile mentioning that the objective value of the model with IGPP reaches slightly lower values, likely due to reduced leakage fluxes through solid regions. This observation is supported by the structures found after the final iteration: the solid wall thicknesses designed into the domain with IGPP are considerably reduced compared to the case without IGPP implemented. This suggests that the model that uses IGPP is able to much more effectively reduce leakage fluxes, even with much thinner solid features.

Also worthwhile mentioning is that, due to the pressure drop having been applied as a constraint, rather than an objective, intermediate density regions of considerable size still remain in the design domain, even after more than 150 iterations.

5.2.1. Discretised design simulation

To evaluate the continuous model once more, figure 5.4 displays the results of the discrete and the continuous optimised design, with respective objective values -13.12 and -14.73 $\mathrm{mm\,s^{-1}}$. Both simulations satisfy the pressure drop constraint, with 76.2 mPa and 100.0 mPa for the thresholded and the continuous-design variable simulations, respectively.

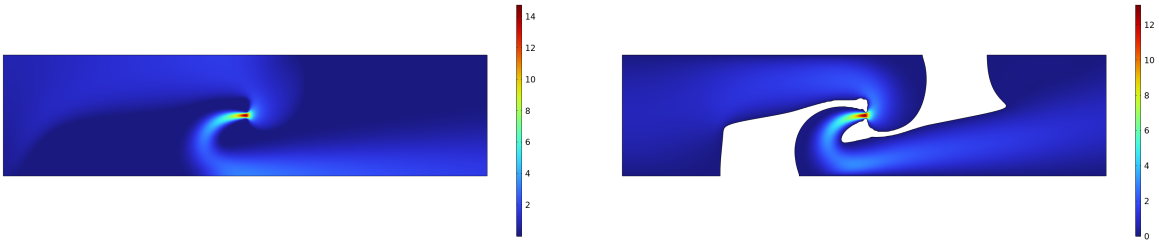

(a) *Velocity field of final iteration (158) with D-(3,2)-990 model.* **(b)** *Velocity field of the discretised design.*

Figure 5.4: *Comparison of the continuous and discrete design after flow inversion optimisation. Discretisation, or thresholding, occurs at 0.5 projected pseudodensity.*

The observed error in the continuous simulation, then, is ca. 31% overestimation on the overall pressure

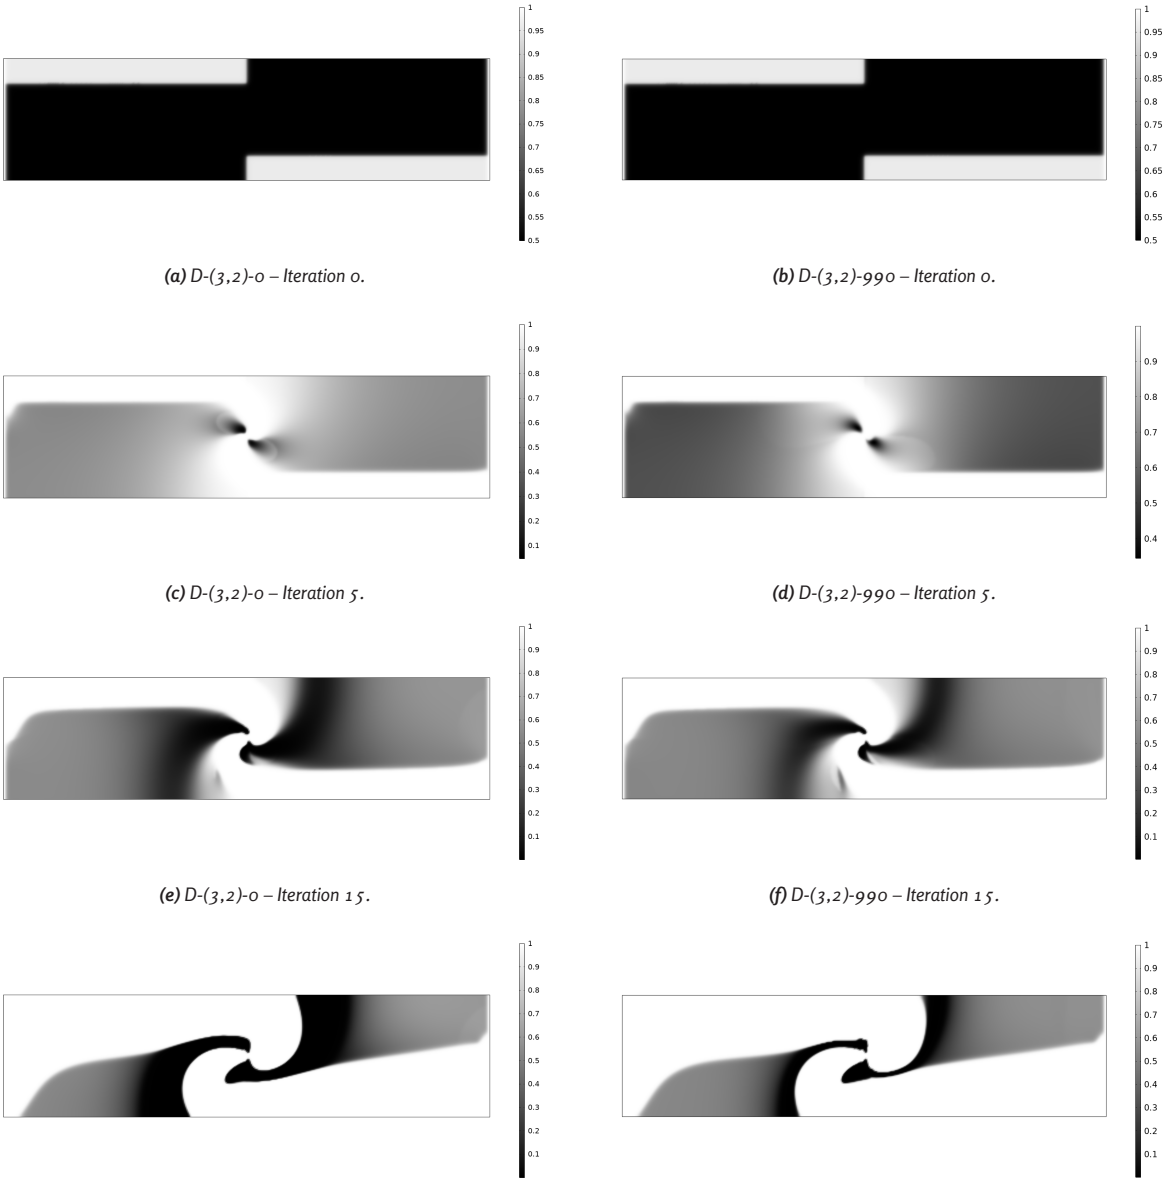

(a) D-(3,2)-0 – Iteration 0.

(b) D-(3,2)-990 – Iteration 0.

(c) D-(3,2)-0 – Iteration 5.

(d) D-(3,2)-990 – Iteration 5.

(e) D-(3,2)-0 – Iteration 15.

(f) D-(3,2)-990 – Iteration 15.

(g) D-(3,2)-0 – Iteration 167.

(h) D-(3,2)-990 – Iteration 158.

Figure 5.2: *Projected pseudodensity of the flow inverter optimisation problem. Optimisation is performed with D-(3,2)-0 and D-(3,2)-990 fluidic penalisation interpolations. Iterations 0, 5, 15 and 167/158 (the final iteration) are shown.*

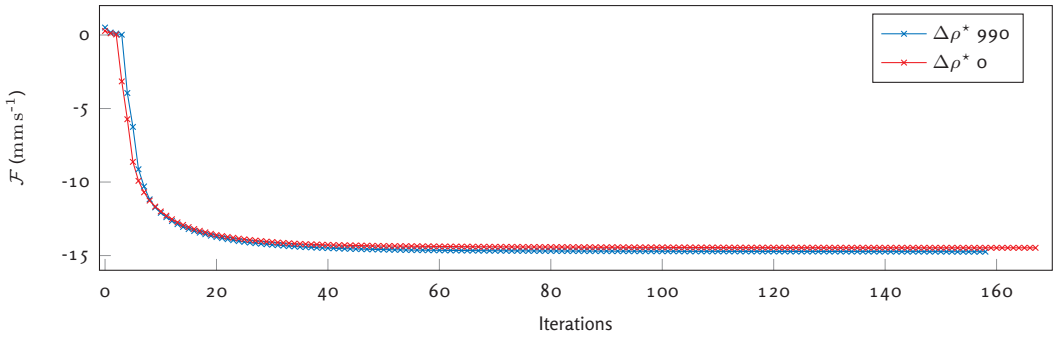

Figure 5.3: *Flow inverter optimisation objective function evolution for D-(3,2)-990 and D-(3,2)-0.*

drop and approximately 12% overestimation on the objective. The particularly large error found on the pressure drop can, in part, be attributed to the intermediate-pseudodensity region upstream of the inversion features. However, the majority of this error will be a direct consequence of the error in the fluid velocity in the restriction at the inversion point.

5.3. Thermofluidic topology optimisation problems

Unfortunately, it has not been achieved to generate meaningful results in heat transfer optimisation problems. For an as of yet unknown reason, the implementation of the algorithm in COMSOL, with the proposed thermofluidic model, is unable to completely define the required Jacobian through sensitivity analysis to make meaningful design iterations. This issue appears to originate specifically at the sensitivity analysis of temperature field.

All of the attempted problem geometries have failed in generating any designs, filling the optimisation domain with a near-uniform intermediate-density. Any changes made to the thermal resistance objective, the pressure drop constraint or the proposed model have not resolved this problem. Also any changes made to implement various parameterisations of continuation methods onto the fluid permeation penalisation and IGPP, have not resulted in successful optimisation.

While the optimisation problems are not intended to serve as further evaluation of model accuracy, they are intended to provide evidence of and insight into the functionality of the model for use in an optimisation case. Since these results are missing, no evidence is available to support or discredit the current model's suitability to topology optimisation problems specifically.

Ultimately, the successful implementation of this or a further-derived model – in COMSOL or in other software – is left as future work.

6

Discussion

6.1. Evaluation of model validity

This section covers the discussion on the validity and accuracy of the fluidic and thermofluidic models constructed and tested in this thesis.

6.1.1. The fluidic model and IGPP

Considering the data from sections 4.1, 4.2 and 5.2, some interesting deductions can be made.

Starting off with the fluid permeation penalisation for fluids in solid elements, utilising the methodology as outlined by Theulings, Noël, et al. (2025) has proven to be very successful. Indeed, as the work suggests and presents, their methods allow for the robust control of fluid velocities in fully-solid regions through the parameterisation of q and $\hat{q}$, where the former realises an order-of-magnitude reduction upon integer increment and the latter is the RAMP-interpolation parameter. After having trialled all three of their presented methods, the Darcy-with-Forchheimer method was excluded on the basis of stability concerns – a problem also outlined in the paper – and the Darcy-with-filtered-Forchheimer formulation was excluded on account of the considerable computational overhead. Though, it should be mentioned that the improved convergence that the DFF-formulation may offer could prove to ultimately yield faster iteration times and ultimately a shorter overall optimisation runtime. This would require further investigation.

All in all, in so far as this thesis serves as an independent replication of the results shown by Theulings, Noël, et al. (2025), the results generated are seemingly in agreement with their assessments: for sufficiently small q, order reduction is indeed achievable, though at values beyond 3, efficacy of the method drops off quite significantly for each of the three formulations. Furthermore, DFF shows promise as a more effective penalisation model, with it leading to generally more accurate results for relevantly large Reynolds numbers, though Darcy, making use of the a-priori elemental Reynolds number estimation, works well for, at least, the tested subset of problems and parameterisations.

Ultimately, based on the simulations performed in section 4.1, the Darcy formulation was selected with parameterisation (4, 5). In hindsight, through further evaluation after incorporation of the gradient-parallel penalisation scheme in section 4.2, this was found to be an error. Thermofluidic simulations performed with parameterisation (5, 4), specifically with $q - \hat{q} = 1$, have generated far superior results.

Nonetheless, even with the quite different in-material flow penalisations D-(3,2) – note that $q - \hat{q} = 1$ – and D-(5,3), the newly introduced implicit gradient-parallel penalisation (IGPP) approach shows promise. While its current linear interpolation is a crude implementation, at best, still it shows its ability in reducing flow into solid regions and should, in theory, function as a hard limit on gradient-parallel flow as it is part

of the numerical residual of the fluid continuity equation. This, however, is also a possible drawback: if implementation is done incorrectly or not matched properly to the utilised fluid permeation penalisation, convergence and stability issues are easily introduced into any problem geometry. Conversely, considering topology optimisation generally relies on being able to explore designs that are not physically realisable, building a fluid model that is too harsh on such designs will also not aid the optimisation process.

Unfortunately, in spite of any effort, no satisfactory approach has been derived successfully that enables constructing robust and stable formulations to the non-physical $\rho^\star$. This holds true for the strength of penalisation, its interpolation as well as any dependencies on any selected in-solid flow penalisation. This means that the selection of any parameterisation in this thesis relies exclusively on the data presented in section 4.2.

Despite all of these issues, still, IGPP has shown great promise; quite effectively penalising flow into (and out of) solid element regions in the reference problems and thereby reducing leakage mass flux for values $\Delta\rho^\star$ of 990 and 999. Even at smaller values, however, its effect is notable for fluid permeation penalisation schemes with high interpolation parameters $\hat{q}$. However, simultaneously, it allows for more realistic fluid velocity profiles near the threshold value of the projected pseudodensity field, whereas this would require intricate tuning of q and $\hat{q}$ and thereby results in lesser accurate flow field solutions. To alleviate possible convergence issues the somewhat more conservative value of 990 is recommended for $\Delta\rho^\star$.

While the selected fluidic penalisation formulation is not intended for use in purely fluidic optimisation, but rather for utilisation in conjunction with a thermal model for thermofluidic topology optimisation problems, it is still interesting to explore the results of the flow inverter optimisation problem in section 5.2. The resulting designs did not present with particularly well-defined design variable fields, nor was grayscaling properly avoided. Likely, the addition of the pressure drop constraint as an objective instead, would have further reduced these intermediate-pseudodensity regions. Nevertheless, the optimisation results have clearly demonstrated that there is significant potential in IGPP regarding fluidic topology optimisation problems.

6.1.2. The thermofluidic model

While the performance of the fluidic penalisation scheme has not been simulated for and evaluated explicitly, its relevant performance metrics are captured in the heat transfer reference problems implicitly. This, in conjunction with the constructed thermal model outlined in section 4.3, also realised the results found in section 5.1.

The thermal model is constructed to require minimal interpolation and, consequently, tuning. With both pseudodensity-independent heat capacity and density, only the thermal conductivity of the elements and the heat generation still require interpolation. As a reasonable compromise between model accuracy and the optimisation behaviour of the model, the RAMP interpolation parameter q_{heatgen} for the heat influx is set at 0. To achieve the most accurate possible results, the thermal conductivity of any element should be as close as possible to the discrete equivalent model, however. To this end, its interpolation is constructed using a hyperbolic tangent interpolation, much like the projection function that is applied to the filtered pseudodensity field. Its parameterisation of β_κ, the slope of the function at the threshold projected pseudodensity 0.5, is taken as 50. While higher values correspond to better results, setting this value too high will result in model behaviour that is unstable due to the localised high sensitivity of the problem to elements near the threshold pseudodensity.

Quite clearly, the results obtained through the reference simulations on the proposed density-based thermal model, especially in comparison to existing thermal models such as the one utilised by Dilgen et al. (2018) for circa 5M-DoF optimisation, are very promising. With errors in the reference simulations found to generally be around 5%, compared to typically up to 30% errors for existing models or roughly 15% in the case of the model of Dilgen et al. (2018), the proposed techniques and implementation would be considered a significant improvement and has exceeded initial expectations of achievable modelling accuracy using this density-based optimisation model.

6.2. Notes on the problems with thermofluidic optimisation

Considering the optimisation problems have not been successfully executed, no further conclusions can be drawn about the accuracy of any optimised structures, nor is any information available on the behaviour of the model, in combination with the objective and constraints, within the optimisation algorithm. However, as alluded to in section 5.3, there are still some meaningful observations to be highlighted.

It should be mentioned, first of all, that the failure to make the thermofluidic optimisations working, is an issue caused by time constraints. The available timeframe of this thesis had simply run out prior to the underlying problem having been identified and resolved.

Regarding this problem, it is most likely one based in user error in the implementation of the model and solver into the software. There are several indications for this to be the problem: for one, since other optimisation and simulation runs have worked without problems, likely, the error is a one-off with the particular implementation of the thermofluidic topology optimisation. While several checks have been performed, including reconstructing the entire model from scratch, the problem has not been successfully identified.

Simultaneously, it is unlikely to be a problem with the choice in objective function or constraints: changing the objective or constraint equations has not been found to enable COMSOL to find sensitivities to the temperature field. Moreover, also changes made to the model itself – reduction of penalisation factors and using interpolation functions with reduced gradients – have not helped. Though comparative inspections of the entire algorithm implementation to, for instance, the fluidic optimisation problem have been conducted, no discrepancies have been observed. Even so, this issue is expected to be an implementation error through the principle of parsimony: it offers the simplest explanation of all observations, despite the inspection of the models against one that is known to work. Realistically, an implementation error has been overlooked.

Finally, given the demonstrated potential gains in modelling accuracy for fluid systems, further study to implement IGPP successfully in thermofluidic optimisation problems could well prove a worthwhile endeavour.

7
Conclusion

7.1. Conclusions on remaining research questions

In section 1.1 the main research question is posed:

How can thermal and fluidic requirements be combined into a sufficiently accurate multidisciplinary topology optimisation method?

In section 2.3, after the literature review, the initial research questions have been revisited and revised. The resulting open-ended sub-questions may be addressed:

- *How may a sufficiently accurate fluid velocity field solution be achieved? Is the work of Theulings, Noël, et al. (2025) adequate, or are additional steps required?*
- *How can the various material properties be interpolated for density-based topology optimisation to achieve accurate temperature field results?*
- *How does the proposed approach perform on reference problems and optimised designs?*

To address the sub-questions, in order of appearance:

- The method introduced by Theulings, Noël, et al. (2025) is a useful starting point towards the construction of a model that enables effective fluid permeation penalisation. To achieve the necessary amount of order-of-magnitudes in fluid velocity in solid regions for accurate temperature field solutions, IGPP has been introduced as a new technique to realise further penalisation to achieve improved velocity field solutions;
- Starting from the thermal governing equation, minimal material properties have had to be interpolated through the analysis of the non-dimensionalised PDE. Through this approach, a simple-yet-effective thermofluidic model has been constructed in which only the thermal conductivity is an interpolated material property;
- The complete thermofluidic density-based model that has been constructed has been verified to reach a roughly 5% error on heat flux values and temperature field solutions. This error is attributed in large part to the 'blurred' heat influx in the intermediate density regions, as the errors on the thermal resistance of the reference simulations are found at 1.5 and 6.5%, for the 2D-Plate-TF and 2D-Island-TF geometries, respectively. For comparison, the thermofluidic interpolation scheme of Dilgen et al. (2018) implemented with the constructed fluid permeation penalisation and IGPP, sees errors of 5.5% and 43.7%, respectively, when compared to the discrete-equivalent model.

Then, to answer the main research question, the results from the various sub-questions may be extended:

For any density-based topology optimisation method or implementation to generate results that are accurate to, at least, results generated by the 'true' simulation of the same design after post-processing, i.e. in its physically manufacturable form, several components are key. First off, the temperature field is calculated, through the advection-diffusion equation and includes the velocity field solved for prior. And so, the emphasis in the pursuit of accurate heat transfer results should lie on the construction of material properties interpolations that are as accurate as possible in the simulation domain to those found in the equivalent discrete problem as well as an accurate solution of the fluid velocity field.

To this end, the interpolation functions have been derived effectively from the thermal governing equation. This has been done using some simplifying assumptions and differentiable approximations to the theoretical (discrete) material properties.

Thereafter, to generate the fluid velocity fields with the necessary accuracy for any given continuous-pseudodensity design, two methods are employed: reduced-tuning fluid permeation penalisation on fluid flow in solid-designed elements and IGPP, implicit gradient-parallel penalisation on fluid flow along gradients in the design, or projected pseudodensity field. IGPP has not so much enabled the more accurate solving of the velocity field independently, but, specifically, IGPP has been instrumental in enabling higher penalisation factors in the fluid permeation penalisation, reducing flow in solid regions of the design. This way, the non-physical advection in solid elements has been effectively mitigated in the reference simulations.

So, then, with the formulations as outlined in chapter 4, it is clear that, despite a lack of data on optimised thermofluidic topologies, the constructed model yields accurate heat transfer results for the devised reference problems. In fact, errors between the thermal resistance modelled in the optimisation simulation and the post-processed design simulation are reduced to circa 5% across simulations. Errors on the model's thermal resistance – an arguably better metric for thermofluidic model accuracy – are found to be in the range of 1.5 to 6 % for the 2D-Plate-TF and 2D-Island-TF reference simulations.

Specifically on the techniques for fluidic penalisation covered in this thesis, some additional remarks are in order.

Theulings, Noël, et al. (2025) have provided, with their paper, work that was instrumental to this research. This thesis has managed to implement their work and has not made any observations of stability or robustness issues beyond those relayed by the authors of the paper.

Then, regarding the implementation of Implicit Gradient-Parallel Penalisation of fluid flow, this thesis has only explored the possibility of incorporating a linear interpolation to the non-physical density $\rho^\star$. While the potential of this technique has been demonstrated, at least to some degree, through the reference problem results, further exploration of improved interpolations could yield significant returns in terms of computational effort, numerical stability, model robustness and model accuracy. As a further note, the flow inverter optimisation case has demonstrated the desirable effects that IGPP seeks to realise compared to a fluidic optimisation model consisting only of a permeation penalisation scheme. This means that the efficacy of IGPP is demonstrated, despite the lack of implementation into a thermofluidic optimisation problem in this thesis.

7.2. Recommendations for future research

Despite any potential progress made in this thesis to the field of thermofluidic topology optimisation research, several open-ended questions and unsolved problems with this methodology remain. Moreover, several clear extensions to the presented implementation and techniques can be considered in any future research.

First and foremost, the thermofluidic model as it is presented in this thesis should be applied to a thermofluidic optimisation problem. This might well turn out to require modifications to the interpolation function used in the IGPP scheme, or a variation of the parameters used in the Darcy penalisation scheme. While the

discussion in section 6.2 possibly presents an argument against the need to change the parameterisation of the model, this is no guarantee. It would be advisable to simply rebuild the model and an optimisation problem at first, from scratch and possibly in a different software, to verify if it can be made to work, rather than trying to identify what is likely a specific implementation error.

Furthermore, a thermofluidic optimisation implementation of the model should be prepared to verify its accuracy in an optimisation setting: the usefulness of the model is solely based on its ability to be utilised in practical applications. After optimisation results have been found to be successfully generated, a further study on their accuracy to their equivalent post-processed discrete-density design would be ideal in an effort to further validate the claims of accuracy made in this thesis.

Concerning the model and approach as they are presented in this thesis: in some cases, convergence issues have been observed with higher penalisation factors. The mitigation of these issues requires further investigation, but lies beyond the scope of this thesis. The implementation of continuation methods into the proposed model formulation is likely necessary to circumvent these issues.

Besides this, quite some extensions to the model as it is presented in this thesis as well as its techniques can be thought of. Perhaps most obviously, a framework that enables proper interpolation of the non-physical density used in IGPP would be a good addition. Instead of the current linear interpolation, an interpolation that is constructed to match a fluid permeation penalisation scheme such as Darcy-only or the Darcy-with-filtered-Forchheimer, possibly even be parameterised using the same variables q and $\hat{q}$, might prove very effective in reducing convergence problems and reducing fluid velocity through 'solid' elements.

Concerning the penalisation schemes against in-material flow, a study into the additional computational complexity of the DFF- and possibly DF-implementations, when compared to the more conventionally used Darcy penalisation, would be useful to increase understanding of these methods and of the feasibility of adopting them into a larger-scale design case.

The implementation of IGPP in turbulent flow problems and in compressible flow problems also might be worthwhile investigating. Considering the compressible Navier-Stokes equations are utilised to enable IGPP, the construction of a multivariate function for the 'non-physical' density used that depends on the element pseudodensity, as well as the fluid pressure (and possibly temperature), is necessary, but might prove to be a challenge.

If the scope was further expanded, a study into the functionality and performance of the constructed model in (thermally) transient problems, too, could prove to be a valuable endeavour. Specifically in higher-performance applications, the effect of thermal load variations might prove to have a large impact on the designed optimal topologies, as the design iteration algorithm could, for instance, strategically place material to increase thermal inertia, effectively modifying the thermal frequency response of any given design.

7.3. Closing remarks

Throughout this study on thermofluidic modelling and its implications for topology optimisation, some fairly fundamental issues have been observed.

Starting off with penalisation of flow in 'solid' elements. While the method by Theulings, Noël, et al. (2025) presents a leap in the effective penalisation of flows for moderate Reynolds numbers and for laminar flows, in particular, it, much like other penalisation schemes for fluid flow optimisation, still relies on some minimal leakage velocity for penalisation to occur. As penalisation factors increase, this required leakage velocity (and the resulting mass flux) decreases, but does not disappear.

While also IGPP can help reduce this leakage flux, still this increase in fluid permeation penalisation brings with it at least three problems: for one, through this method, achieving the 5 - 7 orders of magnitude of leakage flux reduction in solid elements compared to fluid elements for accurate thermal modelling and realistic optimised topologies, is nearly, if not fully, unachievable. Secondly, in increasing the penalisation factors and with reducing velocities, numerical errors become large compared to the 'real' velocity field. This

leads to convergence and stability issues that are inherent to the model. Finally – IGPP is particularly susceptible to this – harsh penalisation may easily become too harsh: especially without proper continuation methods to gradually increase penalisation during the optimisation routines, the algorithm will converge immediately on a local 'optimum' or the solver will simply not converge.

To mitigate some of these problems, it is possible to imagine the construction of a penalisation that does not force higher penalisation factors on terms linearly dependent on the local fluid velocity. Instead, perhaps, a penalisation that consists of lower-order terms of the velocity will prove more effective: a penalisation of the square root or even cube root of the local absolute velocity or filtered velocity, could help reduce the size of penalisation terms and reduce numerical instability and divergence issues in the process in solid regions, when velocities become particularly small.

Finally, some remarks also pertain to topology optimisation for use in thermofluidic design problems and the selection of density-based methods for these problems. In general, density-based approaches are fundamentally poorly suited to heat transfer problems where heat is transferred across solid-fluid interfaces, given its innate inability to define crisp interfaces. Perhaps with the exception of implementations that make use of advanced FEA techniques, such as XFEM or equivalent XFVM implementations, optimisation methods for which such interfaces are inherently defined, level-set methods, for example, will always be more accurate in modelling heat transfer.

Moreover, should flow velocities increase in any design problem, and Reynolds numbers with it, or when dealing with a different fluid medium, the introduction of turbulence into the model may become unavoidable to achieve model accuracy. Considering the high sensitivity of turbulent boundary layer modelling accuracy to simulation domain meshing, specifically also considering the criticality of capturing the viscous sublayer with sufficient accuracy to achieve desired model heat transfer accuracy levels, density-based methods are likely never going to be able to model heat transfer within any reasonable margins. At least, not unless utilising extensive re-meshing every design iteration with an exceptionally robust algorithm to do so.

Methodologies to account for the knock-on effects – consider for instance the effect of re-meshing on sensitivity analysis for gradient-based design iteration schemes such as (GC)MMA – are not ready and, frankly, computational power, as it stands, would likely not be able to optimise such structures with any reasonable degree of model accuracy regardless. Considering that the computational power of modern-day clusters still leads to weeks of optimisation time even on modestly-sized two-dimensional problems with isothermal flow, albeit using discrete evolutionary methods, the computational overhead of the required additional steps is, with the state of current computing power, likely unmanageable within any reasonable timeframe.

References

Alexandersen, J. and C. S. Andreasen (2020). "A Review of Topology Optimisation for Fluid-Based Problems". In: *Fluids* 5.1, p. 29. DOI: 10.3390/fluids5010029.

Bendsøe, M. P. and N. Kikuchi (1988). "Generating optimal topologies in structural design using a homogenization method". In: *Computer Methods in Applied Mechanics and Engineering* 71.2, pp. 197–224. DOI: 10.1016/0045-7825(88)90086-2.

Bruns, T.E. (2007). "Topology optimization of convection-dominated, steady-state heat transfer problems". In: *International Journal of Heat and Mass Transfer* 50.15–16, pp. 2859–2873. DOI: 10.1016/j.ijheatmasstransfer.2007.01.039.

COMSOL AB (2025). *About the Stationary Solver*. URL: https://doc.comsol.com/5.5/doc/com.comsol.help.comsol/comsol_ref_solver.27.096.html (visited on 08/27/2025).

Dilgen, S. B. et al. (2018). "Density based topology optimization of turbulent flow heat transfer systems". In: *Structural and Multidisciplinary Optimization* 57.5, pp. 1905–1918. DOI: 10.1007/s00158-018-1967-6.

Duan, X., Y. Ma, and R. Zhang (2008a). "Optimal shape control of fluid flow using variational level set method". In: *Physics Letters A* 372.9, pp. 1374–1379. DOI: 10.1016/j.physleta.2007.09.070.

— (2008b). "Shape-topology optimization for Navier–Stokes problem using variational level set method". In: *Journal of Computational and Applied Mathematics* 222.2, pp. 487–499. DOI: 10.1016/j.cam.2007.11.016.

— (2008c). "Shape-topology optimization of stokes flow via variational level set method". In: *Applied Mathematics and Computation* 202.1, pp. 200–209. DOI: 10.1016/j.amc.2008.02.014.

Gersborg-Hansen, A., O. Sigmund, and R.B. Haber (2005). "Topology optimization of channel flow problems". In: *Structural and Multidisciplinary Optimization* 30.3, pp. 181–192. DOI: 10.1007/s00158-004-0508-7.

Guest, J. K. (2009). "Topology optimization with multiple phase projection". In: *Computer Methods in Applied Mechanics and Engineering* 199.1–4, pp. 123–135. DOI: 10.1016/j.cma.2009.09.023.

Guest, J. K., A. Asadpoure, and S. Ha (2011). "Eliminating beta-continuation from Heaviside projection and density filter algorithms". In: *Structural and Multidisciplinary Optimization* 44.4, pp. 443–453. DOI: 10.1007/s00158-011-0676-1.

Guest, J. K., J. H. Prévost, and T. Belytschko (2004). "Achieving minimum length scale in topology optimization using nodal design variables and projection functions". In: *International Journal for Numerical Methods in Engineering* 61.2, pp. 238–254. DOI: 10.1002/nme.1064.

Haertel, J. H.K. and G. F. Nellis (2017). "A fully developed flow thermofluid model for topology optimization of 3D-printed air-cooled heat exchangers". In: *Applied Thermal Engineering* 119, pp. 10–24. DOI: 10.1016/j.applthermaleng.2017.03.030.

Holka, Q. et al. (2022). "Density-based topology optimization of a surface cooler in turbulent flow using a continuous adjoint turbulence model". In: *Structural and Multidisciplinary Optimization* 65.8. DOI: 10.1007/s00158-022-03311-z.

Li, X. et al. (2022). "Multi-material topology optimization of transient heat conduction structure with functional gradient constraint". In: *International Communications in Heat and Mass Transfer* 131, p. 105845. DOI: 10.1016/j.icheatmasstransfer.2021.105845.

Lohan, D. J., E. M. Dede, and J. T. Allison (2019). "A study on practical objectives and constraints for heat conduction topology optimization". In: *Structural and Multidisciplinary Optimization* 61.2, pp. 475–489. DOI: 10.1007/s00158-019-02369-6.

Matsumori, T. et al. (2013). "Topology optimization for fluid–thermal interaction problems under constant input power". In: *Structural and Multidisciplinary Optimization* 47.4, pp. 571–581. DOI: 10.1007/s00158-013-0887-8.

Noël, L. and K. Maute (2022). "XFEM level set-based topology optimization for turbulent conjugate heat transfer problems". In: *Structural and Multidisciplinary Optimization* 66.1. DOI: 10.1007/s00158-022-03353-3.

Stolpe, M. and K. Svanberg (2001). "An alternative interpolation scheme for minimum compliance topology optimization". In: *Structural and Multidisciplinary Optimization* 22.2, pp. 116–124. DOI: 10.1007/s001580100129.

Svanberg, K. (1987). "The method of moving asymptotes—a new method for structural optimization". In: *International Journal for Numerical Methods in Engineering* 24.2, pp. 359–373. DOI: 10.1002/nme.1620240207.

— (2002). "A class of globally convergent optimization methods based on conservative convex separable approximations". In: *SIAM J. Optim.* 12.2, pp. 555–573. DOI: 10.1137/S1052623499362822.

Theulings, Langelaar, et al. (2023). "Towards improved porous models for solid/fluid topology optimization". In: *Structural and Multidisciplinary Optimization* 66.6. DOI: 10.1007/s00158-023-03570-4.

Theulings, Noël, et al. (2025). "Reducing parameter tuning in topology optimization of flow problems using a Darcy and Forchheimer penalization". In: *Computer Methods in Applied Mechanics and Engineering* 443, p. 118027. DOI: 10.1016/j.cma.2025.118027.

Trillet, D., P. Duysinx, and E. Fernández (2021). "Analytical relationships for imposing minimum length scale in the robust topology optimization formulation". In: *Structural and Multidisciplinary Optimization* 64, pp. 2429–2448. DOI: 10.1007/s00158-021-02998-w.

Wu, S., Y. Zhang, and S. Liu (2019). "Topology optimization for minimizing the maximum temperature of transient heat conduction structure". In: *Structural and Multidisciplinary Optimization* 60.1, pp. 69–82. DOI: 10.1007/s00158-019-02196-9.

Yaji, K. et al. (2015). "A topology optimization method for a coupled thermal–fluid problem using level set boundary expressions". In: *International Journal of Heat and Mass Transfer* 81, pp. 878–888. DOI: 10.1016/j.ijheatmasstransfer.2014.11.005.

Yoon, G. H. (2010). "Topological design of heat dissipating structure with forced convective heat transfer". In: *Journal of Mechanical Science and Technology* 24.6, pp. 1225–1233. DOI: 10.1007/s12206-010-0328-1.

Zhou, M., Y. K. Shyy, and H. L. Thomas (2001). "Checkerboard and minimum member size control in topology optimization". In: *Structural Multidisciplinary Optimization* 21.2, pp. 152–158. DOI: 10.1007/s001580050179.

Zillober, C. (1993). "A globally convergent version of the method of moving asymptotes". In: *Structural Optimization* 6.3, pp. 166–174. DOI: 10.1007/bf01743509.

A

Derivations and arguments

A.1. Argument for the thermal resistance independencies

Consider the conduction-advection governing equation for the temperature field, see equation (3.4):

$$\nabla \cdot (\kappa \nabla T) - \rho c_p \mathbf{u} \cdot \nabla T + \dot{q}_{\text{gen}} = 0. \tag{A.1}$$

Note that this equation is linear in the temperature field T and gradient ∇T and heat influx (field) q_{gen}. Assume, then, that a design remains unchanged, but heat influx is increased by some factor ξ in its existing distribution, both by the linear scaling of the field and by the linear scaling of the relevant heat flux boundary conditions on the domain, both with ξ. Note, then, that through

$$\{\kappa \nabla + \nabla \kappa - \rho c_p \mathbf{u}\} \cdot \nabla T = -\dot{q}_{\text{gen}} \tag{A.2}$$

shows a linear relation between the temperature gradient vector field and the heat influx field (and / or influx boundary conditions). As such, scaling factor ξ, will simply result in scaling ∇T to $\xi \nabla T$. By extension, since the domain is unchanged, the temperature field is scaled, such that $T - T^{\text{ref}}$ becomes $\xi \left(T - T^{\text{ref}}\right)$. This also implies that any linear, even non-differentiable, function of T, denoted $T^\star$ as the control temperature, such as the global maximum temperature or a region average temperature are scaled to $\xi \left(T^\star - T^{\text{ref}}\right)$.

Finally, one finds that any scaling factor does not affect the quantity

$$\mathrm{R}_{\text{th,eff}} = \frac{\xi \left(T^\star - T^{\text{ref}}\right)}{\xi \dot{q}_{\text{gen}}} = \frac{T^\star - T^{\text{ref}}}{\dot{q}_{\text{gen}}} \tag{A.3}$$

so long as boundary conditions are scaled accordingly (to allow for identical thermal load distribution) and the design is unchanged. For this argument, the heat influx distribution may also be integrated over Ω to yield the total influx $\dot{Q}_{\text{tot}}$, to arrive at the same conclusion, if the influx distribution is kept constant implicitly as a condition to the total influx $\dot{Q}_{\text{tot}}$.

A.2. Derivation of the filtered pseudodensity field in the 2D-Island-F/TF reference problems

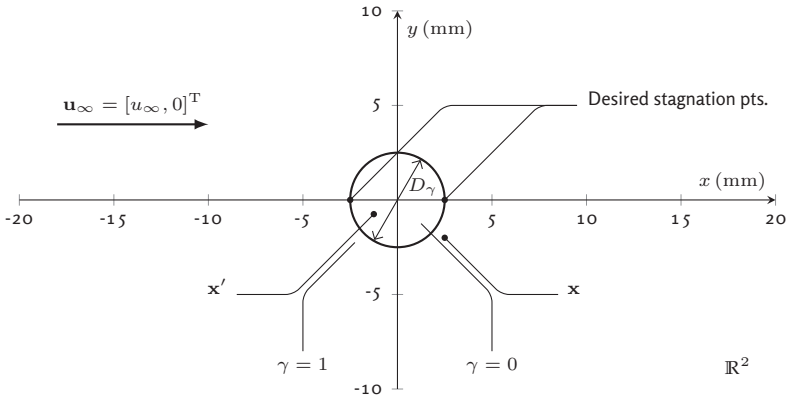

Figure A.1: *Schematic representation of the 2D-Island reference problem specification. The design, characterised by control variable field γ, is axisymmetric and makes a discrete step, to result in the sharpest possible interface. The desired positions of stagnation points along the x-axis are also depicted.*

From the axisymmetric nature of γ in the 2D-Island reference problems, as depicted in figure A.1, the control variable field γ can be defined:

$$\gamma(r) = H\left(r - \frac{D_\gamma}{2}\right). \tag{A.4}$$

with H the Heaviside step function. For the PDE filter:

$$-R_\gamma^2 \nabla^2 \tilde{\gamma} + \tilde{\gamma} = \gamma. \tag{A.5}$$

R_γ, the filtering radius that is also responsible for the minimum feature size control, is typically presented in the form $\frac{N_\gamma h}{2\sqrt{3}}$. In this expression, N_γ is some number of elements and h the characteristic mesh element length scale. This thesis makes use of $N_\gamma = 8$, unless differently specified.

To ensure the inhomogeneous second-order PDE is readily solvable, the inhomogeneous term's function, defined as $f(r)$, will need to have compact support: be non-zero in only a finite region of $\mathbb{R}^2$. $\gamma(r)$ in its current form does not have compact support. To this end, one redefines the PDE problem in terms of unknown field $A(\mathbf{x})$ and finds:

$$\nabla^2 A(\mathbf{x}) - \left(\frac{1}{R_\gamma}\right)^2 A(\mathbf{x}) = \frac{1}{R_\gamma^2}\left(1 - H\left(|\mathbf{x}'| - \frac{D_\gamma}{2}\right)\right) = f(\mathbf{x}').$$

In this description, f indeed has compact support and field $A(\mathbf{x})$ may be solved for using a Green's function G for a screened Poisson equation:

$$A(\mathbf{x}) = \int_{\mathbb{R}^2} G(\mathbf{x}, \mathbf{x}') f(\mathbf{x}') \, \mathrm{d}\mathbf{x}' \,;$$
$$\text{where } G(\mathbf{x}, \mathbf{x}') = -(2\pi)^{-n/2} \left(\frac{1}{|\mathbf{x} - \mathbf{x}'| \; R_\gamma}\right)^{n/2-1} K_{n/2-1}\left(\frac{|\mathbf{x} - \mathbf{x}'|}{R_\gamma}\right),$$

where n equals 2, for the spatial dimensionality of the problem and K_ν is a modified Bessel function of the second kind. Alternatively, this problem can also be solved using the Green's function for the Helmholtz

differential operator. This leads to multiplication of R_γ with the imaginary unit and the use of a Hankel function of the first kind with purely imaginary input, which ultimately yields the same result.

Utilising the screened Poisson equation's Green's function, for the field $A(\mathbf{x})$ follows:

$$A(\mathbf{x}) = \int_{\mathbb{R}^2} -\frac{1}{2\pi} K_0\left(\frac{|\mathbf{x}-\mathbf{x}'|}{R_\gamma}\right) \frac{1}{R_\gamma^2} \left(1 - H\left(|\mathbf{x}'| - \frac{D_\gamma}{2}\right)\right) \mathrm{d}\mathbf{x}'$$
$$A(\mathbf{x}) = \int_{|\mathbf{x}'| \leqslant \frac{D_\gamma}{2}} -\frac{1}{2\pi} K_0\left(\frac{|\mathbf{x}-\mathbf{x}'|}{R_\gamma}\right) \frac{1}{R_\gamma^2} \mathrm{d}\mathbf{x}'$$

Expressing $\mathbf{x}$ and $\mathbf{x}'$ in polar coordinates and with the insight that the problem is axisymmetric and therefore A is exclusively dependent on r, meaning θ may be omitted and θ' could be interpreted as the angle of $\mathbf{x}$ to $\mathbf{0}$ to $\mathbf{x}'$ in this particular geometry:

$$A(r) = -\frac{1}{2\pi\, R_\gamma^2} \int_0^{\frac{D_\gamma}{2}} \int_0^{2\pi} K_0\left(\frac{\sqrt{r^2 + r'^2 - 2rr'\cos\theta'}}{R_\gamma}\right) r' \,\mathrm{d}\theta' \,\mathrm{d}r'$$

Finally, from the conversion to A, $\tilde{\gamma} = A + 1$, and utilising this result, one may calculate the reduction of the contour diameter at which $\hat{\gamma} = 0.5$, compared to D_γ. These numerically simulated results are depicted in figure 3.7. Note that, with a hyperbolic tangent projection as it is used throughout this thesis, $\hat{\gamma} = \tilde{\gamma}$ at 0.5.

A.3. Argument for the monotonicity of the effective thermal resistance at near-fluid projected pseudodensity values during island formation

Consider the control pseudodensity field γ to be of the form

$$\gamma(r) = 1 - (1 - \psi)\, H\left(r - \frac{D_\gamma}{2}\right), \tag{A.6}$$

where D_γ is the island diameter in the control variable field and ψ the pseudodensity in the island region, with $\psi \in [0, 1]$.

Considering that, in the reference problems, $\dot{q}_{\mathrm{gen}}(\hat{\gamma})$ is dependent solely on the local projected pseudodensity $\hat{\gamma}$, its distribution remains unchanged with uniform scaling by changing ψ. The conduction-advection governing equation for the temperature field reads:

$$\{\kappa\nabla + \nabla\kappa - \rho c_p \mathbf{u}\} \cdot \nabla T = -\dot{q}_{\mathrm{gen}}. \tag{A.7}$$

Should ψ approach 1, leading to a near-fluid island region, the evaluation of the effective thermal resistance of the design becomes non-trivial. However, suppose we may take the derivative of the governing equation with respect to ψ:

$$\left\{\frac{\mathrm{d}\kappa}{\mathrm{d}\psi}\nabla + \nabla\left\{\frac{\mathrm{d}\kappa}{\mathrm{d}\psi}\right\} - \frac{\mathrm{d}\{\rho c_p\}}{\mathrm{d}\psi}\mathbf{u} - \rho c_p \frac{\mathrm{d}\mathbf{u}}{\mathrm{d}\psi}\right\} \cdot \nabla T + \{\kappa\nabla + \nabla\kappa - \rho c_p \mathbf{u}\} \cdot \nabla\left\{\frac{\mathrm{d}T}{\mathrm{d}\psi}\right\} = -\frac{\mathrm{d}\dot{q}_{\mathrm{gen}}}{\mathrm{d}\psi} \tag{A.8}$$

Realising that the properties in this equations depend on $\hat{\gamma}$ rather than ψ, use the chain rule to find:

$$\frac{\mathrm{d}\hat{\gamma}}{\mathrm{d}\psi} = \frac{\partial\hat{\gamma}}{\partial\tilde{\gamma}} \frac{\mathrm{d}\tilde{\gamma}}{\mathrm{d}\psi}. \tag{A.9}$$

Known is the expression for the filtered field for this problem, given as

$$\tilde{\gamma}(r) = 1 - \frac{1-\psi}{2\pi\, R_\gamma^2} \int_0^{\frac{D_\gamma}{2}} \int_0^{2\pi} K_0 \left(\frac{\sqrt{r^2 + r'^2 - 2rr'\cos\theta'}}{R_\gamma} \right) r' \,\mathrm{d}\theta' \,\mathrm{d}r' \,. \tag{A.10}$$

Should one differentiate this expression with respect to ψ, the resulting derivative of the field will be found to be a *not globally-zero*, even a ψ-independent field, as

$$\frac{\mathrm{d}\tilde{\gamma}}{\mathrm{d}\psi} = \frac{1}{2\pi\, R_\gamma^2} \int_0^{\frac{D_\gamma}{2}} \int_0^{2\pi} K_0 \left(\frac{\sqrt{r^2 + r'^2 - 2rr'\cos\theta'}}{R_\gamma} \right) r' \,\mathrm{d}\theta' \,\mathrm{d}r' \tag{A.11}$$

Since a hyperbolic tangent projection is used to project the filtered field $\tilde{\gamma}$ and under $a = 0.5$, one finds, so long as $\beta \in \mathbb{R}$ and finite:

$$\left.\frac{\partial\hat{\gamma}}{\partial\tilde{\gamma}}\right|_{\tilde{\gamma}=1} = \frac{\beta}{\sinh\beta} > 0 \tag{A.12}$$

As a consequence, even if $\psi = 1$, representing a fully fluid domain, any reduction in ψ – remember that 1 is its maximum permissible value due to the bounds on the pseudodensity fields – leads to a *non-zero* change in the projected pseudodensity field $\hat{\gamma}$.

Furthermore, with a non-zero change in $\hat{\gamma}$, one may note that κ, $\mathbf{u}$ and $\dot{q}_{\mathrm{gen}}$ all change by some non-zero amount locally in Ω. Note also that ρ and c_p are constructed constant in this thesis and therefore will evaluate to zero-value derivatives with respect to $\hat{\gamma}$.

As evidence of these prepositions, the following: under infinitesimal change $\mathrm{d}\hat{\gamma}$, $\mathrm{d}\kappa$ is non-zero in a manner fully equivalent to pseudodensity projection, since it also uses a hyperbolic tangent projection, only is the parameter β to be exchanged for $\beta_\kappa \in \mathbb{R}$. For $\mathbf{u}$, a finite change in the velocity field $\mathrm{d}\mathbf{u} \neq 0$ is observed as the Darcy, or equivalently DF and DFF, penalisation scheme against flow permeation is interpolated using a RAMP function. Since RAMP interpolation functions are constructed explicitly to feature to a non-zero gradient for $\hat{\gamma} \in [0, 1]$, a non-zero change will be observed in the resulting velocity field as well under design change $\mathrm{d}\hat{\gamma}$. Also $\dot{q}_{\mathrm{gen}}$ is interpolated using a RAMP interpolation function, so the same argument applies.

Finally, it is understood that, at uniform design $\hat{\gamma} = 1$ due to $\psi = 1$, representing a fully-fluid design in Ω, the heat influx equals 0, and the temperature field will be constant at $T = T_{\mathrm{ref}} = T_{\mathrm{inlet}}$. This implies that any change to ψ away from 1 results in non-zero changes to all material properties following the chain rule. As such, even though $\nabla T = \mathbf{0}$ in all of Ω in the fully-fluid design, it is found:

$$\left\{ \frac{\mathrm{d}\kappa}{\mathrm{d}\psi}\nabla + \nabla\left\{\frac{\mathrm{d}\kappa}{\mathrm{d}\psi}\right\} - \cancelto{0}{\frac{\mathrm{d}\{\rho c_p\}}{\mathrm{d}\psi}}\mathbf{u} - \rho c_p \frac{\mathrm{d}\mathbf{u}}{\mathrm{d}\psi} \right\} \cdot \nabla T + \{\kappa\nabla + \nabla\kappa - \rho c_p \mathbf{u}\} \cdot \nabla\left\{\frac{\mathrm{d}T}{\mathrm{d}\psi}\right\} = -\frac{\mathrm{d}\dot{q}_{\mathrm{gen}}}{\mathrm{d}\psi} < 0 \tag{A.13}$$

where, due to the linearity of all differential operators in the resulting terms:

$$\frac{\mathrm{d}\dot{q}_{\mathrm{gen}}}{\mathrm{d}\psi} > 0 \implies \frac{\mathrm{d}T}{\mathrm{d}\psi} > 0 \text{ globally.} \tag{A.14}$$

Then:

$$\frac{\mathrm{d}T}{\mathrm{d}\dot{q}_{\mathrm{gen}}} > 0 \implies \mathrm{R}_{\mathrm{th,eff}} \tag{A.15}$$

B

Problem specifications

B.1. Reference problem dimensional specifications and process to obtaining discrete equivalent designs

B.1.1. 2D-NW-F geometry

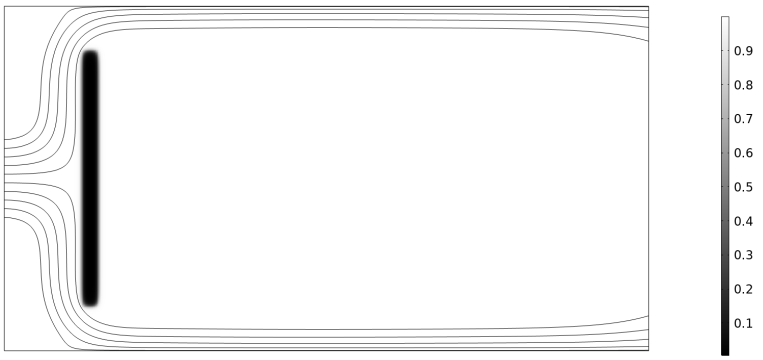

Figure B.1: *Design geometry of the 2D-NW-F reference problem. Depicted is the projected pseudodensity field and streamlines of the velocity field for illustrative purposes.*

The domain is 150 mm wide and 80 mm tall. The fluid inlet is specified on part of the left wall of the domain, is centred and has a length of 20 mm. A symmetric parabolic profile is specified on the horizontal velocity component at the inlet, with a maximum velocity of 10 mm/s. The entire right wall is a zero-pressure outlet and the upper and lower walls are slip walls.

The internal wall is defined with a rectangular region of 4 mm wide and 60 mm tall where zero-control pseudodensity is applied. The center of the wall is placed 20 mm behind the inlet face. This field is then still filtered and projected using the specified methods in this thesis. The discrete equivalent problem geometry is obtained by filtering $\Omega$by the condition $\hat{\gamma} \geqslant \alpha_{\mathrm{th}} = 0.5$.

B.1.2. 2D-Island-F/TF geometry

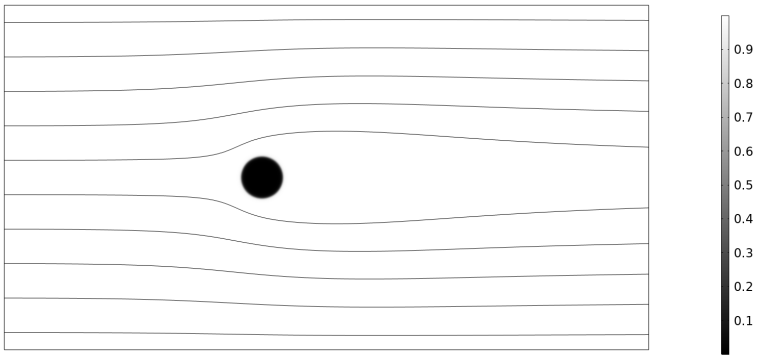

Figure B.2: *Design geometry of the 2D-Island-F and -TF reference problems. Depicted is the projected pseudodensity field and streamlines of the velocity field for illustrative purposes.*

The domain is 150 mm wide and 80 mm tall. The fluid inlet is specified on the left wall of the domain, with uniform horizontal velocity of 5 mm/s. The entire right wall is a zero-pressure outlet and the upper and lower walls are slip walls.

The island is defined with a circular region of diameter 10 mm, centred around the point at the average height of the domain and 60 mm behind the inlet face. In the circular region zero-control pseudodensity is enforced. This field is then still filtered and projected using the specified methods in this thesis. The discrete equivalent problem geometry is constructed with the same specification, but with an island diameter of 9.70 mm to adjust for the shrinkage effect induced by the filtering of the control pseudodensity field.

B.1.3. 2D-Plate-TF geometry

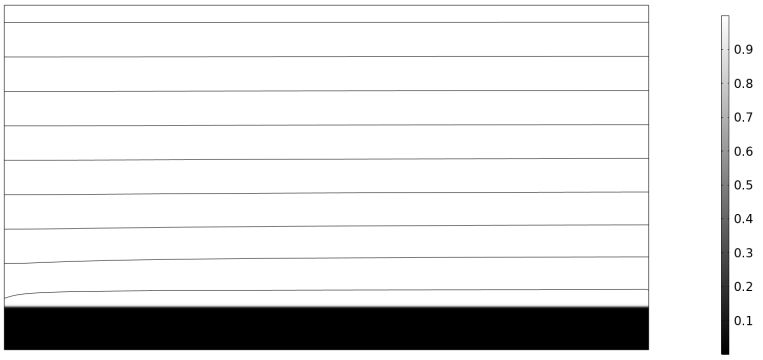

Figure B.3: *Design geometry of the 2D-Plate-TF reference problem. Depicted is the projected pseudodensity field and streamlines of the velocity field for illustrative purposes.*

The domain is 150 mm wide and 80 mm tall. The fluid inlet is specified on the left wall of the domain, as a fully developed flow with an average velocity of 5 mm/s. The entire right wall is a zero-pressure outlet and the upper and lower walls are slip walls.

The plate region is defined with a rectangular region of height 10 mm in the bottom 10 mm of the total domain. It spans the entire width of the simulation domain and in it, zero-control pseudodensity is enforced. This field is then still filtered and projected using the specified methods in this thesis. The discrete equivalent problem geometry is constructed with the same specification as it does not need to be adjusted for shrinkage effects due to filtering, considering the control variable field contours have no curvature.