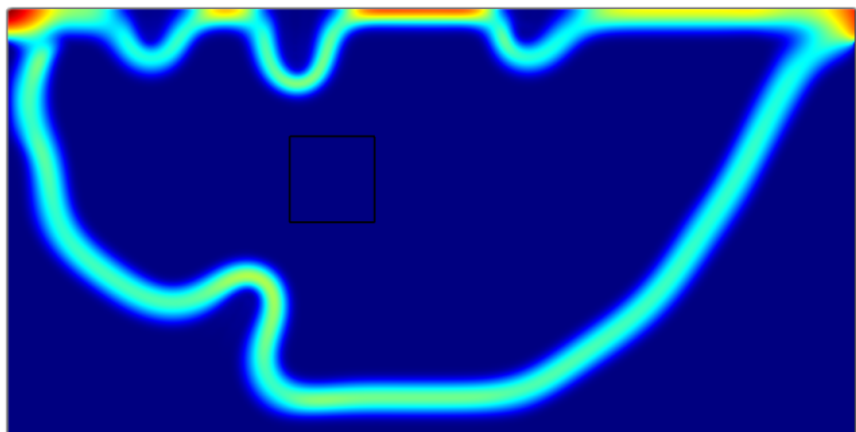
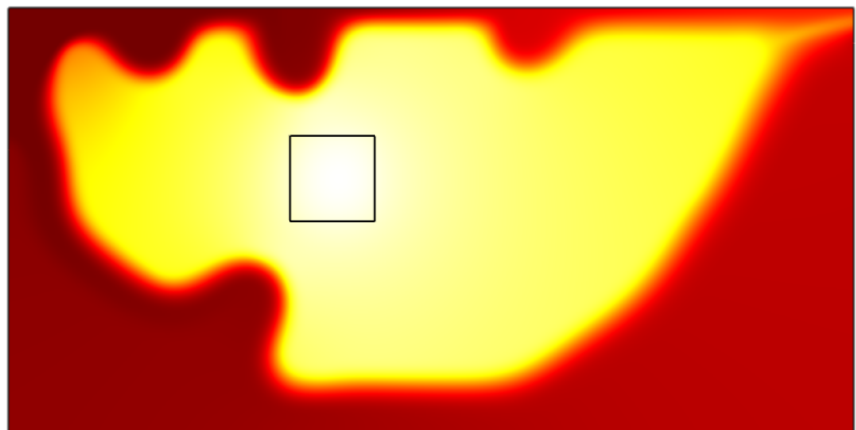


Department of Precision and Microsystems Engineering

Mitigating weaknesses of density-based thermo-fluid topology optimization using meta-optimization of modeling parameters

Thomas Bram van der Hout

Report no : 2021.020
Coach : Dr. ir. M. Langelaar
Specialisation : Mechatronic System Design
Type of report : Master Thesis
Date : March 30, 2021



Mitigating weaknesses of density-based thermo-fluid topology optimization

Using meta-optimization of modeling parameters

by

T.B. van der Hout

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Tuesday March 30, 2021 at 10:00 AM.

Student number:	4218817	
Project duration:	September 1, 2019 – March 30, 2021	
Thesis committee:	Dr. ir. M. Langelaar,	TU Delft, PME, supervisor
	Dr. Ir. L. Noël,	TU Delft, PME
	Ir. M.J.B. Theulings,	TU Delft, PME
	Dr. ir. M.J.B.M. Pourquoi,	TU Delft, P&E
	Ir. F.J.P. Hoeven,	Prodrive Technologies, supervisor
	Ir. W. van Dijk,	Prodrive Technologies, supervisor

This thesis is confidential and cannot be made public until March 30, 2022.

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

Preface

Downward is the only way forward.

Inception

This quote sums up the timeline of my graduation research fairly well. Where it started as a ‘simple’ investigation to compare laminar and turbulent thermo-fluid Topology Optimization, the path towards accurate turbulent Topology Optimization was found to be largely unpaved. I had to dig down deep in the search for definitive answers, often finding even more questions on the way down. This thesis answers some but not all of the questions found while digging, partially filling the hole of questions with solid ground. I hope that my work contributes to a strong foundation to build upon the shiny tower that is turbulent Topology Optimization.

I want to sincerely thank my supervisors dr. ir. Matthijs Langelaar, ir. Wessel van Dijk and ir. Frank Hoeven for their guidance and support during these turbulent times (pun intended). Furthermore, my thanks go to ir. Maarten Theulings for his involvement and extensive reviewing of this report and to dr. ir. Mathieu Pourquoi for his advise on turbulence modeling at the beginning of my research.

Last but certainly not least, I would like to thank Sabien, my family, friends and new colleagues for their support. You all made the digging a lot more bearable.

*T.B. van der Hout
Helmond, March 2021*

Summary

Power electronic systems are reaching higher efficiencies as their technology advances, which often results in components of smaller size with higher power densities. Cooling these components becomes increasingly challenging as high power densities require cooling with large heat fluxes. Topology Optimization (TO) of thermo-fluids can be used to find cooling interface geometries which achieve high heat transfer with realistic pumping power. However, current methods for thermo-fluid TO show several issues. Firstly, the fluid models used for thermo-fluid TO show weaknesses that cause deviating behavior compared to conventional fluid models, which can result in large under- or overestimation of heat transfer especially when applied to turbulent flow. Secondly, although this deviating behavior is known to result in inferior modeling accuracy, the achievable accuracy of thermo-fluid TO has never been quantified. Lastly, TO currently requires many modeling parameters to be specified manually. Since these parameters largely affect the accuracy of the thermo-fluid solver, tedious parameter tuning is part of the TO design process.

The goal of this thesis is to quantitatively analyze and improve the accuracy of (turbulent) density-based thermo-fluid TO. The focus lies on turbulent flows in 2D, using Reynolds-averaged Navier-Stokes (RANS)-based turbulence modeling and the Spalart-Allmaras (S-A) closure model. In this thesis, a framework is firstly presented to quantify the modeling accuracy and the errors caused by weaknesses of the density-based fluid models. This framework consists of a verified reference simulation and a density-based simulation of the same scenario: a developing internal boundary layer in a straight channel. The density-based simulation is set-up by applying the density-based thermo-fluid model to a manually specified density distribution and thus decoupling the density-based fluid model from the TO algorithm. As a measure of the modeling accuracy, the average Root Relative Squared Error (RRSE) of the velocity, temperature and eddy viscosity distributions is introduced. The analysis framework is employed to provide insight into the effect of several modeling parameters on the modeling accuracy. Next, a method is introduced to automatically tune the modeling parameters of these fluid models: meta-optimization. It uses the analysis framework to obtain the average RRSE and optimizes all modeling parameters simultaneously with a gradient-free optimization algorithm. Lastly, two methods to mitigate the errors caused by the weaknesses are discussed. The first mitigation method aims to reduce modeling errors by tuning the modeling parameters using meta-optimization and the second method includes the thermal conductivity of the solid material as a variable alongside the modeling parameters. This material property is chosen because it can alter the temperature gradient in a region of solid material which inherently has slight porosity when using density-based fluid models, with the aim to compensate for convective heat transfer caused by fluid penetrating this region. The results of the two mitigation methods are first compared to two reference cases applied to the density-based simulation of a straight channel. After that, the optimized modeling parameters are applied to a different geometry (a turbulent flow in a U-channel) to test how well their optimality is maintained. Finally, a version of the first mitigation method adapted for laminar flow is applied to a laminar density-based thermo-fluid TO, to investigate its influence on the achievable performance and accuracy.

The developed framework is shown to be functional as it provides a good quantitative match to an in-depth qualitative analysis of the errors visible in a turbulent density-based fluid model. It provides insights into the effects of the various modeling parameters on the accuracy and initiates a novel method of post-processing the resulting density variables of a TO which is based on a Darcy force threshold. The meta-optimization procedure is capable of automatically tuning the modeling parameters with relatively good efficiency compared to a manual tuning cycle, as it drastically reduces manual labour but has significant computational cost. The first mitigation method shows promising accuracy both from the turbulent U-channel test case and from the laminar thermo-fluid TO. The second mitigation method shows inferior performance compared to the first method and is therefore not considered feasible at this point.

Contents

1	Introduction	1
1.1	Background and Motivation	1
1.2	Goal	3
2	Topology Optimization of thermo-fluid systems with the density method	4
2.1	Topology Optimization methods and software packages	4
2.2	Thermo-fluid modeling theory	5
2.2.1	Modeling laminar flow	5
2.2.2	Modeling turbulent flow.	5
2.3	The density method for Topology Optimization	8
2.3.1	Density variables	9
2.3.2	Physics model	9
2.3.3	Material interpolation	10
2.3.4	Objective function	10
2.3.5	Constraint functions	11
2.3.6	Optimization algorithm	11
2.3.7	Filtering functions.	11
2.3.8	Mathematical problem formulation.	12
2.3.9	Post-processing results	13
2.4	Applying the density method to thermo-fluid Topology Optimization	13
2.4.1	Adapting Navier-Stokes equations.	13
2.4.2	Adapting Reynolds-averaged Navier-Stokes equations and Spalart-Allmaras turbulence model	14
2.4.3	Adapting energy equation	15
2.4.4	Relevant objective function and constraints.	16
2.5	Weaknesses of density method in combination with thermo-fluid Topology Optimization	17
2.6	Research questions	19
2.7	Research plan	19
3	Defining modeling accuracy of the density method	21
3.1	Setting up a reference simulation	21
3.1.1	Geometry selection.	22
3.1.2	Verification strategy	23
3.2	Boundary conditions	23
3.2.1	Inlet Reynolds number range	24
3.2.2	Heated wall condition.	25
3.3	FEM discretization	26
3.4	Material selection	26
3.5	Setting up a density-based verification simulation	26
3.6	Defining a measure of modeling accuracy	28
3.7	Meta-optimization set-up	28
4	Detailed analysis of weaknesses in the density method	30
4.1	Porous solid material	30
4.2	Blurred walls	31
4.2.1	Erroneous heat transfer	31
4.2.2	Error in velocity and thermal boundary layer	32
4.2.3	Darcy force profiles.	33
4.2.4	Interpolation function profiles	35
4.2.5	Density variable gradient.	36
5	Mitigation methods for weaknesses in the density method	38
5.1	Mitigating blurred walls by tuning modeling parameters	38
5.1.1	Turbulent optimization problem	38
5.1.2	Turbulent results	39

5.2	Mitigating porosity of fully solid material by altering thermal conductivity	41
5.2.1	Turbulent optimization problem	42
5.2.2	Turbulent results	42
6	Assessment of mitigation methods for weaknesses of the density method	45
6.1	U-channel test case set-up.	45
6.2	U-channel test case results	46
7	Thermo-fluid Topology Optimization of a Prodrive Power Module cooling interface with meta-optimized modeling parameters	48
7.1	Laminar Topology Optimization	48
7.1.1	Problem definition	48
7.1.2	Manually tuned Topology Optimization	50
7.1.3	Meta-optimized Topology Optimization	51
7.2	Discussion of turbulent Topology Optimization	53
7.2.1	Unstable optimizer with thermo-fluid Topology Optimization	53
7.2.2	Unstable thermo-fluid solver with thermo-fluid Topology Optimization	53
7.2.3	Results of continuation strategy	54
8	Conclusions and outlook	56
8.1	Conclusions.	56
8.2	Future work	57
	Acronyms	59
A	Accuracy of reference simulations	60
A.1	Verifying accuracy of developing external flow along flat plate reference simulation	60
A.1.1	Fluid mechanics correlation functions.	60
A.1.2	Thermal correlation functions.	62
A.1.3	Flat plate reference simulation results.	63
A.2	Verifying accuracy of developing internal flow through channel reference simulation.	72
A.2.1	Correlation functions for fully developed channel.	72
A.2.2	Channel reference simulation results	73
B	Blurred wall mitigation method applied to laminar flow	74
B.1	Laminar optimization problem	74
B.2	Laminar results	74
C	Analyzing effect of modeling parameters on modeling accuracy	77
D	Overview of density method variables for thermo-fluid Topology Optimization	80
E	Set-up of turbulent U-channel test case	81
E.1	Dimensions	81
E.2	Boundary conditions	82
E.3	Turbulence model	82
E.4	Material properties and interpolations	82
E.5	Mesh	83
E.6	Stabilization.	83
E.7	Measure of accuracy	84
	Bibliography	85

Introduction

1.1. Background and Motivation

Prodrive Technologies (prodrive-technologies.com) is a company that produces advanced technology solutions in six different technology programs: Embedded Computing Systems, Motion & Mechatronics, Power Conversion, Controls & Connectivity, Vision & Sensing and Industrial Automation. One of the core businesses of Prodrive Technologies (PT) is the design and manufacturing of power electronics and power electronics can be found in nearly all of their products. A cornerstone of electronic systems is the power semiconductor device. Whether it is in a voltage converter for battery systems or a signal amplifier for an MRI machine, they all feature these active switching elements. Traditionally, pre-packaged power semiconductor devices (power modules) are used which often house multiple power semiconductors in an electrical topology which is suited for its application. Such modules are delivered with a housing that features pre-defined electrical, mechanical and thermal connections. PT has taken up the challenge to design its own power modules (Fig. 1.1), ensuring full design freedom in packaging and electrical topology for a seamless integration in its products.




Figure 1.1: A custom Prodrive power module

Technological advancements in the field have gradually increased the efficiency of power semiconductors, which has led to a decreasing size of such devices relative to the power they can convert. With this increased power density comes an increased heat flux which should be transferred to a coolant to keep temperatures as low as possible. The lifetime of a power semiconductor device is greatly dependent on its operating temperature (a 10 °C decrease can lead to a doubling in lifetime depending on the scenario and packaging geometry), which leads to the wish of minimizing the operating temperature. For the last decades, Silicon has been the element of choice for power semiconductors. However, the recent development of Silicon Carbide and Gallium Nitride semiconductors have pushed heat flux values to the MW/m² range, which lies beyond the capability of single phase forced convection when directly cooling the source [30]. This indicates that the heat from the source must be spread to a larger ‘cooling interface’ surface such that the heat flux is low enough to be cooled by single phase forced convection. Spreading heat to a larger surface requires adding solid material between the heat source and the cooling interface (referred to as a heat spreader), adding thermal resistance to the system and resulting in higher heat source temperatures. Literature suggests that only a boiling fluid is relevantly capable of transferring higher fluxes than single phase forced convection [30], which is a concept that PT is also researching. PT desires to investigate methods which can minimize the heat source temperature of a high heat flux power semiconductor using single phase forced convection.

Figure 1.2 displays a typical schematic cross-section of a power module with a conventional cooling interface. The power semiconductor at the top of the power module produces a heat flow Q on an area A_s . This heat is transferred

to a coolant through a heat spreader and a cooling interface (assuming that heat transfer from the top surface of the semiconductor is negligible). The figure displays how the heat is spread from the source area A_s to a larger effective cooling area $A_{c,eff}$ at the cooling interface, where it can be transferred to the coolant by a convective heat transfer coefficient h_c . The heat spreads in a spherical pattern from the source (indicated in Fig. 1.2 by the temperature contours), which results in a non-uniform heat flux at the planar cooling interface. To minimize the heat source temperature, the thermal resistance $R_{th}[K/W]$ between the heat source and the coolant should be reduced. This thermal resistance for a Prodrive power module can be simplified to the sum of the convective resistance and all conductive thermal resistances. Both thermal resistances can be reduced individually.

Reducing the convective thermal resistance equates to maximizing the convective heat transfer coefficient h_c between the cooling interface and the coolant. This coefficient is dependent on the velocity of the coolant, the coolant's material properties and the geometry of the cooling interface [23]. Firstly, increasing the coolant velocity U for a fixed geometry and material properties increases the heat transfer coefficient, especially when the velocity is increased such that the flow becomes turbulent. Secondly, some fluids are capable of transferring heat more efficiently than others. However, in the relevant industries for PT the usage of coolants is mostly restricted to water. Lastly, the interaction between the coolant and the cooling interface greatly influences the heat transfer coefficient but can also add friction to the flow, resulting in a cooling system which requires high pumping power. Finding a geometry that provides a high heat transfer coefficient without causing too much power losses is a major challenge.

The sum of all conductive resistances is partly defined by the fixed layers of the power module and in part by the variable heat spreader material. The thermal resistance between the source and the cooling interface is proportional to the heat spreader thickness and inversely proportional to $A_{c,eff}$. The non-uniformity of the heat flux at the cooling interface, as displayed in Fig. 1.2, indicates that this cooling interface is not transferring heat as efficiently as it could. Ideally a (spherical) cooling interface is generated which has a constant heat flux and therefore maximizes $A_{c,eff}$ while keeping the heat spreader thickness constant. However, changing the geometry of the cooling interface influences the convective heat transfer coefficient h_c , which couples the optimal shape of the heat spreader to the fluid behavior. Other methods of minimizing the conductive thermal resistance such as the material selection and layer composition of the Prodrive power module are subjects of separate research projects within PT.




Figure 1.2: Illustration of the heat from a heat source spreading through a Prodrive power module and heat spreader material

Minimizing the heat source temperature by generating a cooling interface which has a high convective heat transfer coefficient and a large effective cooling area poses a challenging three-dimensional design problem. To this extent, numerous structural optimization methods are available today which range in complexity from optimizing a set of modeling parameters to Topology Optimization (TO), a technique which is able to generate optimized 3D designs in a given design space. TO discretizes the design space into a set of ‘pixels’ and determines, based on an optimization objective and constraints, whether each pixel should become solid or fluid. This pixel-based approach gives TO the highest design freedom of all optimization methods and allows it to converge closest to a global optimal 3D structure for a given problem. Various methods exist within TO, with the density method being the most established due to its age and mathematical simplicity. Density-based TO has been proven to be effective at optimizing thermo-fluid systems with the goal to maximize heat transfer, mostly for laminar [21] but recently also for turbulent scenarios, both with a ‘frozen turbulence’ assumption [25] and with a continuous adjoint formulation [16]. An implementation of turbulent TO with other methods than the density method is yet to be published [4]. From literature it is clear that the problem of thermo-fluid TO can be simplified to a two-dimensional scenario [16] to reduce computational cost, although this limits its capability to converge to a global optimum.

In a recent review paper of TO for fluid-based problems [4], Alexandersen and Andreasen recommend that future work in the field should focus on i.a. “accuracy of the geometric representation” and “precision of solution and/or optimality” instead of “merely reproducing old examples using a new method”. However, to date no works have been found which *quantitatively* analyze the accuracy of density-based fluid models. Any user-induced errors or errors caused by inherent weaknesses of the method can only be analyzed on a relative scale, which makes any useful improvement in

accuracy and/or precision difficult to substantiate. Furthermore, it is found that the geometry generated by the TO relies heavily on the values of modeling parameters, which are always tuned manually on a trial-and-error basis.

If a framework can be developed which allows the quantitative analysis of the accuracy of density-based fluid models, the errors introduced by weaknesses of the density method can be properly analyzed. Furthermore, possible mitigation of these weaknesses can be based on quantitative trade-offs and optimized for their maximum effect. Lastly, tuning of the modeling parameters can be performed in an efficient automated way. These improvements could increase the accuracy of density-based thermo-fluid TO and can reduce their computational cost by eliminating tedious parameter tuning.

1.2. Goal

The over-arching goal of this research is

To quantitatively analyze and improve the accuracy of (turbulent) density-based thermo-fluid Topology Optimization.

First, in Chapter 2 the theory of thermo-fluid TO is discussed and a set of research questions is defined based on literature research. The subsequent chapters of this thesis will discuss the research questions following a research plan defined in Section 2.7.

The scope of this research is bounded by computational cost of simulations and by a restricted project duration. Therefore, a simplification from a 3D problem to a 2D scenario is made. The extension to 3D scenarios can be made after feasibility is proven of the quantitative accuracy analysis and the mitigation methods. Secondly, a fluid flow and thermal behavior are modeled as steady-state which is not often the case, especially for turbulent flows. Finally, this thesis focuses on *turbulent* density-based thermo-fluid TO, but during its research some observations are also made for *laminar* flows. These observations will be discussed in the appendices.

2

Topology Optimization of thermo-fluid systems with the density method

In this chapter the theory behind Topology Optimization (TO) of thermo-fluid systems using the density method will be introduced. In Section 2.1, the choice for TO is justified and a software package to be used in this work is selected. To understand the concept of density-based TO of thermo-fluid systems, TO and modeling thermo-fluids will be introduced separately before combining them. Section 2.2 introduces the basics of modeling thermo-fluids and Section 2.3 covers the classical density method for TO. Section 2.4 combines the two components into the concept of TO of thermo-fluid systems with the density method. It also explains some of the issues present in the density method when applied to thermo-fluid problems. Finally, Section 2.6 formulates a set of research questions and hypotheses from the gaps which are found in literature and Section 2.7 introduces the research plan, which functions as a guide for the further chapters of this report.

2.1. Topology Optimization methods and software packages

The choice for TO over other types of optimization (size- or shape optimization) has been made in a literature study preceding this research [30]. Opposite to the other types of optimization, TO requires no initial design input from the user to generate geometries [7]. Shapes are generated by determining for a set of 'pixels' which ones are best black (solid) and which need to be white (air / fluid) in order to minimize a user-selected objective function. In other words, with TO the algorithm is completely free to create a very detailed optimized shape given a small set of user-defined parameters. This exceptional level of optimizer freedom makes TO the method of choice for this project, where ideally an optimized geometry is determined with minimal influence from the user.

Various methods are distinguished within TO, which define the interaction between a physics simulation (e.g. mechanical loading of a structure) and the algorithm that computes the topological change. The selection of the density method over other methods for TO (ground structure, level-set, etc.) is not trivial as the trade-offs lie in detailed issues, such as stability when applied to thermo-fluid simulations, achievable accuracy of the method, ease of implementation and availability of software with built-in TO algorithms. The choice for the density method is also explained in the aforementioned literature study [30]. The literature study found that most of the relevant research in the field of TO of thermo-fluid systems selected to use the density method. This indicates that the density method is applicable to such thermo-fluid optimizations. Since the goal of the research is not to invent a new optimization method, the density method was deemed the most relevant approach. An added benefit of this choice is that the density method is one of the earliest implementations of TO. The extensive history of research being done on this method has led to many FEM software packages implementing TO modules using the density method. The ability to use a built-in model to run TOs for this research greatly reduces initial set-up time and increases user friendliness, which results in a larger chance of applying the TO of thermo-fluid systems in future projects within Prodrive Technologies. A more detailed account of the strengths and weaknesses of this method will be given in Section 2.4. The use of the level-set method is emerging in the field of thermo-fluid TO because it solves some of the largest issues emerging from the density method. However, due to the young age of the method no CFD or FEM software packages have implemented it as a ready-to-use tool yet (at the start of the research). A custom implementation requires a large time investment as the method requires advanced features (e.g. adaptive mesh refinement, hole nucleation schemes) to reach its full potential, which makes it less suited for this thesis.

After choosing the optimization method, a suitable software package needs to be selected. Unfortunately, the number of software packages that allow the user to set up a multiphysics TO (such as thermo-fluid TO) using a built-in density method model is still very limited. The two main competitors that have this capability are Altair Optistruct and COMSOL

Multiphysics, of which the latter is selected for its ability to perform automatic differentiation of analytical functions as well as the author's previous experience with this software and the availability of support from TU Delft.

2.2. Thermo-fluid modeling theory

Before a TO on a thermo-fluid system can be performed, the fluid flow and heat transfer behavior must first be modeled. This is done by solving a numerical finite element Computational Fluid Dynamics (CFD) simulation. A CFD simulation approximates the continuous problem of a fluid flow by solving a set of governing equations on a discrete grid [8]. In this research the discretization is done on a Finite Element Method (FEM) grid opposed to a Finite Volume Method (FVM) grid [1] due to the integration of FEM in COMSOL.

As is discussed in Chapter 1, a significant increase in cooling performance is expected when moving to the turbulent flow regime. Therefore, both a laminar and turbulent flow situations are considered in this research. A different set of governing equations is required when modeling a laminar flow compared to a turbulent flow. The laminar governing equations are introduced in Section 2.2.1 and the turbulent governing equations are introduced in Section 2.2.2.

2.2.1. Modeling laminar flow

A laminar flow occurs when viscous friction in a fluid dampens out any small perturbations which may be caused by imperfections of a wall or other irregularities. Because of this, a laminar flow has a predictable behaviour. It can therefore be modeled exactly, as long as the grid is small enough to resolve the smallest details of the flow. The governing equations for this exact solution of an incompressible fluid are the Navier-Stokes (N-S) equations, which consist of an equation for the conservation of mass (Eq. (2.1)) and an equation for the conservation of momentum (Eq. (2.2)). To model heat transfer in a fluid and between solid and fluid domains, a conjugate heat transfer model is used where it is assumed that the fluid density is constant (independent of temperature) and viscous heating is neglected. This adds a third governing equation for the conservation of energy (Eq. (2.3)). The resulting set of governing equations, following the differential notation defined by the COMSOL CFD module documentation [11] is stated in Eqs. (2.1) to (2.3).

$$\rho \nabla \cdot \vec{u} = 0 \quad (2.1)$$

$$\underbrace{\rho \frac{\partial \vec{u}}{\partial t}}_{\text{Local acceleration}} + \underbrace{\rho(\vec{u} \cdot \nabla)\vec{u}}_{\text{Convective acceleration}} = \underbrace{-\nabla P}_{\text{Pressure gradient}} + \underbrace{\nabla \cdot \mu(\nabla \vec{u} + (\nabla \vec{u})^T)}_{\text{Diffusion}} + \underbrace{\vec{F}}_{\text{Body force}} \quad (2.2)$$

$$\underbrace{\rho C_p \frac{\partial T}{\partial t}}_{\text{Energy change with time}} + \underbrace{\rho C_p \vec{u} \cdot \nabla T}_{\text{Convection}} + \underbrace{\nabla \cdot \vec{q}}_{\text{Diffusion}} = \underbrace{Q}_{\text{Source}} \quad (2.3)$$

where \vec{u} is the velocity vector, t is time, ρ is the material density, P is the pressure, μ is the dynamic viscosity, \vec{F} is an external force, C_p is the thermal capacity at constant pressure, T is the temperature, \vec{q} is the conductive heat flux (as defined by Fourier's law) and Q is the volumetric heat source term.

The energy equation applies both to the fluid and solid domains of a CFD simulation. In the fluid domain the heat transfer is governed by a combination of conductive- and convective heat transfer, of which the latter is directly proportional to the fluid's velocity field. At the walls, the fluid velocity is forced to be 0 (no-slip boundaries), causing a boundary layer to develop (for the fundamentals of fluid flow and boundary layers, refer to [33]). The shape and thickness of this boundary layer defines the heat transfer between the solid and fluid domains. In solid domains the convective term of the energy equation is disabled (the fluid velocity is 0), simplifying the equation to Fourier's law.

In all simulations that are performed for this research, the assumption is made that the flow shows a steady-state behavior (this is true for laminar flow if the boundary conditions are time-independent), which removes all time-dependent terms from the governing equations.

2.2.2. Modeling turbulent flow

In contrast to a laminar flow, a turbulent flow shows chaotic behavior. The fluid carries enough kinetic energy that a small perturbation can cause it to become locally unstable, creating a turbulent eddy. These eddies interact with each other at various length and timescales. They break up into smaller eddies and swirl around each other. The N-S governing equations for turbulent flow are exactly the same as for a laminar flow. However because of its chaotic behavior at multiple scales, solving the governing equations directly for a turbulent flow, a concept called Direct Numerical Simulation (DNS), requires a grid resolution which is high enough to compute the smallest eddies in the flow. For this finest resolution the Kolmogorov microscale [19] is often used, which can be several million times smaller than the largest eddies. Furthermore, the temporal scale of the eddies needs to be resolved which makes DNS only possible in time-dependent simulations with small time-steps. If any of these scales is not properly resolved or if too much numerical diffusion is added to the simulation to stabilize it [2], the turbulence might not be visible although it is present in real-life scenario's.

Performing DNS simulations for anything but very simple geometries is too computationally expensive using current technology.

This Section will first introduce the Reynolds-averaged Navier-Stokes (RANS) equations: an adaptation of the N-S equation to allow for the approximation of turbulent effects. Secondly, the concept of turbulent closure models is introduced which perform this approximation of the turbulence. Thirdly, the selected Spalart-Allmaras (S-A) turbulence closure model is explained. Finally, the modeling of turbulent heat transfer is discussed.

Reynolds-averaged Navier-Stokes equations. A less computationally expensive approach to solving turbulent flow is to use the RANS equations [26], where the original velocity \vec{u} is split up into a time averaged velocity \vec{U} and a fluctuating velocity \vec{u}' component (Fig. 2.1 displays the relation between these three velocity components). The original pressure p is also split up into a time averaged pressure P and a fluctuating pressure p' . The RANS equations are defined again in differential form by the COMSOL CFD module documentation [11] in Eq. (2.4) (Conservation of mass) and Eq. (2.5) (Conservation of momentum).

$$\rho \nabla \cdot \vec{U} = 0 \quad (2.4)$$

$$\rho \frac{\partial \vec{U}}{\partial t} + \rho \vec{U} \cdot \nabla \vec{U} + \underbrace{\nabla(\rho \vec{u}' \times \vec{u}')}_{\text{Reynolds stress tensor}} = -\nabla P + \nabla \cdot \mu(\nabla \vec{U} + (\nabla \vec{U})^T) + \vec{F} \quad (2.5)$$

where the velocity and pressure terms in the original N-S equations are replaced with the time-averaged components. The last term on the left-hand side of Eq. (2.5) is added and represents the interaction between the fluctuating velocity and the mean velocity. This term is called the Reynolds stress tensor. In conclusion, using the RANS equations the mean velocity of a turbulent flow can be solved as if it were a laminar flow, but the detailed information of the fluctuating velocity is still needed, i.e. the RANS equations need to be ‘closed’ with the Reynolds stress tensor.




Figure 2.1: Definition of mean and fluctuating turbulent variables: (a) velocity; (b) pressure, source: [33]

Turbulence closure models. By approximating the Reynolds stress tensor without actually computing it, the computational burden can be greatly reduced with respect to DNS simulations. The models that approximate the Reynolds stress tensor are so-called ‘turbulence closure models’, often simply referred to as ‘turbulence models’ of which many are available [27].

The earliest method of closing the RANS equations is the Boussinesq approximation [36], in which the Reynolds stress is approximated using the behavior of the mean flow and a ‘turbulent eddy viscosity’ μ_T . Following the notation of COMSOL [11], the deviating part of the Reynolds stress is defined in Eq. (2.6).

$$\underbrace{\rho(\vec{u}' \times \vec{u}')}_{\text{Reynolds stress}} - \frac{\rho}{3} \text{trace}(\vec{u}' \times \vec{u}') \mathbf{I} = \underbrace{-\mu_T(\nabla \vec{U} + (\nabla \vec{U})^T)}_{\text{Boussinesq approximation}}. \quad (2.6)$$

The spherical part of the Reynolds stress (second term in Eq. (2.6)) can be written as

$$\frac{\rho}{3} \text{trace}(\vec{u}' \times \vec{u}') \mathbf{I} = \frac{2}{3} \underbrace{\rho k}_{=0}$$

where k is the turbulent kinetic energy and \mathbf{I} is the unity matrix. In this research only incompressible flow is simulated, in which case the above term is included in the pressure field and can therefore be neglected. The result is the Boussinesq

approximation for incompressible flow as is visible on the right-hand side of Eq. (2.6). The turbulent viscosity used in this model can either be assumed constant throughout a flow or it can be modeled. The turbulence closure models that model the turbulent viscosity can be categorized by the number of Partial Differential Equation (PDE) that are solved alongside the RANS equations.

Generally, the selection of a turbulence model is based on the applicability of the turbulence model for relevant geometries, its robustness and its accuracy. Additionally, for this research a turbulence model is preferred which has previously been adapted to the density method for TO, and the number of parameters which are added to the equations for its adaptation should be minimized (since these parameters need to be tuned for accuracy).

Spalart-Allmaras turbulence model. Based on literature research [30] of previous work in the field of TO of turbulent flows, the S-A [29] model is selected. The S-A model has been adapted for the density method of TO first by Yoon [34] and later by Dilgen, Dilgen et al. [15]. The model is a one-equation ‘low Reynolds number’ model, meaning that it solves one PDE which is accurate enough to solve the flow field throughout the whole complex turbulent boundary layer. Non-‘low-Reynolds number’ turbulence models resort to ‘wall functions’ which interpolate the boundary layer flow from an analytical velocity profile defined by the ‘law of the wall’ [31]. These models therefore switch between two solutions locally (numerical solution from CFD and analytical solution near walls). Switching to an analytical solution near the wall results in lower resolution requirements compared to low-Reynolds number models, as the latter require a resolution near walls that is fine enough to resolve the strong velocity gradient which is present in the turbulent boundary layer.

The S-A model was initially developed for aerodynamic applications, but after a modification [5] the model became applicable to more situations. The model has moderate resolution requirements and is relatively robust [3]. “Experience shows that this model does not accurately compute fields that exhibit shear flow, separated flow, or decaying turbulence.” The lack of wall functions in this model is beneficial for its adaptation to the density method, as wall functions introduce issues with the sensitivity analysis of the turbulence variables. This latest issue will be explained in more detail in Section 2.4. The PDE solved with the S-A model is stated in Eq. (2.7) following the definition of an adapted version of the model [27] in the notation defined by COMSOL [11].

$$\frac{\partial \tilde{\nu}}{\partial t} + \vec{U} \cdot \nabla \tilde{\nu} = c_{b1} \tilde{S} \tilde{\nu} - c_{w1} f_w \left(\frac{\tilde{\nu}}{l_w} \right)^2 + \frac{1}{\sigma_{\tilde{\nu}}} \nabla \cdot ((\nu + \tilde{\nu}) \nabla \tilde{\nu}) + \frac{c_{b2}}{\sigma_{\tilde{\nu}}} \nabla \tilde{\nu} \cdot \nabla \tilde{\nu} \quad (2.7)$$

where the auxiliary variables are defined as

$$\begin{aligned} c_{w1} &= \frac{c_{b1}}{\kappa_{\tilde{\nu}}^2} + \frac{1 + c_{b2}}{\sigma_{\tilde{\nu}}}, & \chi &= \frac{\tilde{\nu}}{\nu}, & f_{v1} &= \frac{\chi^3}{\chi^3 + c_{v1}^3} \\ f_{v2} &= 1 - \frac{\chi}{1 + \chi f_{v1}}, & f_w &= g \left(\frac{1 + c_{w3}^6}{g^6 + c_{w3}^6} \right)^{1/6}, & g &= r + c_{w2}(r^6 - r) \\ r &= \min \left(\frac{\tilde{\nu}}{\tilde{S} \kappa_{\tilde{\nu}}^2 l_w^2}, 10 \right), & \tilde{S} &= \max \left(\Omega + C_{Rot} \min(0, S - \Omega) + \frac{\tilde{\nu}}{\kappa_{\tilde{\nu}}^2 l_w^2} f_{v2}, 0.3 \Omega \right) \\ S &= \sqrt{2 S_{ij} S_{ij}}, & \Omega &= \sqrt{2 \Omega_{ij} \Omega_{ij}} \\ S_{ij} &= 0.5(\nabla \vec{u} + \nabla \vec{u}^T), & \Omega_{ij} &= 0.5(\nabla \vec{u} - \nabla \vec{u}^T) \end{aligned}$$

where S_{ij} and Ω_{ij} are the mean strain rate and mean rotation rate tensors, $\nu = \mu/\rho$ is the kinematic viscosity and l_w is the distance to the closest wall which is computed by

$$l_w = \frac{1}{G} - \frac{l_{ref}}{2}.$$

To compute the reciprocal wall distance G , a modified Eikonal PDE equation with non-homogeneous boundary condition $G = G_0 = 2/l_{ref}$ on walls is solved along with the S-A PDE (Eq. (2.7)). This PDE is stated in Eq. (2.8).

$$\nabla G \cdot \nabla G + \sigma_w G (\nabla \cdot \nabla G) = (1 + 2\sigma_w) G^4. \quad (2.8)$$

Finally, the S-A equation computes the undamped turbulent kinematic viscosity $\tilde{\nu}$ where the RANS equations require the input of turbulent eddy viscosity μ_T , so the two are related by Eq. (2.9).

$$\mu_T = \rho \tilde{\nu} f_{v1}. \quad (2.9)$$

The default values of the coefficients as defined by COMSOL [11] are used which are:

$$\begin{aligned} c_{b1} &= 0.1355, & c_{b2} &= 0.622, & c_{v1} &= 7.1, & \sigma_{\tilde{\nu}} &= 2/3, \\ c_{w2} &= 0.3, & c_{w3} &= 2, & \kappa_{\tilde{\nu}} &= 0.41, & C_{Rot} &= 2.0, \\ & & \sigma_w &= 0.2, & & & & \end{aligned}$$

and l_{ref} is automatically computed as one tenth of the shortest length of the bounding box of the geometry.

In conclusion, with the S-A turbulence model the undamped turbulent kinematic viscosity $\tilde{\nu}$ can be approximated using one turbulence equation as displayed in Eq. (2.7). It can be converted into the turbulent eddy viscosity μ_T using Eq. (2.9) and inserted into the RANS equations (Eqs. (2.4) and (2.5)) to solve the mean turbulent flow without having to numerically compute the detailed turbulent eddies. A second PDE is solved to compute the distance to the closest wall (Eq. (2.8)) as that parameter is required for the S-A turbulence model.

Turbulent heat transfer. The aforementioned turbulence models only approximate the fluid-mechanical effects of the turbulence and therefore apply to isothermal simulations. In the case of this research, an additional model is required that accounts for the turbulent effects on the heat transfer behavior of a fluid. Due to mixing of such a turbulent fluid, its heat transfer is increased compared to a laminar flow. However, because most of this mixing occurs at sub-grid level when using the RANS equations, the conjugate heat transfer model (Eq. (2.3)) is incapable of accounting for this increased heat transfer. To solve this, the thermal conductivity of the fluid is compensated by adding a turbulent thermal conductivity term to Fourier's law. The resulting equation is displayed in Eq. (2.10) following the notation defined by COMSOL [11].

$$\vec{q} = -(k + k_T)\nabla T \quad (2.10)$$

where the turbulent thermal conductivity k_T is defined as

$$k_T = \frac{\mu_T C_p}{Pr_T}$$

where Pr_T is the turbulent Prandtl number. This turbulent Prandtl number is approximated with the Kays-Crawford model according to

$$Pr_T = \left(\frac{1}{2Pr_{T\infty}} + \frac{0.3}{\sqrt{Pr_{T\infty}}} \frac{C_p \mu_T}{\lambda} - \left(0.3 \frac{C_p \mu_T}{\lambda} \right)^2 (1 - e^{-\lambda/(0.3C_p \mu_T \sqrt{Pr_{T\infty}})}) \right)^{-1}$$

where $Pr_{T\infty} = 0.85$ is the Prandtl number at infinity and λ is the thermal conductivity of the fluid [11].

2.3. The density method for Topology Optimization

With the fluid models for laminar and turbulent thermo-fluids introduced, the focus shifts to the second component of TO of thermo-fluid systems: the density method for TO. The density method was originally proposed by Bendsoe and Kikuchi [6]. A TO of heat conduction (from hereon referred to as 'thermal TO') with the density method will be used as an introductory case in the following subsections, because it is the most simple form of TO with the density method when applied to a thermal problem (conduction of heat from a source to a sink).

This section discusses all the standard elements of the density method following the flowchart which is displayed in Fig. 2.2, starting with the density variables in Section 2.3.1 and continuing with the physics model formulation in Section 2.3.2, the interpolation functions in Section 2.3.3, the objective function in Section 2.3.4, the constraint functions in Section 2.3.5, the optimization algorithm in Section 2.3.6, and finally the filtering functions in Section 2.3.7. A full mathematical problem formulation is given in Section 2.3.8. After the optimization is converged, the density variables must be post-processed which is discussed in Section 2.3.9. More advanced concepts that do not have a direct application to the introductory case but which will be implemented in this thesis will be stated in *italic*.

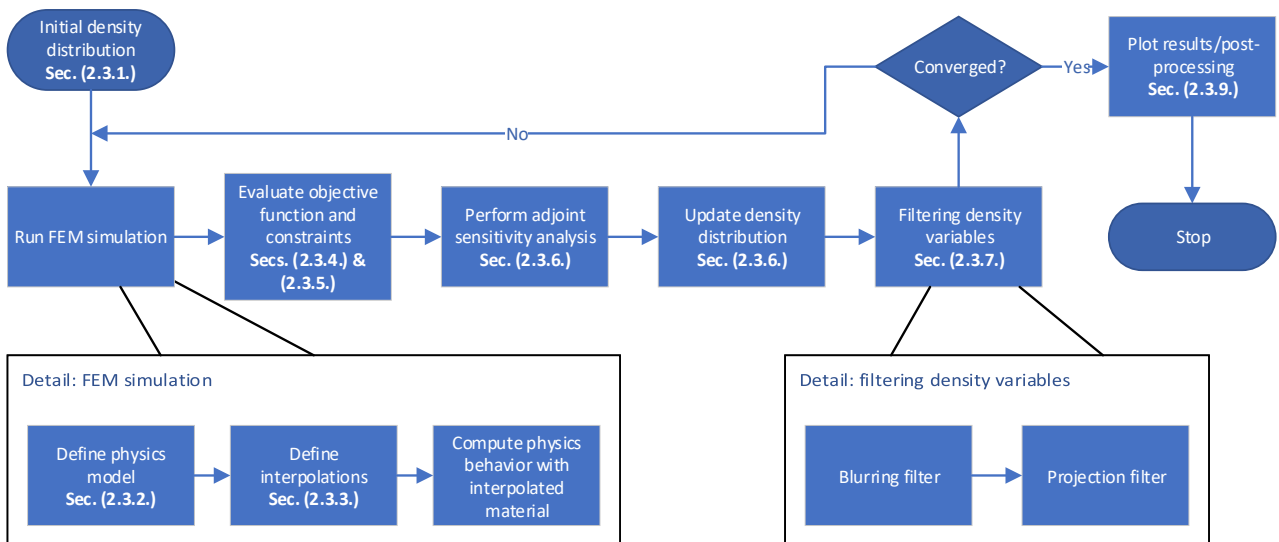


Figure 2.2: Flowchart of computations following the density method, as adapted for this research

2.3.1. Density variables

The density method defines a set of virtual density variables γ which are discretized through a finite-element mesh as is displayed in Fig. 2.3a, such that each mesh element or each mesh node is assigned one continuous variable ranging from 0 (no material, void) to 1 (solid material). In case the variables are specified on mesh nodes, the density variables are interpolated within each mesh element by a linear interpolation function such that no discontinuities are found between elements. Any type of mesh can be used (triangular elements or quad elements, structured or unstructured), and local mesh refinement results in a locally higher resolution of density variables. Figure 2.3a uses a uniform unstructured triangular mesh.


2.3.2. Physics model

A FEM simulation must be set-up with a physics model (such as the ones explained in Section 2.2) which is adapted for the density method. One set of governing equations must be implemented in such a way that they apply to the whole design domain. Solid, void, and intermediate density regions must be represented by this set, as the density method does not recognize a clear boundary between solid and void. In the case of thermal TO Fourier's law is altered by interpolating the thermal conductivity of the solid material, resulting in Eq. (2.11).


$$\nabla \cdot (-k_{int}(\gamma)\nabla T) = Q \quad (2.11)$$

where $k_{int}(\gamma)$ is the interpolated thermal conductivity as a function of the density variables. For this interpolation, an interpolation function is used as will be explained in Section 2.3.3.

The simulation needs boundary conditions such that a unique physical solution exists. For a thermal TO, a small part of the domain is set to a defined temperature and acts as a heat sink, while the rest of the domain is subjected to a constant transverse heat flux (displayed in Fig. 2.3b). This represents an electronic component which has a constant heat loss, on top of which the optimizer shall create conductive paths to connect the heat sink to this hot plate. This way, the simulation always converges because all elements have finite thermal conductivity and therefore must have a finite temperature to conduct the locally generated heat. The outer borders are defined as adiabatic, such that no heat enters or leaves the domain through the sides.



(a) Mesh used for example problem. Each mesh node is associated with a density variable



(b) Boundary conditions of example problem: in the blue region a heat sink with defined temperature is defined. In the gray region a heat source with defined heat flux is specified

Figure 2.3: Two parts of a thermal TO problem definition

2.3.3. Material interpolation

The properties of the void and solid material are interpolated based on the density variables using one or multiple interpolation functions. The thermal TO uses a so-called Solid Isotropic Material with Penalization (SIMP) interpolation (Eq. (2.12) and displayed in Fig. 2.4a) to interpolate the thermal conductivity from a low value to that of the conducting solid, ensuring that little to no heat is transferred between void elements. Setting the lower limit of thermal conduction to a finite value instead of 0, numerical instabilities are avoided. The SIMP function penalizes intermediate densities by giving them lower performance relative to their effective volume compared to the extremes of the density (0 or 1). Combined with a volume constraint (Section 2.3.5) this drives the optimization towards clear 0/1 solutions.

Other interpolation functions can be used to either penalize intermediate density values or make them more advantageous. A relevant interpolation for this research is Darcy interpolation (Eq. (2.13) and displayed in Fig. 2.4b). The Darcy interpolation works opposite compared to SIMP, by penalizing extreme densities (0 or 1). Penalizing extreme density values can help the optimizer to generate initial geometries by generating larger regions of intermediate density which enhance sensitivities (the concept of sensitivities will be explained in Section 2.3.6).

$$k_{int}(\gamma) = k_{min} + \gamma^p (k_{max} - k_{min}) \quad (2.12)$$

$$k_{int}(\gamma) = k_{min} + (k_{max} - k_{min})q \frac{1 - \gamma}{q + \gamma} \quad (2.13)$$

where k_{min} and k_{max} are the lower and upper limits of thermal conductivity, respectively; p is the SIMP interpolation penalty parameter and q is the Darcy interpolation penalty parameter.

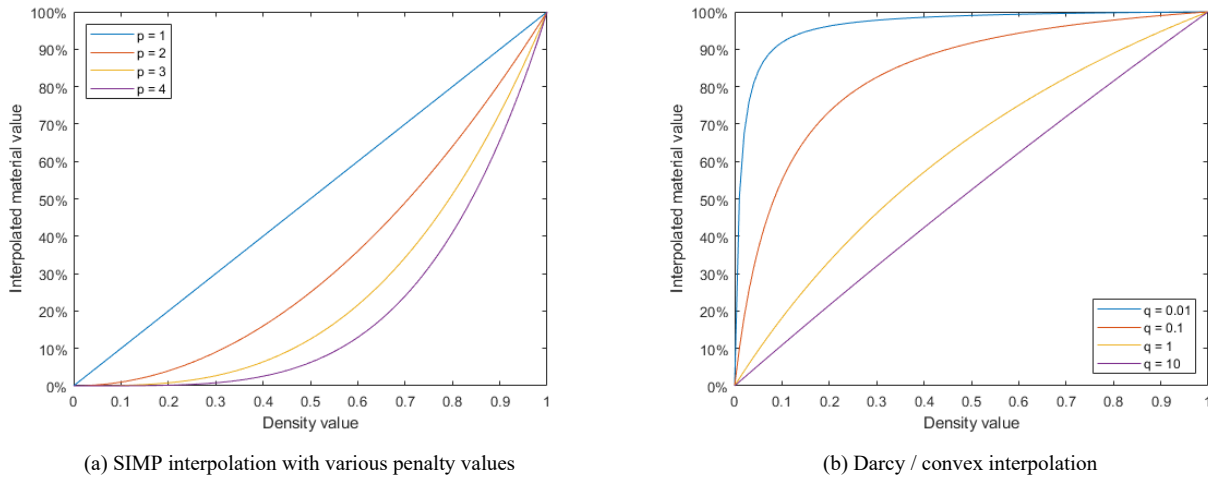


Figure 2.4: Two different interpolation functions

Borrvall and Petersson have proven that for Stokes flow a linear interpolation of the inverse permeability of solid material, α , is required to ensure existence of discrete-valued solutions [9]. However, they note that with linear interpolation their optimization often converged to local optima instead of the global optimum. They proposed a convex interpolation (here referred to as Darcy interpolation) to reduce the penalization of the inverse permeability and “to control the level of ‘gray’ in the optimal design” by altering the interpolation penalty q . Here, it is inferred that the level of gray material is inversely proportional to q . However, in two papers by Dilgen, Dilgen et al. [15],[16] a relation is discussed between q and the sharpness of the solid-fluid transition in the resulting geometry of a turbulent TO. They state that a lower q results in sharper walls (which infers less gray material is present) and opt to initially use extremely low values of q . This conflicts with the behavior described by Borrvall and Petersson. Dilgen, Dilgen et al. introduce a continuation strategy which gradually increases q towards a linear interpolation during the optimization, however it is unclear for which reason a higher penalization is desired if the low values of q already provide sharp results.

2.3.4. Objective function

The optimization algorithm is provided with an objective function which needs to be minimized in order to improve the performance of the optimized geometry. This objective function can be a summation of multiple functions describing different aspects of performance. Often, these functions are defined as integrals of some form of energy throughout the whole design domain (e.g. thermal energy, friction- and viscous losses in a fluid, strain energy in mechanically loaded structures). Other global quantities can also be used as well as local functions (e.g. temperature at specified region or point).

The objective function of the thermal TO can be selected in multiple ways that achieve a similar goal, depending on the type of heat source [22]. The most common objective function for thermal TO is to minimize the thermal compliance

[17], [7] as is defined in Eq. (2.14).

$$C = \int_{\Omega} \vec{q} T d\Omega \quad (2.14)$$

which is a product of the incoming heat flux q in each element (which is constant in this case) and its temperature T . For the example, minimizing C is thus equivalent with minimizing the average temperature in the domain. Others have also adopted different objectives such as minimizing temperature gradients [13].

2.3.5. Constraint functions

Without any constraints (aside from the upper and lower limit of the density variables), many objective functions would force the optimizer to generate domains which are completely solid or completely void (for thermal TO the best thermal conduction takes place when the whole domain is fully conductive). However, in these cases a large fraction of the domain does not significantly contribute to improving the objective function. Furthermore, in industrial applications the optimization often lies in the trade-off between performance of design and its cost or weight. Because of this, TO is most often set-up as a ‘performance-to-weight’ optimisation, to find an optimised geometry that maximizes performance for a given weight (or vice-versa). To achieve this a volume constraint is imposed in Eq. (2.15):

$$\frac{\int_{\Omega} \gamma(\vec{x}) d\Omega}{\int_{\Omega} 1 d\omega} - V \leq 0 \quad (2.15)$$

where V is the allowed volume fraction of solid volume with respect to the total volume.

Other types of constraints are also possible to ensure that the solution meets specified requirements. The thermal TO can, for instance, include a maximum temperature constraint when minimizing the thermal compliance to ensure a certain temperature in a critical region or point. Such a temperature constraint can be formulated as is displayed in Eq. (2.16).

$$\max(T) - T_{max} \leq 0 \quad (2.16)$$

where T_{max} is the maximum allowed temperature. However, the maximum operator is not differentiable which will cause issues with the sensitivity analysis. A ‘smooth maximum’ can be used to solve this, by computing the p-norm of the temperature with sufficiently high p-value (this concept is introduced by Bendsøe [7] as a global stress constraint for compliance minimization).

2.3.6. Optimization algorithm

An optimization algorithm or ‘optimizer’ is used to iteratively compute an optimized distribution of the density variables which minimizes an objective function while meeting a set of constraints. Classically a heuristic optimizer is used with the density method, which determines independently for each individual element in every iteration how much its density needs to increase or decrease based on a set of ‘optimality criteria’ [7].

For more complex (i.e. non-linear) problems, a mathematical gradient-based optimizer can be used, which determines the iterative update of all density variables simultaneously based on the local gradient of the density variables with respect to the objective function and constraints. For this, a sensitivity analysis is performed for each element to determine its local gradient. This can be done simply by a ‘finite difference analysis’, slightly perturbing the density of one element and re-computing the FEM solution to determine the new objective value. While simple, this process is extremely computationally intensive as a new FEM solution needs to be computed for each element sensitivity to even obtain one update of the density variables. A more efficient but more complex method is to approximate the gradient using a ‘discrete adjoint sensitivity analysis’ which uses the analytic derivative of the discretized governing equations and objective function to determine the gradient of the density variables with respect to the objective function and constraints. In many cases, these analytic derivatives are computed manually, which involves differentiation of extensive equations with many ‘chain rule’ steps. Fortunately COMSOL has the capability to perform these derivations automatically, of which the implementation is discussed by Dilgen, Dilgen et al. [15].

2.3.7. Filtering functions

The inherent freedom of TO can result in geometries which are physically impossible to manufacture. Two of the most common causes for these infeasible results are areas of intermediate density and checkerboard structures. Regions of intermediate density represent a porous material with altered properties, in case of thermal TO the material has reduced thermal conductivity. Checkerboard structures can occur when the optimizer mimics a region of intermediate density while being forced to generate discrete density values. Neighboring elements each receive alternating density values, hence the name checkerboard.

Currently, materials cannot be manufactured with controlled spatially varying material properties, and therefore the regions of intermediate density or checkerboard structures must be minimized. The material interpolation with penalization already aims to achieve this but in some physics it fails, because regions of intermediate density are too advantageous.

Blurring filter. Checkerboarding can be avoided by blurring the generated geometry such that ‘single-cell’ structures are transformed into regions of intermediate density. Blurring can be done by averaging a cell’s density value with values of neighboring cells. COMSOL performs this blurring with a Laplace equation [20] as is given in Eq. (2.17).

$$\boldsymbol{\gamma}_f = R_{min}^2 \nabla^2 \boldsymbol{\gamma}_f + \boldsymbol{\gamma}_c \quad (2.17)$$

where $\boldsymbol{\gamma}_f$ is the blurring-filtered density field, R_{min} is the filter radius and $\boldsymbol{\gamma}_c$ is the unfiltered density field which is obtained by the optimizer. By setting the filter radius larger than the mesh element size, checkerboarding is eliminated by averaging an element’s density value with its directly neighboring elements. By setting a larger filter radius more elements are averaged out and the minimum length scale of details, including the length of density gradient from 0 to 1, can be controlled independently from the mesh size. The filter radius R_{min} of the Laplace equation can be converted into a equivalent support domain radius R_{reg} of a ‘regular’ blurring filter with a linear decaying weighting function, by $R_{reg} = 2\sqrt{3}R_{min}$. A regular blurring filter will blur a step in the unfiltered density variables over a length of $2R_{reg}$ and therefore a parameter L_{trans} is introduced to control the 0-1 transition length. L_{trans} is related to R_{min} by Eq. (2.18):

$$R_{min} = \frac{L_{trans}}{4\sqrt{3}}. \quad (2.18)$$

However, due to the exponential weighting function of the Laplace filter the filtered density variables will approach 0 and 1 exponentially and the transition length will likely not hold exactly. It will therefore serve as an inexact measure to control the transition length.

Projection filter. ‘Projection filters’ can be utilized that project the intermediate density values towards 0 or 1 by, for example, using hyperbolic tangent functions [32], as is displayed in Eq. (2.19).

$$\boldsymbol{\gamma} = \frac{\tanh(\beta(\boldsymbol{\gamma}_f - \gamma_\beta)) + \tanh(\beta\gamma_\beta)}{\tanh(\beta(1 - \gamma_\beta)) + \tanh(\beta\gamma_\beta)} \quad (2.19)$$

where $\boldsymbol{\gamma}$ is the projection-filtered density field which is the input for the FEM simulation, $\boldsymbol{\gamma}_f$ is the blurring-filtered density field, β is the projection slope and γ_β is the projection threshold. The steeper the projection slope, the more the intermediate values are shifted towards 0 or 1. The projection threshold is typically set to 0.5, such that a balanced projection occurs.

The example problem of thermal TO uses both the blurring filter and projection filter to force 0/1 designs.

2.3.8. Mathematical problem formulation

For the example thermal TO problem, the following mathematical problem formulation is composed by combining the equations introduced in previous sections:

$$\min_{\boldsymbol{\gamma}_c} C = \int_{\Omega} \tilde{q} T d\Omega \quad (2.20)$$

$$\text{s.t.} \quad \nabla \cdot (-k_{int}(\boldsymbol{\gamma}) \nabla T) = Q, \quad (2.21)$$

$$k_{int}(\boldsymbol{\gamma}) = k_{min} + \boldsymbol{\gamma}^p (k_{max} - k_{min}), \quad (2.22)$$

$$\boldsymbol{\gamma} = \frac{\tanh(\beta(\boldsymbol{\gamma}_f - \gamma_\beta)) + \tanh(\beta\gamma_\beta)}{\tanh(\beta(1 - \gamma_\beta)) + \tanh(\beta\gamma_\beta)}, \quad (2.23)$$

$$\boldsymbol{\gamma}_f = R_{min}^2 \nabla^2 \boldsymbol{\gamma}_f + \boldsymbol{\gamma}_c, \quad (2.24)$$


$$\frac{\int_{\Omega} \boldsymbol{\gamma}(\vec{x}) d\Omega}{\int_{\Omega} 1 d\Omega} - V \leq 0 \quad (2.25)$$

where Eq. (2.20) is the thermal compliance objective function of the thermal TO, Eq. (2.21) is Fourier’s law with interpolated thermal conductivity, Eq. (2.22) is the SIMP interpolation of the thermal conductivity, Eq. (2.23) is the hyperbolic projection filter of the blurred density variables, Eq. (2.24) is the blurring filter of the control density variables and Eq. (2.25) is the volume constraint on the output (filtered) density variables. To perform this thermal TO, the modeling parameters $p, \beta, \gamma_\beta, R_{min}, V$, the material properties k_{min} and k_{max} and boundary conditions for Q and T are to be selected by the user.

2.3.9. Post-processing results

When the optimization is converged, the resulting density distribution can be plotted, which is displayed in Fig. 2.5a for the thermal TO. Because a blurring filter is usually implemented, the resulting density distribution will contain some areas of intermediate density. Before the design can be used for further analysis or manufacturing, it needs to be post-processed.

Post-processing is done by including all density values above a certain threshold in the solid domain(s) and drawing a smoothed boundary around the solid. The resulting geometry is displayed in Fig. 2.5b. This boundary can later be used to generate a body-fitted mesh or to generate a 3D geometry for manufacturing. The position of the boundary is dependent on the selection of the threshold but when the regions of intermediate density are small this effect is minimal.



(a) Resulting density distribution after optimization is converged

(b) Resulting geometry after post-processing, threshold $\gamma > 0.5$

Figure 2.5: Results of thermal TO before and after post-processing.

2.4. Applying the density method to thermo-fluid Topology Optimization

Now that the two pillars of this research have been introduced separately, they can be combined to form the theory of the density-based thermo-fluid TO. The density variables in the thermo-fluid TO are defined such that a density of 0 represents solid material and 1 represents fluid. Section 2.4.1 explains how the N-S governing equations for a fluid flow are adapted to include a penalization based on the density variables. This ensures that the fluid flow is slowed down in regions of low density (solid material).

Section 2.4.2 shows how the RANS equations are adapted to include the Brinkman penalization. This Section also explains how the S-A turbulence equation is adapted to include a penalization on the turbulent viscosity, and how the wall distance equation is adapted with a penalty. These penalties respectively ensure that the turbulent energy approaches 0 in solid regions and that the boundaries between solid and fluid are seen as walls in the S-A model. Next, in Section 2.4.3 the energy equation is adapted to correctly model heat flow through the fluid and solid regions. The combination of objective function and constraints are explained in Section 2.4.4. Section 2.5 shortly introduces the weaknesses of the density method. A graphical overview of the variables used with the density method for thermo-fluid TO is given in Appendix D.

2.4.1. Adapting Navier-Stokes equations

In the case of a thermo-fluid TO, the thermo-fluid model needs to be altered to interface correctly with the density-based formulation. Due to its lack of explicitly defined walls between solid and fluid, the fluid model adapted for the density method requires a single formulation that applies to the complete design domain. The different behaviors of the fluid and solid regions need to be interpolated purely based on the density variables.

Focusing on the fluid velocity field, it is quite logical that no fluid should flow through the solid regions of the design domain. The mainstream approach adopted for achieving this is referred to as the ‘Brinkman penalization method’, which was originally proposed by Brinkman as an alternative to Darcy’s law to calculate the viscous force exerted by a flowing fluid on a dense swarm of particles [10]. The Brinkman penalization method was pioneered for Stokes flow in combination with the density method by Borrvall and Peterson [9], further extended to N-S flow by, among others, Olesen et al. [24] and first applied to thermo-fluid TO by Dede [13]. Many implementations of the method followed (e.g. [14], [21], [18]) with varying boundary conditions, objective functions and constraints.

The Brinkman penalization method utilizes a porous damping force (also referred to as ‘Darcy force’) defined by a locally variable permeability which is governed by the density variables. The porous damping force with unit N/m^3 , defined in Brinkman’s work as $\mu \vec{q}/\kappa$ is derived from Darcy’s law for flow through porous media which is given in

Eq. (2.26).

$$\vec{q} = -\frac{\kappa}{\mu}\nabla p \quad (2.26)$$

where \vec{q} [m s^{-1}] is the flow rate through a surface of unit area (which can also be interpreted as the mean velocity in an infinitesimally small element of porous medium), κ [m^2] is the permeability of the porous medium, μ [Pa s] is the dynamic viscosity and p [Pa] is the pressure.

Darcy's law describes that a porous medium with low permeability poses a high resistance to flow, resulting in a lower flow rate with a fixed pressure gradient. Essentially the cross-section area of a channel filled with porous medium is reduced proportionally to the permeability, compared to its 'open' configuration. In the Brinkman penalization method the porous damping force replaces the body force term which is included in the original N-S equations (Eqs. (2.1) and (2.2)). Instead of specifying the dimensional permeability of the porous medium, Li et al. [21] and Dilgen, Dilgen et al. [15] suggest to correlate the porous damping force based on the dimensionless Darcy number. In their approaches the steady-state incompressible N-S equations are adapted as given in Eqs. (2.27) and (2.28).

$$\rho\nabla \cdot \vec{u} = 0 \quad (2.27)$$

$$\rho(\vec{u} \cdot \nabla)\vec{u} = \nabla \cdot [-p\mathbf{I} + \mu(\nabla\vec{u} + (\nabla\vec{u})^T)] - \underbrace{\alpha\rho\vec{u}}_{\text{Brinkman penalty}}. \quad (2.28)$$

This approach introduces the 'inverse permeability' term, α which is equivalent to ν/κ , where ν is the fluid's kinematic viscosity. The inverse permeability will be locally interpolated based on the density variables. This is done with Darcy interpolation to enhance the sensitivities in early iterations, as is explained in Section 2.3.6. The inverse permeability is interpolated by Eq. (2.13) where k_{int} is replaced with α , α_{min} is set to 0 and q controls the shape of the interpolation. The upper limit of inverse permeability α_{max} would ideally be infinite to create fully solid walls, however this will cause numerical instability and therefore a finite number must be used. Both Li et al. [21] and Dilgen, Dilgen et al. [15] use different definitions to determine the maximum inverse permeability, which are displayed in Eqs. (2.29) and (2.30) respectively.

$$\alpha_{max} = \left(1 + \frac{1}{\text{Re}}\right) \frac{1}{\text{Da}} \quad (2.29)$$

$$\alpha_{max} = \frac{\nu}{\text{Da}L^2} \quad (2.30)$$

where Re is the Reynolds number of a fluid flowing through a channel and L is a selected length scale. The definition in Eq. (2.29) scales with the flow Reynolds number, but this effect levels out when the Reynolds number becomes large ($\text{Re} \gg 1$). The second definition (Eq. (2.30)) seems more appropriate for high Reynolds number flows because it is used for turbulent flow in its defining work. When using this definition, the characteristic length scale should be correctly selected especially when attempting to compare geometries of varying scales. In this thesis, the characteristic length scale is set similar to that of the Reynolds number for different geometry types:

- Tube: internal diameter D
- Channel or duct: hydraulic diameter D_h
- Flat plate: plate length L

2.4.2. Adapting Reynolds-averaged Navier-Stokes equations and Spalart-Allmaras turbulence model

When a fluid flow is laminar or when DNS simulations are performed on a turbulent flow, the Brinkman penalization suffices to generate correct behavior of the fluid interacting with solid regions. However, when a turbulence model is included to approximate the turbulent behavior of the fluid, the turbulence equations must be adapted to interface with the density method formulation. This is needed to provide the optimizer with gradient information with respect to the turbulent behavior of the fluid, such that it can determine how the turbulent energy influences the performance of the optimized geometry. If the adaptations are omitted but a turbulence model is still used to compute the fluid velocity, the assumption of 'frozen turbulence' is made, where the turbulence equations are assumed to not alter the gradient of the objective function [15]. Furthermore, an improper interface with the density method formulation can lead to numerical instabilities when the boundary conditions of the flow velocity field do not match with those of the turbulence model.

Adapting Reynolds-averaged Navier-Stokes equations. Firstly, instead of the classic N-S equations the steady-state RANS equations (Eqs. (2.4) and (2.5)) are penalized with the Brinkman term, resulting in Eqs. (2.31) and (2.32).

$$\rho \nabla \cdot \vec{U} = 0 \quad (2.31)$$

$$\rho \vec{U} \cdot \nabla \vec{U} + \nabla(\rho \vec{u}' \times \vec{u}') = \nabla \cdot [-p\mathbf{I} + \mu(\nabla \vec{U} + (\nabla \vec{U})^T)] - \underbrace{\alpha \rho \vec{U}}_{\text{Brinkman penalization}} \quad (2.32)$$

where the interpolation of α is according to Eq. (2.13) and the maximum inverse permeability is determined by Eq. (2.30).

Adapting Spalart-Allmaras equation. Next, the S-A equation is adapted by penalizing the turbulent viscosity (Eq. (2.7)) according to the definition introduced by Dilgen, Dilgen et al. [15], as well as penalizing the wall distance equation (Eq. (2.8)) following the definition of Yoon [34]. Similarly to the Brinkman penalization, the penalty of the S-A equation aims to bring the turbulent viscosity down to zero in solid regions. The wall distance penalty forces the wall distance to a small number in solid regions such that it approaches the behavior of a solid wall.

The definition by Dilgen, Dilgen et al. of the turbulent viscosity penalty is selected because they use the same interpolation function as the Brinkman penalization, in contrast to Yoon who uses different interpolation functions for the two. However, scaling parameters $n_{\tilde{\nu}}$ and n_G are added for the respective penalties, as these will be used in Section 3.7 for tuning the density method formulation. For brevity, re-stating the auxiliary equations and coefficients of the S-A model is omitted since they are unaltered. The S-A equation is adapted, resulting in Eq. (2.33).

$$\vec{u} \cdot \nabla \tilde{\nu} = c_{b1} \tilde{S} \tilde{\nu} - c_{w1} f_w \left(\frac{\tilde{\nu}}{l_w} \right)^2 + \frac{1}{\sigma_{\tilde{\nu}}} \nabla \cdot ((\nu + \tilde{\nu}) \nabla \tilde{\nu}) + \frac{c_{b2}}{\sigma_{\tilde{\nu}}} \nabla \tilde{\nu} \cdot \nabla \tilde{\nu} - \underbrace{\alpha n_{\tilde{\nu}} \tilde{\nu}}_{\text{Added penalty}}. \quad (2.33)$$

Adapting wall distance equation. The reciprocal wall distance equation is adapted resulting in Eq. (2.34).

$$\nabla G \cdot \nabla G + \sigma_w G (\nabla \cdot \nabla G) = (1 + 2\sigma_w) G^4 + \underbrace{\alpha n_G (G - G_0)}_{\text{Added penalty}}. \quad (2.34)$$

Because the fluid solver solves for the reciprocal wall distance G instead of the actual wall distance l_w , the boundary condition at walls is non-homogeneous. Ideally the reciprocal wall distance will go to infinity at a wall, but this causes numerical instabilities. Therefore the reciprocal wall distance is forced to approach the default reciprocal wall distance at a wall G_0 which is computed automatically by COMSOL as

$$G_0 = \frac{2}{l_{ref}}$$

where the geometric reference length l_{ref} is computed as one tenth of the shortest length of the geometry's bounding box.

Note on wall functions. All literature that implement a turbulent TO of a fluid flow have thus far avoided using a turbulence closure model which uses wall functions. Yoon [34] and Dilgen, Dilgen et al. select the S-A model [15]. Furthermore Dilgen, Dilgen et al. implement the $k - \omega$ model, which is also a low-Reynolds number model. Yoon implemented the $k - \epsilon$ model (which is a model that uses wall functions) in another work [35], but mentions that the wall functions have been disabled due to 'difficulties in the implementation and the sensitivity analysis.' For this reason, the use of a turbulence model with wall functions has been avoided in this thesis.

The adaptation of the RANS equations and the S-A turbulence model enables a realistic boundary layer to form along walls defined by the density variables. However, the heat transfer between the solid and fluid regions is not behaving correctly yet because only one set of thermal properties (thermal conductivity, thermal capacity, density) is applied to the whole simulation domain. To correct this, the energy equation will be adapted as well.

2.4.3. Adapting energy equation

With the flow velocity, turbulent viscosity and wall distance forced to the correct values in solid regions, all that is left to adapt is the energy equation (Eq. (2.3)) such that the heat flow through solid and fluid regions can be governed correctly by this equation. The approach is not to alter the energy equation, but to interpolate the relevant material properties based on the density variables, such that the energy equation becomes implicitly dependent on the density variables. The following material properties are interpolated:

1. Mass density $\rho(\gamma)$
2. Thermal capacity $C_p(\gamma)$

3. Thermal conductivity $k(\boldsymbol{\gamma})$

With the interpolation of the material properties the steady-state energy equation results in Eq. (2.35).

$$\rho(\boldsymbol{\gamma})C_p(\boldsymbol{\gamma})\vec{u} \cdot \nabla T + \nabla \cdot \vec{q} = Q \quad (2.35)$$

where \vec{q} is determined by Eq. (2.10) with all material properties interpolated as stated. The interpolation of these material properties is done with a SIMP interpolation (Eq. (2.12)) using one common interpolation penalty p . The minimum and maximum values of each property are based on material properties for the solid and fluid, respectively.

2.4.4. Relevant objective function and constraints.

Many different combinations of objective functions and constraints have been found in literature, some of which have been summarized in the literature study preceding this work [30]. The relevant functions which have been implemented in this work will be introduced in this section. Note that any of these functions can be implemented either as (part of) the objective function or as a constraint. Their implementation will be discussed in Chapter 6.

Average heat source temperature. Because the TO of the power module will contain local heat sources, the average temperature must be determined locally. An average local temperature is used instead of a maximum temperature, because the $\max()$ function is not differentiable and will cause issues with the adjoint sensitivity analysis. The average heat source temperature is stated in Eq. (2.36).

$$T_{avg} = \frac{1}{A_{src}} \int_{A_{src}} T dA \quad (2.36)$$

where T_{avg} is the average heat source temperature, A_{src} is the area of the heat source and T is the local temperature.

Fluid power losses (global). Following the definition proposed by Dede [13], but with similar implementations by Koga et al. [18] and Li et al. [21], the fluid power losses in the simulated domain are formulated in Eq. (2.37)

$$P_{loss} = \int_{\Omega} \left[\frac{1}{2} \mu \sum_{ij} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)^2 + \sum_i \alpha(\boldsymbol{\gamma}) u_i^2 \right] d\Omega \quad (2.37)$$

where P_{loss} is the total fluid power loss, u_i is the velocity component in the i direction and $\alpha(\boldsymbol{\gamma})$ is the inverse permeability.

Fluid power losses (local). Instead of using a global integral to determine the fluid power losses, Dilgen, Dilgen et al. [16] propose to use a boundary integral on the inlet and outlet to determine the difference in fluid power. Li et al. also suggest a similar approach at first, but dismiss this because the global integral is deemed more useful. The local fluid power losses are computed in Eq. (2.38).

$$P_{loss} = - \int_{\Gamma} \left[\vec{u} \cdot \vec{n} \left(\frac{1}{2} \vec{u} \cdot \vec{u} + p \right) - 2\nu_{eff}(\boldsymbol{S}\vec{n}) \cdot \vec{u} \right] d\Gamma \quad (2.38)$$


where Γ are the inlet and outlet boundaries, \vec{u} and \vec{n} are the velocity- and boundary normal direction vector, respectively; p is the local pressure, ν_{eff} is the effective kinematic viscosity given by $\nu_{eff} = \nu_t + \nu$; \boldsymbol{S} is the mean strain rate tensor defined in Eq. (2.39).

$$\boldsymbol{S} = \frac{1}{2} (\nabla \vec{u} + \nabla \vec{u}^T). \quad (2.39)$$



2.5. Weaknesses of density method in combination with thermo-fluid Topology Optimization

Among other TO strategies, the density method is likely the most robust approach for many implementations and provides satisfactory accuracy. However, the method has some issues which arise in edge-cases. When referring to thermo-fluid TO there are two weaknesses which can cause problems, especially in turbulent flows.

Porous solid material. Firstly, the Brinkman penalty cannot reach infinity, thus solid regions are never modeled as fully solid but instead have a certain porosity. This results in pressure diffusion through these regions. Figure 2.6a displays a reference pressure distribution of an external flow past a cylinder, and Fig. 2.6b displays the pressure distribution of the same scenario with a fluid model adapted to the density method, where it is visible that the pressure propagates smoothly through the solid region when using the density method.





(a) Reference simulation, conventional fluid model





(b) Simulation with density-based fluid model,
 $Da_D = 10^{-3}$

Figure 2.6: Comparison of the pressure distributions of an external flow past a cylinder when modeled conventionally and with the density method

Furthermore, a measurable velocity magnitude through these porous solids is visible. Figures 2.7a and 2.7b display the comparison of the velocity distribution between conventional modeling and the density method for aforementioned scenario. Since heat transfer through solids should be governed only by thermal conduction, the additional convective heat flow caused by a fluid flowing through the solid is unwanted. The error caused by this effect can be studied, and possibly a method can be found to reduce this error.



(a) Reference simulation, conventional fluid model



(b) Simulation with density-based fluid model,
 $Da_D = 10^{-3}$

Figure 2.7: Comparison of the velocity distributions of an external flow past a cylinder when modeled conventionally and with the density method

Blurred walls. Secondly, the transition from solid to fluid can never be perfectly ‘sharp’ with the density method, as a blurring filter must be used to avoid checkerboarding and to enforce a minimum length scale. These ‘blurred walls’ cause the velocity and thermal boundary layers to have a different shape compared to the reference, which can result in a large over- or under-estimation of the heat transfer from solid to fluid.

Fig. 2.8 displays the velocity distribution of a developing laminar boundary layer along a flat plate, modeled with the density method. Here, fluid enters the simulation domain with a uniform velocity at the left vertical external boundary. It travels uninterrupted until it crosses the internal vertical boundary, where it meets the leading edge (defined as $x = 0$) of a solid density-based flat plate which is placed at the bottom of the fluid region. The interface between the solid and fluid domain is modeled with a transition of the density variables, resulting in a ‘blurred wall’. From the leading edge of the plate, a boundary layer develops in the streamline direction of the flow.


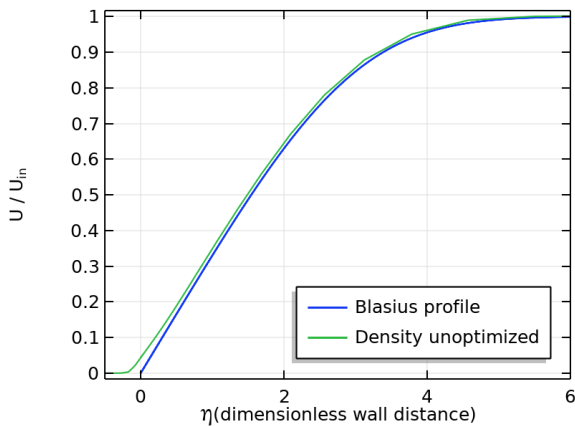
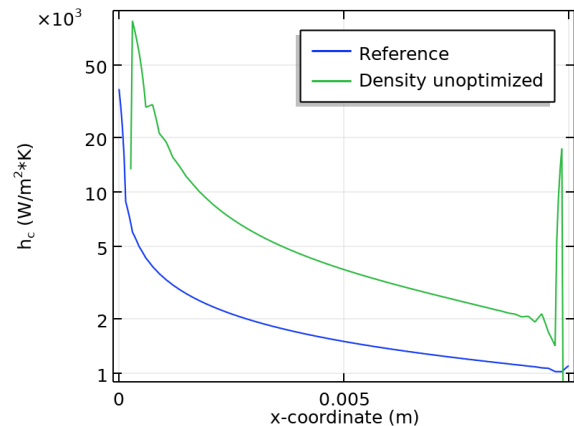


Figure 2.8: Density-based simulation velocity distribution. Fluid enters at the left vertical external boundary with uniform velocity (white arrows) and meets the leading edge of the flat plate (defined as $x = 0$) at the internal vertical boundary (black line). From the leading edge, a boundary layer develops in streamline direction.

The developing boundary layer looks ordinary when viewing the velocity distribution, but when a cross-section velocity profile is plotted in dimensionless coordinates (Fig. 2.9a) a significant error is visible compared to the Blasius profile (an analytical solution for laminar boundary layers, refer to Appendix A.1 for its formulation). Especially close to the porous wall the profile of the density method shows a large deviation, following an ‘exponential decay’ curve towards zero velocity instead of the expected constant slope. Even through a portion of the wall the velocity is significant. The velocity profile in the region close to the wall governs the heat transfer behavior from the solid to the fluid, as the balance between conductive and convective heat transfer is dependent on the fluid’s velocity. An erroneous velocity profile will therefore introduce an error in the heat transfer behavior along the length of the wall (Fig. 2.9b). The error can be reduced by choosing a smaller blurring filter radius, resulting in sharper transitions. However, the mesh resolution must scale with the filter radius to ensure accurate convergence of the filter equation. This poses a computational cost limit to the sharpness of gradients. An analysis can be made on how much error is introduced by the blurred walls, and methods can be discovered to minimize the error.



(a) Density-based simulation dimensionless velocity profile compared to Blasius profile



(b) Density-based simulation heat transfer coefficient compared to reference simulation

Figure 2.9: Plots displaying an erroneous boundary layer of a density-based simulation. A dimensionless velocity profile plot (a) shows that a significant error exists compared to the analytical solution of the Blasius profile. This erroneous velocity profile causes an extreme oscillation and a net over-estimation of heat transfer (b). Parameters used: $Da = 10^{-7}$, $q = 0.01$, $p = 3$, $\beta = 8$, $Re_{D_h} = 1000$

2.6. Research questions

The aim of this research is to gain insight in the accuracy and performance of thermo-fluid TO with the density method and to find methods to mitigate some of the weaknesses of this method. The Prodrive Power Module is a good case for turbulent thermo-fluid TO, as it pushes the heat flux of single phase forced convection to its limit. From the researched literature a number of issues and gaps have been found, which lead to a set of research questions that this thesis aims to answer.

The accuracy of the density method in combination with thermo-fluid TO has never been studied quantitatively. Because of this, the weaknesses of the method have not been quantified yet. The issue of porous solid material is mentioned in some literature, but there it is assumed to be insignificant when a sufficiently high Brinkman penalty is selected. The issue of blurred walls is usually assumed to be mitigated by minimizing large regions of intermediate density values with higher penalization of the interpolation functions, or by adding a projection filter. Both causes for error have never been studied quantitatively, and therefore no advanced mitigation methods have been found for these issues. Furthermore, the manually tuned parameters of the density method formulation (Darcy number, interpolation penalties) have not been quantitatively proven to show convergence. The assumption of ‘the higher the penalty, the better accuracy’ is therefore not proven to be valid. The currently applied heuristic countermeasures are therefore not guaranteed to work, which can lead to inferior optimization results and the need for tedious parameter-tuning.

Based on these observations, a main research question and a number of sub-questions are defined:

How can the errors caused by the two main weaknesses of the density method when applied to thermo-fluids, being ‘porous solid material’ and ‘blurred walls’, be minimized?

RQ.1 How can the errors caused by these weaknesses be analyzed?

RQ.1a How can the modeling accuracy of the density method when applied to thermo-fluid TO be defined?

RQ.1b How do the modeling parameters of the density-based fluid models affect the modeling accuracy?

RQ.2 To what extent can mitigation strategies minimize this error?

RQ.2a By adapting the thermal conductivity of the solid?

RQ.2b By tuning the modeling parameters of the density method?

2.7. Research plan

The research plan of this work is explained following the chart in Fig. 2.10. The chart starts with the first milestones that answer research sub-questions, and works towards answering the main question in the end.

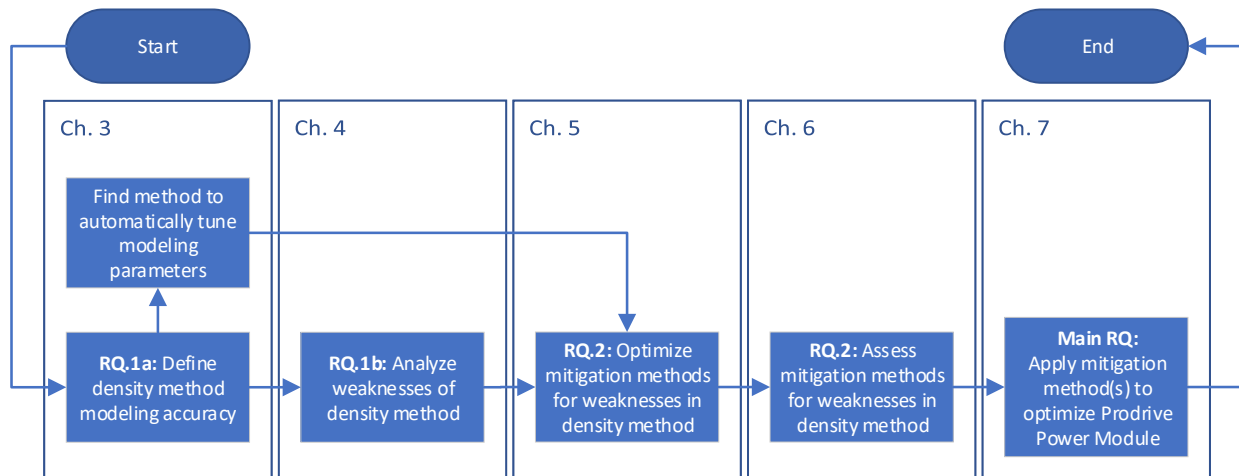


Figure 2.10: Research plan flowchart

RQ.1a: Define density method modeling accuracy. As the first step, research question RQ.1a will be answered. The ‘modeling accuracy’ of the density method has not been quantitatively researched yet, and therefore no clear definition has been made. A reference must be created or found that symbolises the real-life behavior of the fluid. The choice is made to build a high-fidelity reference CFD simulation that is verified by a number of correlation functions to ensure a high accuracy. Chapter 3 discusses the approach to define the modeling accuracy of the density method.

Find method to automatically tune modeling parameters. From literature it is known that the modeling parameters of a thermo-fluid TO should be tuned to achieve desirable results. Possibly, an optimal combination of the modeling parameters can be found that maximizes the modeling accuracy. The tuning of the modeling parameters can be done manually by trial-and-error, but preferably a method is found to tune them automatically. Note that this optimal combination is likely dependent on the flow conditions (Reynolds number) and the material properties. Therefore the tuning of the modeling parameters should ideally be performed for every change of these dependencies. Chapter 3 also discusses how the modeling parameters can be optimized automatically.

RQ.1b: Analyze weaknesses of density method. To answer research question RQ.1b, in Chapter 4 an in-depth analysis is performed on the weaknesses of density-based turbulent thermo-fluid TO. The analysis dissects the weaknesses from their global effect down to the influence of individual modeling parameters on causing them to occur. In Appendix C a ‘black box’ approach is taken to analyze the effect of certain modeling parameters on the modeling accuracy by generating a large random dataset of modeling parameter combinations.

RQ.2: Optimize mitigation methods for weaknesses in density method. The mitigation methods for the weaknesses in the density method that are proposed in research questions RQ.2a and RQ.2b are dependent on some modeling parameters which need to be tuned. The mitigation methods will be optimized with the automated tuning method introduced in Chapter 3. After optimizing the mitigation methods, the accuracy of the density-based fluid models can be quantified, answering research question RQ.2. In Chapter 5, the optimized mitigation methods will be introduced and their effect on the modeling accuracy will be discussed.

RQ.2: Assess mitigation methods for weaknesses in density method. The mitigation methods for the weaknesses in the density method are assessed in Chapter 6. Based on a test-problem, a comparison is made between a density-based fluid model with and without the mitigation strategies. Here it can be determined how the modeling accuracy resulting from the mitigation methods translate to a different geometry which is more relevant to the TO problem.

Main RQ: apply mitigation methods to optimize a Prodrive Power Module cooling infrastructure using TO. The developed method with mitigation strategies can be applied to the thermo-fluid TO (both laminar and turbulent) of a Prodrive Power Module cooling infrastructure, which is discussed in Chapter 7. A comparison is made between a manually tuned TO and a meta-optimized TO to determine whether the mitigation methods can result in increased cooling performance. In the process of implementing the TO, notes will be made on the stability of the optimization. The goal is to provide a first indication on the effect that the mitigation methods can have on the thermo-fluid TO.

Defining modeling accuracy of the density method

This chapter discusses the methods used to quantify the accuracy of the density method. To quantify the modeling accuracy, a comparison is made between reference data of a conventional flow model (walls 100 % solid, body fitted mesh), and the density method formulation where walls are modeled by a gradient of the density variables. Section 3.1 explains how this reference is constructed. The modeling accuracy will initially be analyzed on one specific scenario that is relevant to the implementation in a Prodrive Power Module. Therefore the boundary conditions of the reference simulation are determined based on applicable ranges set by Prodrive Technologies. The selection of boundary conditions is explained in Section 3.2. The FEM discretization and material selection are discussed in Sections 3.3 and 3.4, respectively. Section 3.5 introduces the simulations used to generate data of the density method which can be compared to the reference data. Finally, a formal measure of ‘modeling accuracy’ is proposed in Section 3.6.

3.1. Setting up a reference simulation

Different geometries can be used to analyze the accuracy of the fluid models which are adapted to the density method. Ideally the accuracy of the fluid models is directly obtained by comparing their results to the real-world behavior of a fluid flowing along / through the same geometry. Such real-world information can be obtained through two methods: by generating experimental data for the specific geometry, or by using previous experimental data obtained for the specific geometry through the use of correlation functions. Obtaining sufficient experimental data to verify a fluid model requires a sophisticated test setup and many experiments, which is beyond the scope of this thesis and therefore this option is eliminated. For some geometries more detailed correlation functions are available than for others and ideally as much detailed data as possible is desired to reliably assess of the accuracy of the adapted fluid models. Furthermore, for reasons that will become apparent in Section 3.7, a high fidelity reference dataset is desired to obtain a proper method that automatically tunes the modeling parameters of the adapted fluid models towards high accuracy (referred to as meta-optimization).

However, no geometries have been found that have sufficiently detailed correlation functions which can be used as reference for an accurate meta-optimization. The available correlation functions are often 0D functions (a global evaluation of some quantity, such as an average heat transfer coefficient) or at most 1D functions (a line-evaluation of a local quantity along a predefined boundary, such as the local heat transfer coefficient along a wall). Working with 2D or even 3D reference data is preferred for the meta-optimization, and therefore the choice is made to create a reference CFD simulation that matches the available correlation data as accurately as possible. This reference simulation creates much more detailed data to use as reference for the meta-optimization. Sufficient correlation data is still required for the selected geometry to ensure that the reference CFD simulation is accurate.

A second aspect for selecting the geometry of the reference simulation is the relevance of the geometry to the expected geometry in the TO. The resulting parameter set from the meta-optimization should be applied to the TO, to ensure an accurate TO result. However, the expectation is that the fluid flow behaves differently when the geometry is changed (for instance comparing an external flow with internal flow) and because of this difference the most accurate parameters which are meta-optimized for a certain geometry might not be the most accurate when applied to another geometry. Therefore the meta-optimization should be performed on a geometry that is relevant for the expected resulting geometry from a TO.

Finally, a geometry which brings out the weaknesses of the density method is sought, being firstly that a solid wall in the density method formulation can never be fully solid and some fluid will penetrate through it. The regions of high penetration are found where the velocity gradient - and with that the shear stresses - at a wall are highest. This may be

a location where a channel becomes narrower, a location where a sharp bend occurs, or when a jet of fluid impinges onto an object (such as a flat wall, cylinder or splitting channel). Essentially anywhere where the boundary layer of the fluid flowing along a wall is forced to decrease in thickness, a peak in shear stress will occur and some amount of fluid penetrating solid regions of a density-based fluid model is expected. Therefore a geometry is desired that exhibits a (local) region of high shear stress to evaluate the amount of fluid penetrating the solid wall. The second weakness, where the heat transfer behavior is affected by the shape of the density gradient at walls, is likely to appear regardless of the selected geometry.

The combination of the three aforementioned wishes for a certain geometry (1: have sufficiently detailed correlation data available, 2: be relevant to the resulting geometry of a TO and 3: exhibit a (local) region of high shear stress) results in a trade-off between different geometry options, which will be made in Section 3.1.1. A strategy is introduced in Section 3.1.2 to verify the accuracy of the reference simulation.

3.1.1. Geometry selection

There are two relevant geometries within the scope of this research:

1. Internal flow through channel with developing boundary layer
2. External flow along a flat plate with developing boundary layer

These two cases will be shortly introduced in the next paragraphs, followed by a strategy to verify their accuracy using correlation functions.

Internal flow through channel. The most relevant geometry to the TO result is that of a tube or channel (Fig. 3.1a). This is because the flow through the optimized cooling infrastructure will be internal. The choice is made to study a developing internal boundary layer instead of a fully developed flow, because the former of the two exhibits a region of high shear stress at the leading edge of the channel which decreases as the boundary layer develops. This effect is not visible with a fully developed flow, since the boundary layer thickness, and therefore shear stress, does not vary along the wall in this case. High wall shear stresses can still be enforced in a fully developed flow by increasing the Reynolds number of the flow, however these Reynolds numbers may become much higher than what is relevant for the TO of a cooling infrastructure. When simulating a channel in 2D, it is modeled as two plates that extend infinitely out-of-plane and the characteristic length scale being the hydraulic diameter $D_h = 2H_{chan}$.

A number of correlation functions exist for a fully developed flow through a channel or tube, whereas none exist for a developing boundary layer. This makes it difficult to verify the accuracy of a developing flow through a channel.

External flow along flat plate. An external flow along a flat plate with developing boundary layer (Fig. 3.1b) shows relatively similar behavior compared to a developing internal boundary layer. It exhibits a region of high wall shear stress at the leading edge of the plate, as well as decreasing wall shear stresses as the boundary layer develops. However, there are some key differences. The boundary layer thickness of an external flow along a flat plate can develop to infinite thickness whereas in a channel it can only become half of the channel height until the flow is fully developed. An external boundary layer along a flat plate has an outer layer that always remains at the inlet velocity, whereas an internal boundary layer has an ‘inviscid core’ which accelerates due to the conservation of mass (the velocity at the wall decreases with respect to the inlet velocity, so the core velocity must increase). The characteristic length scale of this geometry is the length of the plate L .

Many correlation functions are available for the flat plate geometry, such as: average heat transfer coefficient, local heat transfer coefficient along the plate, local shear stress along the plate, local boundary layer thickness, velocity profiles, etc. The differences in behaviour with respect to an internal boundary layer in a channel might make the external flow along a flat plate geometry less relevant for the TO of a cooling infrastructure.

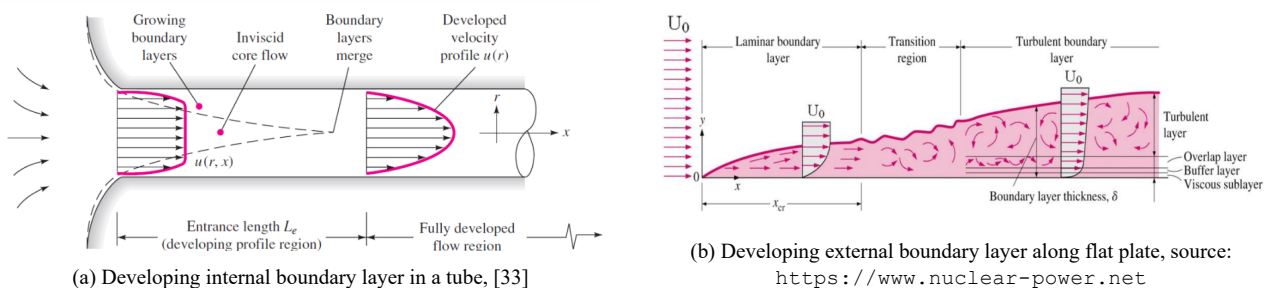


Figure 3.1: Comparison of a developing boundary layer in two different geometries

3.1.2. Verification strategy

Ideally, the channel geometry with a developing boundary layer will be used for the reference simulation because a channel is the only geometry which can be directly applied to the TO problem and a developing boundary layer shows a region of high shear stress. However, due to lack of correlation functions, the accuracy of a reference simulation based on this geometry cannot be directly verified. If the simulation features a channel which is long enough to fully develop the flow, the correlations for a fully developed flow can be used to indirectly verify the simulation. However, the entry effects of the channel might still show an error even though the fully developed flow is accurate. Therefore a more accurate verification of the reference simulation is sought.

The boundary conditions for a simulation of an external flow along a flat plate and an internal flow through a channel are largely similar, although the flow behavior is not comparable. Due to the different definition of the characteristic length scale for the two geometries (the plate uses the distance from the leading edge, the channel uses the hydraulic diameter) and because the boundary layers develop differently, the result of one geometry cannot be converted to the other geometry.

A strategy is devised to allow the use of the developing boundary layer in a channel with sufficient verification, see Fig. 3.2. First, a reference simulation of the plate geometry is verified based on the available correlation functions. This process is discussed in Appendix A.1. The boundary conditions of this reference simulation are then adapted to the channel geometry with developing internal boundary layer. By using the configuration of the verified reference simulation of the plate geometry, it is assumed that a similar accuracy is achieved for the channel reference simulation. The correlation functions for a fully developed flow are used to check the local accuracy of the channel flow. This check is discussed in Appendix A.2. Although the results of the two reference simulations cannot be compared and the channel geometry cannot be directly verified, this strategy should reduce the possible error of the channel reference simulation.

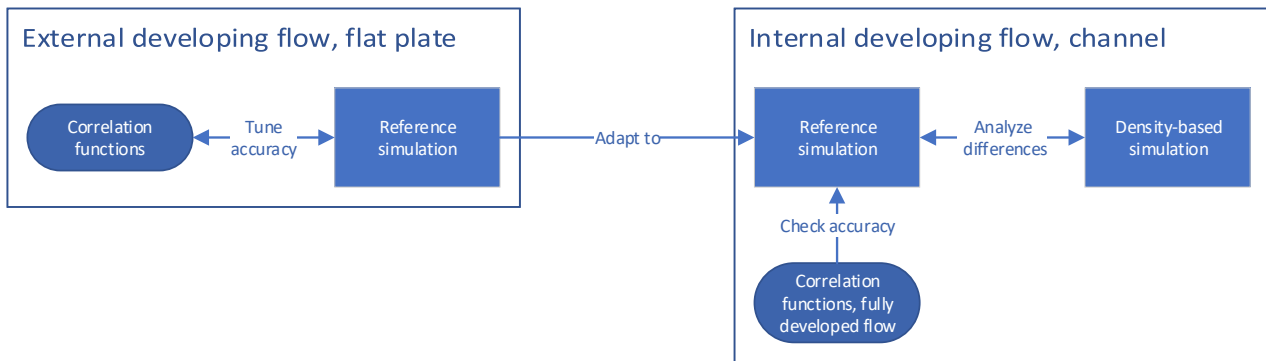


Figure 3.2: Verification strategy, using the verified result of an external flow along a flat plate to ensure high accuracy of a developing boundary layer in a channel without directly verifying the latter geometry

3.2. Boundary conditions

Both the reference simulation of the external flow along a flat plate and the internal developing flow through a channel are based on a “2D Zero Pressure Gradient Flat Plate Verification Case” defined by NASA [27]. First, the reference simulation of the flat plate is constructed from NASA’s case with as little alteration as possible. This results in:

- A uniform inlet velocity
- A zero-pressure outlet
- A slip-wall or symmetry condition in front of the leading edge of the plate
- A no-slip wall to develop the boundary layer
- An open boundary at the top of the domain

One main alteration is necessary: NASA’s definition of the verification case does not feature a heat source, and therefore one is introduced for the reference simulation in this thesis. Adding heat transfer to the reference is realized by specifying a heated boundary at the bottom of the domain.

After the flat plate reference simulation is verified (Appendix A.1), a number of adaptations are made with respect to the flat plate geometry to convert the external flow along a flat plate to an internal developing flow through a channel. This is achieved by switching the top boundary to a symmetry condition, by setting the simulation domain to $1/2$ of the channel height and by defining the inlet velocity based on Re_{D_h} . Figure 3.3a illustrates the various boundary conditions used for the channel reference simulation and Fig. 3.3b displays the boundary conditions for the density-based channel simulation. The following sections will discuss the inlet and heated wall boundary conditions in more detail.

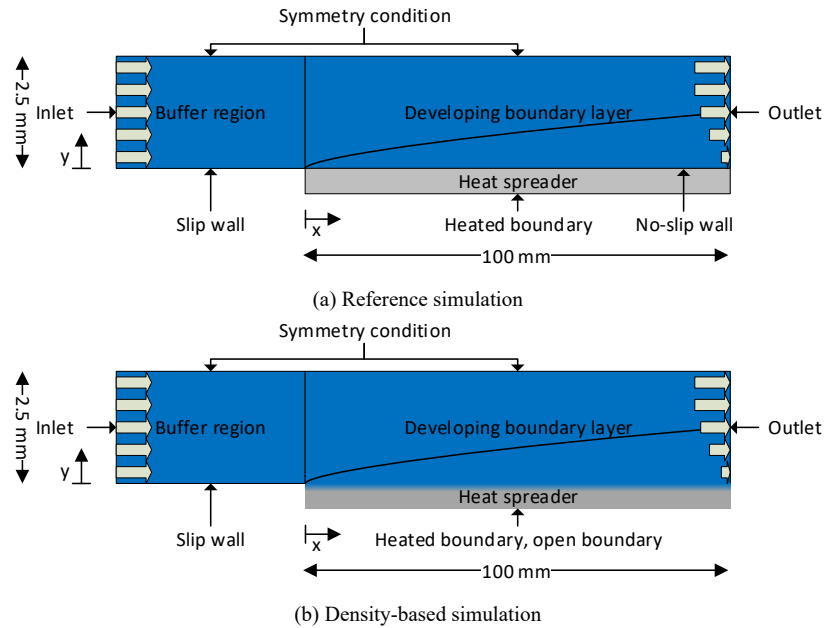


Figure 3.3: Overview of domains and boundary conditions for reference simulation and density-based simulation of developing boundary layer in a channel. The aspect ratio of the simulation domain (length/height) is much higher than illustrated here.

3.2.1. Inlet Reynolds number range

The most important boundary condition for the cooling system of a Power Module is the inlet velocity, which is defined by the inlet Reynolds number based on the inlet channel hydraulic diameter Re_{D_h} . This velocity directly determines the heat transfer capability of the system. For the TO an inlet Reynolds number should be selected that can later be implemented in a prototype and the same Reynolds number should be used to analyze the accuracy of the density method for thermo-fluid TO. A range for the inlet Reynolds number is determined based on ranges of the flow rate (1 L min^{-1} to 10 L min^{-1}) and inlet diameter (5 mm to 20 mm) that are applicable within Prodrive Technologies. The relation between the inlet diameter, flow rate and inlet Reynolds number is displayed in Fig. 3.4. The range of applicable inlet Re_D is rather broad, with a minimum of approximately 1000 and a maximum of 42 500.





Figure 3.4: Range of inlet Reynolds number based on flow rate and inlet diameter ranges of applications in Prodrive Technologies. Red contour lines indicate the lower and upper limit Reynolds numbers of the turbulent transition regime

Within the range of applicable Reynolds numbers, a turbulent inlet Reynolds number should be selected to analyze the accuracy of the adapted turbulent fluid model, and similarly a laminar Reynolds number should be selected for the adapted laminar fluid model. In a tube, the transition from laminar to turbulent flow occurs in a range of Reynolds numbers between 2300 to 10 000 depending on the wall and inlet conditions (wall roughness, steadiness of the incoming fluid). The lower and upper limits of the transition to turbulence are also displayed in Fig. 3.4 as red contour lines. This indicates that the majority of the applications lie in the transitional or turbulent flow regime. However, since this thesis focuses on cooling of small local heat loads the TO-optimized cooling infrastructure will likely have small channels.

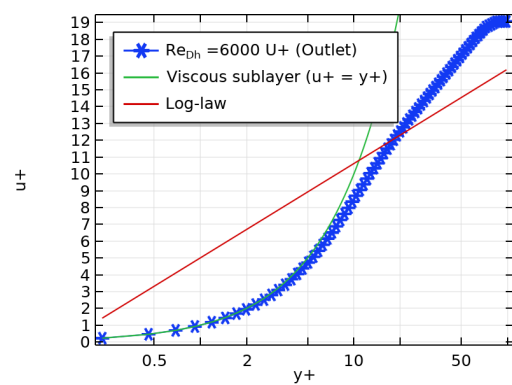
This causes high pressure drops and limits the flow rate. The curved shapes and irregular walls of the resulting geometry will however cause turbulence to trip sooner than the smooth straight wall of the reference simulation. Therefore it is not trivial to assume that the flow of the TO will be turbulent at a Reynolds number above 2300. The result of the reference simulation is thus qualitatively checked for turbulence by two indicators:

- Turbulent viscosity. When the turbulent viscosity starts increasing along the developing boundary layer, this indicates the transition to turbulence. As a rule-of-thumb, the turbulence is fully established when the turbulent viscosity becomes 10 times larger than the physical viscosity of the fluid.
- Dimensionless velocity profile (u^+ with respect to y^+ , see Appendix A.1.1). When the boundary layer shows signs of a ‘log-law’ layer, this indicates transition to turbulence.

The combination of the turbulence indicators and a feasibility study of a turbulent thermo-fluid TO shows that at a Reynolds number of ≈ 6000 the reference simulation shows sufficient turbulence (see Figs. 3.5a and 3.5b), the TO result shows full turbulence and the TO is relatively stable. Therefore this Reynolds number is selected for the turbulent reference simulation. For the laminar reference simulation, the lowest applicable Reynolds number of 1000 is unfortunately too high to ensure a laminar flow throughout a geometry generated by a laminar TO (as will be discussed in Section 7.1). A Reynolds number of 100 is selected, based on numerical experiments and on what is often used in literature [30].



(a) Maximum turbulent viscosity ratio μ_T/μ with respect to x-coordinate, plotted for different values of Re_{D_h} . The maximum turbulent viscosity ratio is defined as the highest value in a vertical cross-section through the channel



(b) Velocity profile at the outlet of the reference simulation at $Re = 6000$

Figure 3.5: Two plots, each displaying a factor used to determine if a flow shows turbulence. At $Re_{D_h} = 6000$ the boundary layer of the turbulent reference simulation is far from fully turbulent ($\mu_T/\mu < 10$ and the velocity profile does not follow the log-law), but it does show sufficient turbulence

3.2.2. Heated wall condition

In Section 3.2 a heated boundary is introduced to create a thermo-fluid reference simulation. The heated wall condition should be common between the reference and density-based simulations. The behavior of the density-based simulation is leading in selecting the heat source type and geometry. A constant heat flux boundary condition is selected in contrast to a constant temperature, because the former is most relevant to the application of cooling electronic components. The ‘blurred boundaries’ effect of the density method causes the heat transfer from solid material to a fluid to occur in a region of interpolated porosity and thermal conductivity (the density variable gradient). The balance between conductive and convective heat transfer through this region is uncertain. The density-based heated wall condition is specified as a constant conductive heat flux because it contacts predominantly solid material. This heated wall is also an open boundary which allows pressure to diffuse through the blurred boundary and fluid to flow out of the domain through the bottom. Because of the open boundary, the total heat flow from the heated boundary is a sum of the conductive and convective heat flow (outflow of fluid through the bottom boundary can reduce the total heat flowing into the domain). To ensure predictable heat transfer from the heated wall it is decided to place it at a distance from the solid-fluid interface. The region of solid material placed between the heated wall and the solid-fluid interface is referred to as a ‘heat spreader’, see Fig. 3.3b. The open boundary at the bottom of the heat spreader is likely an extreme case of the convection caused by pressure diffusion, as the heat spreader is fairly thin compared to geometries generated by TO which suffer from this effect (e.g. branching channels). The heat spreader region can be made thicker to increase resistance to the fluid being forced through the solid material, but this comes with additional computational cost.

The reference simulation should have a heat source which can be reproduced with the density-based simulation, and therefore a heat spreader is also added to the reference, see Fig. 3.3a. The heated boundary is placed at the bottom of this heat spreader. This decreases the match between the reference simulation and the correlation functions, as the heated wall which interfaces with the fluid is neither isothermal nor does it have a constant heat flux. However, the effect occurs

equally in both the reference and the density-based simulation and therefore it is assumed to not affect the analysis of the modeling accuracy. The height of the heat spreader region (i.e. the distance between the solid-fluid interface and the heated wall) is selected such that the blurred boundary of the density-based simulation fits inside the heat spreader with some leeway. This allows the blurred boundary to shift downwards and gives room to increase the density variable transition length L_{trans} (Eq. (2.18)), as this will be part of the accuracy analysis.

3.3. FEM discretization

The reference simulation is discretized by a structured quadrilateral finite element mesh with first-order elements. The mesh is refined in the fluid domain towards the solid-fluid boundary. The turbulent reference simulation at $Re_{D_h} = 6000$ has its first evaluation point at $y^+ = 0.23$ (see Fig. 3.5b), which is accurate enough to fully resolve the viscous sublayer. A pseudo time step solver is used which transforms the steady-state equations into ones that are solved on a virtual time scale until the steady-state solution is obtained [11]. This eliminates the need for an accurate initial condition for the fluid solver. The default consistent stabilization techniques (streamline and crosswind stabilization) are used for the Navier-Stokes, turbulence and energy equations. The Navier-Stokes equations require this consistent stabilization to ensure stability [11]. Consistent stabilization provides numerical diffusion in regions where the numerical accuracy is insufficient to resolve the flow, but is reduced as the solution converges. It therefore does not affect the accuracy of the final solution.

3.4. Material selection

The material properties used for the reference simulation, as well as all other simulations in this thesis, are stated in Table 3.1.

Table 3.1: Material properties used for all simulations in this thesis


Material	Property	Symbol	Value	Unit
Solid	Density	ρ_s	2700	kg/m ³
	Conductivity	k_s	400	W/(mK)
	Thermal capacity	cp_s	910	J/(kgK)
Fluid	Density	ρ_f	1000	kg/m ³
	Conductivity	k_f	0.6	W/(mK)
	Thermal capacity	cp_f	4100	J/(kgK)
	Dynamic viscosity	μu_0	0.001	Pa s

3.5. Setting up a density-based verification simulation

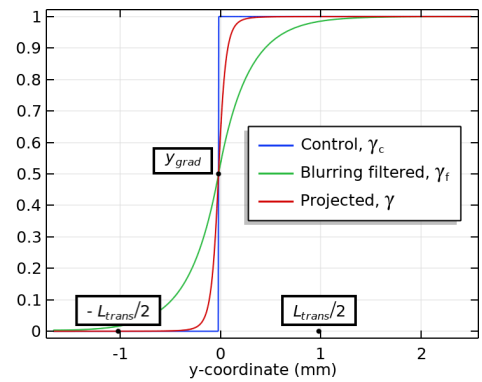
To enable a direct comparison between the reference simulation and a simulation using the density method, a simulation is constructed with the same scenario and boundary conditions as the reference simulation, but which uses a density-based fluid model. A gradient of density variables with transition length L_{trans} and centered about y-coordinate y_{grad} is used to realize the density-based equivalent of a solid wall, see Figs. 3.6a and 3.6b. Anisotropic plots are used to display the channel in this thesis starting from Fig. 3.6a, because the channel is long compared to its height. To generate this gradient of density variables, a step of the control density variables γ_c is generated at y_{grad} (Fig. 3.6b) and a blurring filter is applied that generates a gradient of filtered density variables γ_f with transition length L_{trans} according to Eq. (2.18). A projection filter is then applied with slope β , steepening the gradient and generating the output density variables γ .

In contrast to a thermo-fluid TO, a mesh with refinements is used for the density-based verification simulation. This is done to reduce computational cost and is possible because the position of the blurred wall is roughly known (in contrast to a TO where walls can be formed at any position in the domain). The refinements ensure a sufficient resolution to resolve the viscous sublayer of a turbulent boundary layer without the computational cost of a fine mesh throughout the whole domain. Two requirements dictate the mesh resolution: the height of the first fluid element neighboring the wall in viscous coordinates (y^+) shall stay below unity as a rule of thumb and the element size shall be smaller than the blurring filter's radius to ensure an accurate solution of the filter equation.

The first requirement becomes questionable when applied to the density method, as there is no sharp wall and thus no 'first element'. Therefore the mesh resolution of the density-based verification simulation is chosen to be the same as the reference simulation, for which the 'first element height' is chosen well below unity in viscous units. The assumption is made that the resolution in viscous coordinates is transferable between the reference and density-based simulations.



(a) A manually specified density gradient is used to simulate the boundary between solid (black) and fluid (white)




(b) Cross-section of the density gradient, displaying the relation between the different density variables. The solid-fluid transition length is annotated


Figure 3.6: Graphical representation of the blurred boundary between solid and fluid when using the density method for thermo-fluid TO. In the example the modeling parameters are $L_{trans} = 2$ mm, $y_{grad} = 0.01$ mm, $\beta = 8$

The requirement of the element size with respect to the blurring filter's radius is ideally met throughout the domain. However, this can result in extremely dense meshes if the filter radius is chosen to be small. According to the theory of the blurring filter (Section 2.3.7) and as displayed in Fig. 3.6b the density variables approach zero and unity exponentially without projection filter. Because of this, the Darcy force also tends to zero exponentially. It is expected that the Darcy force does not influence the fluid velocity below a certain threshold, as the viscous forces in the fluid become dominant. This indicates that the required accuracy of the density variables is lower in regions where it has no effect on the fluid, and therefore the element size can become larger than the blurring filter radius in these regions. Because the region where the Darcy force affects the fluid flow is roughly known in the density-based verification simulation, the mesh refinement is tuned such that the filter radius requirement is met only in the influential regions.

When a direct comparison is made between the reference and the density method (i.e. the same geometry and boundary conditions are used), the difference of simulation outputs between the two can be plotted by subtracting one dataset from the other. To illustrate what such a dataset comparison may look like, Figs. 3.7a and 3.7b show the relative velocity difference and relative temperature difference for an arbitrary set of parameters of the density method in a laminar flow. These difference plots are a qualitative representation of the accuracy of the density method for thermo-fluid TO and can be used for analyzing the results of a meta-optimization. A quantitative measure can also be constructed from this comparison between the two datasets, which will be introduced in Section 3.6.



(a) Relative velocity difference as percentage of the inlet velocity



(b) Relative temperature difference as percentage of the inlet temperature (in $^{\circ}\text{C}$)

Figure 3.7: Differences between the reference simulation and a simulation using the density method

Aside from simulation outputs that apply to the whole domain, some outputs can only be determined at the solid-fluid interface boundary (such as the heat transfer coefficient or wall shear stress). To perform a boundary analysis of a density-based simulation, a 'pseudo-boundary' is modeled, which is located at the solid-fluid interface boundary of the reference simulation (y -coordinate of 0). The blurred boundary of the density-based simulation does not necessarily align with the pseudo-boundary in a way that it represents actual boundary behavior accurately. The modeling parameter y_{grad} allows the blurred boundary to shift with respect to the pseudo-boundary to find the best match between the two. It is expected that the boundary behavior of the density-based model converges to that of the reference proportionally to the convergence of the global behavior with respect to the reference (e.g. the difference in velocity).

3.6. Defining a measure of modeling accuracy

To allow for quantitative analysis of the accuracy of the density method for thermo-fluid TO and to enable automatic tuning of the modeling parameters, a measure for the accuracy must be defined. No relevant measures have been found in literature and therefore one will be constructed in this thesis.

Selecting relevant simulation outputs. Since the accuracy analysis will be applied to a thermo-fluid simulation, both the thermal and the fluid dynamics aspects of accuracy need to be considered. This is mainly to avoid ‘loop-hole’ solutions where the thermal behavior may be accurate, but it is caused by erroneous fluid behavior. A selection is made of 2D simulation outputs that apply to the whole simulation domain, as these provide the highest fidelity when computing the modeling error (compared to boundary or point measurements). The following simulation outputs are selected:

- Velocity
- Temperature
- Turbulent viscosity (only applicable when a turbulent density-based fluid model is used)

Selecting an error definition. If the simulation outputs which have different dimensions are to be combined into one measure of modeling accuracy, a relative error should be computed of each of them. For this, the Root Relative Squared Error (RRSE) is selected. The RRSE is defined as Eq. (3.1):

$$E = \sqrt{\frac{\sum_{i,j=1}^{M,N} (P_{ij} - T_{ij})^2}{\sum_{i,j=1}^{M,N} (T_{ij} - \bar{T})^2}} \quad (3.1)$$

where E is the RRSE of a 2D simulation output, i and j are the position indexes in x and y directions, respectively; M and N are the number of samples in aforementioned directions, P_{ij} is a sample of the simulation output from the density method, T_{ij} is a sample of the reference simulation and \bar{T} is the average of all the reference samples. The RRSE normalizes the total squared error of the ‘predictor’ (the density-based data) with respect to the reference by dividing it by the total squared error of the simplest predictor possible, being the average of the reference values. The denominator is similar and proportional to the variance of the reference dataset. An RRSE of 100 % means that the density-based data approaches the reference data equally as accurate as the average of the reference data does. The RRSE has no upper bound. The RRSE reaches its lower bound of 0 % only when the density-based data matches the reference data exactly, because squared errors are used. This also ensures smoothness of the function, which is preferred for optimization. A downside of the RRSE is that it cannot provide information on over- or under-estimation of the density-based data. The RRSE is computed for all three aforementioned simulation outputs, and they are averaged to compute the ‘Average RRSE’.

3.7. Meta-optimization set-up

In Section 2.7 a task is created to tune the modeling parameters of the density-based thermo-fluid models. Given the complex behavior of the accuracy with respect to the modeling parameters (shortly analyzed in Appendix C), a manual tuning procedure would be cumbersome. Therefore an automated meta-optimization procedure is sought to perform this tuning. This section discusses the trade-off for the optimization algorithm used for the meta-optimization. After that, a short note is made on the objective function.

As with the optimization algorithms used for topology optimization (discussed in Section 2.3.6), the optimizer options for the meta-optimization can be categorized in two main ways: heuristic optimizers and mathematical optimizers. Within the latter category both gradient-based and gradient-free optimizers are considered. The trade-off between these three optimizer types is discussed in the following paragraphs.

- **Heuristic optimizers:** Because the physics involved with (turbulent) thermo-fluids are far too complex to determine a set of local criteria that are independent of their surroundings, these optimizer types are deemed infeasible.
- **Gradient-based mathematical optimizers:** Gradient-based optimization is theoretically more efficient than gradient-free optimization if the gradients are efficiently computed. However, with these optimizers every step in the process of the problem definition must be differentiable such that the gradient information can be computed using sensitivity analysis. Attempts to set-up these optimizer types in COMSOL with sensitivity analysis failed and using the finite difference method to compute gradient information resulted in lower efficiency due to the large amount of noise caused by the FEM approach. These optimizers are theoretically well-suited for the type of problem of the meta-optimization, with the downside that they are difficult to implement in their most efficient form. Moreover, these optimizers converge to the nearest local optimum and this may be far from the best solution for this problem.

- **Gradient-free mathematical optimizers** are less efficient than gradient-based optimizers but more easily implementable as they do not require the problem to be differentiable. They are more robust than gradient-based optimizers because of their lower sensitivity to noise in the model response. Several options in this category involve randomness that reduces the chance to end up in an inferior local optimum.

A comparison is made in a previous study between a gradient-based optimizer (SNOPT) which uses the finite difference approach to compute gradient information, and two gradient-free optimizers (BOBYQA and Nelder-mead). A simplified optimization problem is used to analyze the performance and stability of the optimization algorithms. From this analysis the Nelder-Mead optimizer is selected because it displayed the highest accuracy. It is assumed that the characteristics of the Nelder-Mead optimizer also apply to the meta-optimization problem.

Table 3.2 displays the modeling parameters along with their respective ranges and their definition. These parameters are considered as design variables for the meta-optimization procedure. These are all the parameters which define the density-based thermo-fluid model, the filtering of density variables and a variable specific to the meta-optimization. The penalty multipliers of eddy viscosity n_v and wall distance n_G are specific to the turbulent density-based fluid model and are not used for laminar meta-optimization. The vertical position of the density gradient y_{grad} is specific to the meta-optimization. The relevant ranges are based on what can be found in literature. Other optimizer-specific parameters such as convergence criteria or number of iterations are not considered as design variables and should be set manually. Table 3.2 also displays which parameters are converted to logarithmic scale due to their large (expected) ranges. In these cases $\log_{10}(x)$ is considered the design variable instead of x .

Table 3.2: Overview of the modeling parameters considered for the accuracy analysis of the density method. Relevant ranges are based on literature research.

Variable	Symbol	Definition	Relevant range	Linear/logarithmic parameter
SIMP interpolation penalty	p	Eq. (2.12)	2 - 5	linear
Darcy interpolation penalty	q	Eq. (2.13)	0.01 - 10	logarithmic
Length of solid-fluid transition	L_{trans}	Eq. (2.18)	-	linear
Projection filter slope	β	Eq. (2.19)	1-14	linear
Darcy number	Da	Eq. (2.30)	10^{-10} - 10^{-5}	logarithmic
Penalty multiplier eddy viscosity	n_v	Eq. (2.33)	-	logarithmic
Penalty multiplier wall distance	n_G	Eq. (2.34)	-	logarithmic
Vertical position of density gradient	y_{grad}	Section 3.5	-	linear

The objective function of the meta-optimization consists of the average of three RRSE errors, as defined in Eq. (3.1). The same three simulation outputs are used as for the accuracy analysis: velocity, temperature and turbulent viscosity. This objective function combination is a result of a number of experiments with different data types (1D, 2D) and simulation outputs (global variables, local variables, thermal outputs, physical outputs). The average RRSE of global variables is found to be stable and balanced without the need for manual tuning of weighting factors. Furthermore it is easily reproducible as opposed to custom objective functions with complex normalization methods.

To summarize, in this chapter the reference scenario (geometry and boundary conditions) of a developing turbulent boundary layer in a straight channel is introduced. Furthermore, the verification simulation of the adapted fluid models which is used to quantify the modeling accuracy is explained. Finally, the quantitative measure of modeling accuracy is selected being the average of RRSE values of the velocity, temperature and turbulent viscosity. These three steps combined form the framework for analyzing the accuracy of density-based TO and determining possible methods to increase it. Finally, a meta-optimization is set-up which automatically tunes all modeling parameters simultaneously by using a gradient-free (Nelder Mead) optimizer with the average RRSE as its objective. The trade-offs made to construct an efficient meta-optimization are discussed, of which the results will be discussed in the following chapters.


Detailed analysis of weaknesses in the density method

The framework for analyzing the accuracy of the density method introduced in Chapter 3 enables detailed qualitative and quantitative analysis of its weaknesses. In this chapter the effect of the two weaknesses introduced in Section 2.5, i.e. porous solid material and blurred walls, on the modeling accuracy will be analyzed both qualitatively and quantitatively using the accuracy analysis framework.


4.1. Porous solid material

The effects of solid material having a certain porosity are clearly visible in the density-based verification simulation (Section 3.5). In this scenario, the convective heat flow through the solid heat spreader is caused by a pressure gradient from the open boundary at the bottom of the heat spreader to the bulk of the fluid. Figure 4.1a displays this pressure gradient, which is in turn caused by the pressure losses in the channel flow. Since the outlet of the pipe is constrained at zero pressure, the centerline pressure is largest at the leading edge of the developing boundary layer, causing the largest pressure diffusion to occur near this leading edge. To eliminate convolution of this analysis with errors caused by a badly tuned fluid model, the modeling parameters have been optimized using the meta-optimization procedure.

When analyzing the convective heat flux passing through the open boundary at the bottom of the heat spreader, it is visible that a consistent outflow of heat occurs along the length of the channel. This is displayed in Fig. 4.1b along with the temperature distribution of the channel. The magnitude of the heat flux decreases gradually to 0 towards the outlet of the channel. This convective outflow of heat reduces the total heat flow supplied to the boundary layer by the heat source, which is specified as a conductive source.



(a) Pressure contours showing pressure diffusion through the solid heat-spreader region. The arrows indicate flow direction



(b) Temperature distribution along channel. Arrows indicate convective heat flux at open boundary

Figure 4.1: Anisotropic plots displaying the effect of the porous solid material when modeling with the density-based fluid models

A quantitative analysis of the total convective heat flowing through the bottom of the heat spreader (Table 4.1) shows that the effect strongly depends on the chosen Darcy number, with a convective outflow of heat ranging from 0.18 to 9.6 percent of the heat source inflow. For high Darcy numbers this becomes a significant error which is dominant in

decreasing the modeling accuracy. It is also clear that decreasing the Darcy number drastically decreases the convection through solid regions, enabling a much higher modeling accuracy.

Table 4.1: Analysis of (heat) flow through open bottom boundary of heat spreader. A positive sign corresponds to inward flow.

Darcy number	Convective heat flow [W]	Conductive heat flow [W]	Total heat flow [W]	Volume flow [%]
10^{-7}	-95.81	1000.00	904.19	-1.36 %
10^{-8}	-16.85	1000.00	983.15	-0.16 %
10^{-9}	-1.84	1000.00	998.16	-0.02 %

However, a downside of choosing extremely low Darcy numbers is that it causes an error in the fluid domain. The cause for this error will be discussed in Section 4.2. Solely using low Darcy numbers to minimize convection through solid regions is therefore not a feasible strategy, and other mitigation methods should be investigated.

4.2. Blurred walls

In this section the weakness of the density method having blurred walls between solid and fluid regions is analyzed in depth. As introduced briefly in Section 2.5, on a large scale the effect of these blurred walls is that they introduce an error in the heat transfer behavior between the solid and fluid regions. When zooming in on this effect, a chain of abnormalities is distinguished that combines into the cause of the erroneous heat transfer behavior. This section will go through each of these links individually. Firstly, the erroneous heat transfer is related to the velocity boundary layer thickness variation $\delta(x)$ and the thermal boundary layer thickness variation $\delta_T(x)$. Secondly, the profiles of the velocity and thermal boundary layer will be analyzed. Thirdly, the velocity profile is related to the profile of the Darcy force. Next, the interpolation functions that govern the shape of the Darcy force profile and the thermal material properties are investigated. Finally, the density variable profile is observed and the effect of various filtering settings is analyzed. From a sweep of modeling parameters, two cases are selected for visualization of the possible errors caused by a wrong selection of the modeling parameters. The two cases clearly show a physical and thermal error (see also Table 4.2), and their differences in terms of modeling parameter selection are chosen as simple as possible.

Case 1: a case which over-estimates the heat transfer from solid to fluid, resulting from a low Darcy interpolation penalty ($q = 0.01$) and a high Darcy number ($Da = 10^{-7}$).

Case 2: a case which under-estimates the heat transfer from solid to fluid, resulting from a high Darcy interpolation penalty ($q = 1$) and a low Darcy number ($Da = 10^{-9}$).

The other modeling parameters are constant, with values $L_{trans} = 1$ mm, $y_{grad} = -0.01 * L_{trans}$, $p = 3$, $n_G = 10^6$, $n_{\bar{v}} = 10^{-2}$ and $\beta = 8$. A turbulent scenario is analyzed with $Re = 6000$.

Table 4.2: RRSE values resulting from two visualization cases. Refer to Section 3.6 for the definition of RRSE

Case	RRSE Temperature	RRSE Velocity	RRSE Turbulent viscosity	Avg(RRSE)
Case 1	97.52 %	2.48 %	17.84 %	39.28 %
Case 2	192.96 %	82.01 %	243.62 %	172.86 %

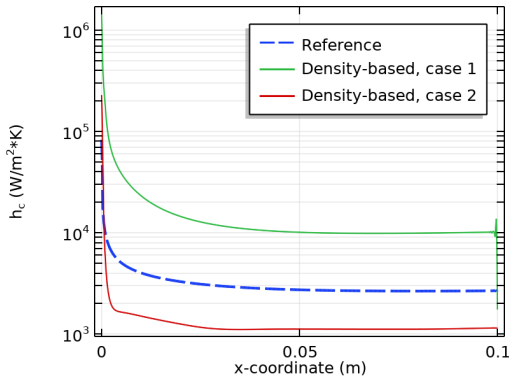
4.2.1. Erroneous heat transfer

This section will link the erroneous heat transfer behavior to the shape of the velocity and thermal boundary layers visible in a density-based simulation. Figure 4.2a shows how the local heat transfer coefficient along the channel of the two density-based cases compare to the reference simulation. It is clearly visible that an over-estimation of the heat transfer occurs for case 1 and an under-estimation occurs for case 2 (see also the high temperature RRSE values for both cases in Table 4.2).

The heat transfer behavior along the channel is closely related to the development of the thermal boundary layer. In a conventional fluid model, the thermal and velocity boundary layer thicknesses are related directly through the Prandtl number, which is a material property of the fluid. This means that the two boundary layers will develop with a similar shape (one being a scaled version of the other). Furthermore, for a given fluid a thin boundary layer provides a higher heat transfer coefficient than a thick one because a thin boundary layer shows a high velocity gradient, providing better convection close to the wall.

When comparing the thermal and velocity boundary layers to the reference simulation (see Figs. 4.2b and 4.2c, two main observations are made:

1. Case 1 shows boundary layer development comparable to the reference (see also the low velocity RRSE in Table 4.2), but there is a mismatch between the thermal and velocity boundary layers at the leading edge. The thermal boundary layer starts developing at a point downstream from the leading edge and develops faster than the reference. This is an indication that the Prandtl number is not constant at the blurred boundary. A varying Prandtl number can explain the increased heat transfer of case 1, but detailed analysis is required to explain this phenomenon.
2. Case 2 shows a significantly faster boundary layer development compared to the reference (see also the high velocity RRSE in Table 4.2). This thick boundary layer can explain the under-estimation of the heat transfer coefficient.



(a) Local heat transfer coefficients along channel



(b) Local thermal boundary layer thicknesses δ_T along channel(c) Local velocity boundary layer thicknesses δ along channel

Figure 4.2: Plots of heat transfer coefficient, velocity and thermal boundary layer thickness for the two density-based simulation cases compared to the reference simulation

4.2.2. Error in velocity and thermal boundary layer


Based on the observations made in Section 4.2.1, a more detailed view of the thermal and velocity boundary layers is required to further analyze the effects that cause the erroneous heat transfer. This section will analyze the profile of the density-based velocity and thermal boundary layers compared to the reference, and it will show how the mismatch between the velocity and thermal boundary layer causes an additional heat transfer error.

Figure 4.3a displays the temperature profile obtained at the outlet of the simulation domain. The solid-fluid boundary of the reference simulation is placed at $y = 0$ and the density gradient is centered around $y = -0.01$ mm, such that the figure displays the issue with case 2 very clearly. Starting at the heat spreader (negative y -coordinates), the temperature remains nearly constant up to an approximate y -coordinate of 0.1 mm. At this coordinate, the ΔT of the reference simulation has already decreased about 50%. The temperature of the heat spreader is significantly higher than the reference, because the heat transfer is lower. Interestingly, the thermal boundary layer is not much thicker than the reference at the outlet, and therefore the average temperature gradient is much greater than the reference.


The cause for the near-constant temperature and high temperature gradient for case 2 is found in the velocity profile (Fig. 4.3b). The fluid moves very slowly ($<1 \times 10^{-2} \text{ m s}^{-1}$) for a large part of the boundary layer thickness (up to the y -coordinate of 0.5 mm) and then accelerates quickly throughout the rest of the boundary layer. From hereon this is referred to as the ‘stagnant fluid effect’. At the y -coordinate of 1 mm and up, the velocity starts to overshoot the

reference velocity profile because the mass flux through the cross-section is constant for each simulation, as set by the inlet boundary condition. The stagnant fluid provides very little convective heat transfer, causing a part of the fluid in the boundary to heat up to the same temperature as the heat spreader. The high acceleration of the fluid causes increased convection at a large distance from the wall, resulting in the high temperature gradient. The profiles of the convective heat transfer for the two density-based cases and the reference simulation are visible in Fig. 4.3c. The stagnant fluid of case 2 is a result of the Darcy force acting on the fluid throughout the domain. The Darcy force of case 2 is likely too high at large distances from the wall. More detailed analysis is performed in Section 4.2.3 to explain the relation between the Darcy force and the velocity boundary layer.


In Figs. 4.3a and 4.3b it is observed that case 1 does not show much out of the ordinary, except that the temperature at the heat spreader is significantly lower than the reference. Figure 4.3c shows that the peak of convective heat transfer of case 1 is slightly shifted and wider compared to the reference. These minor differences still do not explain the major over-estimation of the heat transfer coefficient. When plotting the velocity profile on a logarithmic axis (Fig. 4.3d) and zooming in on the wall region, the cause of the additional heat transfer becomes apparent.




(a) Temperature profiles at outlet, plotted in ΔT values (temperature difference between the local fluid and the inlet temperature)



(b) Velocity profiles at outlet



(c) Convective heat transfer profiles at outlet



(d) Velocity profiles, plotted on a logarithmic vertical axis

Figure 4.3: Plots of various vertical cross-section profiles at the outlet of the simulation domain. Each plot displays a selected quantity at the cross-section for the two density-based cases and the reference simulation

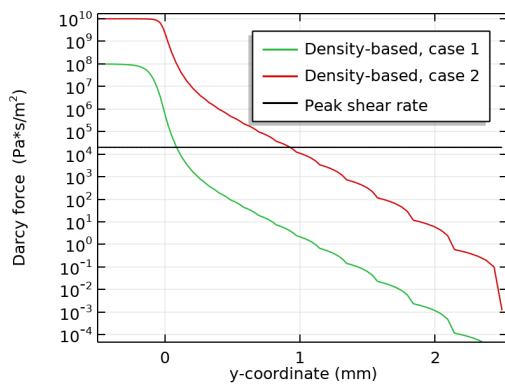
The velocity of case 1 follows the reference profile accurately from the top of the domain towards the wall down to a y -coordinate of approximately 0.05 mm. From this point towards and beyond the wall the density-based velocity decreases more slowly than the reference velocity, resulting in a significant velocity through part of the heat spreader. Likely the Darcy force of case 1 is not high enough in the heat spreader and in the solid-fluid density gradient to decelerate the fluid fast enough towards the wall. It is known that the interpolation of the Darcy force and the thermal material properties are ‘in transition’ in the region where the velocity is still relatively high, meaning that the thermal conductivity and thermal capacity are higher than the fluid’s values. In theory, a small layer of ‘liquid metal’ is present in case 1 at the region where the density gradient is specified. To provide more evidence for this theory, an analysis of the different interpolation functions is performed in Section 4.2.4.

4.2.3. Darcy force profiles


In Section 4.2.2 the observation is made that the Darcy force is too low to decelerate the fluid enough in case 1 and that in case 2 it is too high, causing the fluid in a large part of the boundary layer to remain stagnant. Further analysis on the influence of the Darcy force on the velocity profile is required, which is presented in this section. Figure 4.4a displays

the Darcy force profiles of the two density-based cases. A ‘peak shear rate’ threshold is plotted, which will be explained in the last paragraph of this section. Between case 1 and case 2, the Darcy force profiles differ in two ways: 1) the Darcy number is different for the two cases, resulting in a different maximum Darcy force in the heat spreader. 2) The Darcy interpolation penalty is different, resulting in a different curve throughout the density gradient. Ultimately, the two differences cause the same Darcy force to occur at different points in the boundary layer for the two cases. Another interesting observation is that the Darcy force does not reach 0 but rather shows an exponential decay. This is caused by the density variables which never reach 1 exactly in the fluid domain. This will be further analyzed in Section 4.2.5.


The main question to be answered in this section is: “Does the Darcy force profile directly govern the velocity profile throughout the boundary layer?” When the velocity is plotted with respect to the Darcy force (see Fig. 4.4b), this question can already be partially answered. Up until a Darcy force of approximately 10^5 the velocities are different, indicating that viscous friction is dominant in governing the velocity gradient in this region. At higher Darcy forces the velocities converge strongly, indicating that the Darcy force is dominant above a certain threshold. To determine this threshold, the behavior of the shear rate throughout the velocity boundary layer is analyzed. Figure 4.4c displays the shear rate profiles of the two cases compared to the reference. It is observed that the reference shear rate is monotonically decreasing from the wall up, converging to 0 at the boundary layer thickness. In contrast, the density-based shear rate profiles have a maximum and converge to 0 at the boundary layer thickness as well as within the heat spreader region. The peak of the shear rate is a good indication of the point where the Darcy force and the viscous forces are equal in governing the velocity gradient. When the shear rate of the density-based simulations is plotted with respect to the Darcy force (Fig. 4.4d), the peaks align relatively closely at the same Darcy force between 1.8×10^4 Pa*s/m² to 3.5×10^4 Pa*s/m² (9 cases were tested in total, all lie within the range spanned by case 1 and case 2). For convenience, the value of 2.5×10^4 Pa*s/m² is used for further analysis.




(a) Darcy force profiles



(b) Velocity profile at outlet, plotted with respect to Darcy force



(c) Shear rate profiles at outlet



(d) Shear rate profiles at outlet, plotted with respect to Darcy force

Figure 4.4: Plots of various vertical cross-section profiles at the outlet of the domain. Some plots display the relation between a selected profile and the Darcy force. All plots compare the two density-based cases and where applicable the density-based cases are compared to the reference result. A Darcy force threshold is determined from the peak of the shear rate in Fig. 4.4d

From these observations it can be concluded that a threshold Darcy force can be determined at which the peak shear rate occurs. For the turbulent channel scenario, this threshold lies between 1.8×10^4 Pa*s/m² to 3.5×10^4 Pa*s/m² and indicates the transition between viscous force-dominated and Darcy force-dominated velocity gradients. At Darcy forces above the Darcy force threshold the velocity profiles follow a predictable curve which is independent of the viscous effects in the boundary layer (Fig. 4.4b). *Simply said, this ‘peak shear rate’ Darcy force threshold indicates where the blurred*


boundary starts affecting the boundary layer. The threshold Darcy force can be placed at different distances from the wall by altering the Darcy number, Darcy interpolation penalty and density variable gradient (Fig. 4.4a). Observing the shear rate profile of the reference simulation, the threshold Darcy force is ideally placed as close as possible to the wall. An optimal combination can be found that results in a velocity profile which closely matches the reference (one can argue that case 1 is fairly close to this optimum). For case 2, the combination of a low Darcy number and high interpolation penalty (both deemed to be the ‘better value’ in literature) causes the threshold Darcy force to lie at a large distance from the wall (0.85 mm), slowing down any fluid closer to the wall than this distance.

4.2.4. Interpolation function profiles


In Section 4.2.2, a layer of ‘liquid metal’ is claimed to be present in case 1, which is caused by a mismatch between the two interpolation functions used to interpolate the Darcy force (Darcy interpolation) and the thermal material properties (SIMP interpolation). This section analyzes this mismatch in more detail and relates it to the density variables.

In Fig. 4.5a, the relation between the Darcy force and the interpolated thermal conductivity is displayed for three values of the Darcy number at a fixed Darcy interpolation penalty of 0.01, which is the same as case 1. Furthermore, the plot displays the Darcy force threshold above which the velocity of all cases has decreased below 1 mm s^{-1} . This threshold is used for a rough indication of the y-coordinate where the convective heat transfer and the conductive heat transfer have equal magnitude. It is observed that the position where the velocity threshold is crossed varies with the Darcy number. For the highest Darcy number, which is the same as case 1, the Darcy force crosses this threshold at an approximate y-coordinate of -0.06 mm . At this coordinate, the thermal conductivity is already at 99.99 % of its maximum value (400 W/(mK)) and 666 times higher than the fluid’s thermal conductivity (0.6 W/(mK)), meaning that the convective heat transfer is occurring at a region of very high thermal conductivity. This observation strongly supports the theory of the ‘liquid metal’ layer.


Decreasing the Darcy number shifts the Darcy force profile upwards such that the aforementioned Darcy force threshold is crossed at approximately 0.04 mm . Now the convective heat transfer occurs at a maximum thermal conductivity which is 162 times higher than the fluid’s thermal conductivity. This is still higher than desired, resulting in an over-estimation of the heat transfer coefficient. Other methods to decrease the mismatch between the Darcy force and the interpolated thermal material properties are to increase the Darcy interpolation penalty q (Fig. 4.5b) or to decrease the SIMP interpolation penalty p (Fig. 4.5c).



(a) Left axis: Darcy force profile plotted in logarithmic scale, right axis: interpolated thermal conductivity plotted in linear scale



(b) Darcy interpolation function profiles with varying interpolation penalty q , compared to density variable gradient



(c) SIMP interpolation function profiles with varying interpolation penalty p , compared to density variable gradient

Figure 4.5: Plots that visualize the interpolated variables and the different interpolation functions at a vertical cross-section of the domain

When $q \approx 10$ and $p = 1$ the interpolated variables directly follow the density variables. It is unclear whether altering the density variables to generate a steeper gradient has a positive effect on minimizing the mismatch between the interpolation functions. A steeper gradient moves the point where the thermal conductivity of the fluid starts increasing closer to the point where the Darcy force is high enough to minimize convective heat transfer but the (unpredictable) viscous effects in the velocity boundary layer will determine whether this reduces the mismatch between the interpolated variables.


To conclude, the same factors which improve the accuracy of the velocity boundary layer of case 1 (decrease Darcy interpolation penalty, increase Darcy number), also cause the inaccuracy of its thermal boundary layer. A combination of the Darcy number, Darcy interpolation penalty, SIMP interpolation penalty and (possibly) the density variables can be found that places the interpolated variables correctly with respect to the wall and therefore eliminates the ‘liquid metal’ effect. Careful consideration must be given to balance the ‘liquid metal’ effect with the ‘stagnant fluid’ effect, as likely neither can be completely eliminated.

4.2.5. Density variable gradient


In Section 4.2.3 the density variables are mentioned as one of the options to place the cross-over coordinate of the Darcy force threshold at which the shear rate peaks. Also, Section 4.2.4 states that a steeper density variable gradient might reduce the mismatch between the interpolated variables. Theoretically, a shorter transition will always provide more accurate results because a steep density gradient represents a sharp ‘real’ wall more accurately than a shallow one. The relation between the density gradient and the Darcy force and the interpolated thermal material properties is analyzed in the previous sections; therefore for brevity the interpolated variables are visualized directly instead of first visualizing the density gradient.

A sweep with three transition lengths L_{trans} and four projection filtering slopes β is performed to analyze how these parameters alter the shape of the interpolated variables. A projection slope of 0.1 is used to effectively disable the projection filter, as the projection function is nearly linear for this slope. The Darcy number and Darcy interpolation penalty are set to the values of case 2 ($Da = 10^{-9}$, $q = 1$) to focus on the error causing stagnant fluid near the wall.

In Fig. 4.6a, the effect of changing the solid-fluid transition length L_{trans} (and with that the blurring filter radius) is displayed in terms of the obtained Darcy force profile. For these cases, the projection filter is effectively disabled. The two Darcy force thresholds (peak shear rate, velocity $< 1 \text{ mm s}^{-1}$) are also plotted in black, giving a prediction of the shape and position of the velocity boundary layer with respect to the wall. Reducing the transition length is quite effective in bringing the coordinates at which the Darcy force thresholds are crossed closer to the wall. However, to bring these cross-over coordinates to within a reasonable distance from the wall by only reducing the transition length would be practically infeasible, because the mesh size required to enable such short transition lengths is too extreme. By using projection filtering (Fig. 4.6b) with a reasonable value of L_{trans} , the cross-over coordinates are shifted closer to the wall without the computational burden of decreasing the mesh resolution. This is a strong indication that such a filter can be used to improve accuracy, alongside its known ‘contrast-enhancing’ effect on a TO.



(a) Darcy force profiles for varying solid-fluid transition lengths L_{trans} . Projection filtering is effectively disabled in these cases by selecting $\beta = 0.1$




(b) Darcy force profiles for varying projection filter slopes β . L_{trans} is selected to be the same as case 1 and case 2 in previous sections


Figure 4.6: Plots displaying the effect of varying solid-fluid transition lengths and projection filter slopes on the Darcy force profiles. The Darcy force thresholds where the fluid velocity drops below 1 mm s^{-1} and where the shear rate peaks are displayed as black solid and black dashed lines, respectively

The interpolation of the thermal material properties (in this case the thermal conductivity) shows a similar dependence on both the transition length (Fig. 4.7a) and the projection filter slope (Fig. 4.7b). In both cases the largest change is visible on the low end of the interpolated thermal conductivity. This is beneficial for the accuracy of the thermal boundary layer, as the y-coordinate where the fluid starts showing increased heat transfer is moved significantly closer to the wall with decreasing transition length or projection slope. As explained in Section 4.2.4, it is unclear whether a steeper density gradient will inherently decrease the mismatch between the Darcy force and the interpolated thermal material properties.

Because a blurring filter limits the minimum feature size of geometries in a topology optimization relative to its filter radius, a value for L_{trans} can be selected based on the required minimum feature size of the optimized design and the mesh size constraints. The projection filter slope can then be tuned to generate a sufficiently steep gradient.



(a) Interpolated thermal conductivity profiles for varying solid-fluid transition lengths L_{trans} . Projection filtering is effectively disabled in these cases by selecting $\beta = 0.1$



(b) Interpolated thermal conductivity profiles for varying projection filter slopes β . L_{trans} is selected to be the same as case 1 and case 2 in previous sections

Figure 4.7: Plots displaying the effect of varying solid-fluid transition lengths and projection filter slopes on the interpolated thermal conductivity profiles

To summarize, this chapter discussed the two main weaknesses of turbulent TO using the density method. Section 4.1 analyzes the error caused by porous solid material. It concludes that the convection through solid material can be minimized by decreasing the Darcy number at a cost of introducing different errors. It is suggested that other mitigation methods should be investigated to minimize the errors caused by porous solid material.

Section 4.2 provides an extensive analysis of the errors caused by blurred walls. Following the ‘chain of abnormalities’, starting at an over- or under-estimated heat transfer and ending at the parameters governing the density variable gradient, the effect of each chain-link on the total effect is discussed. It is concluded that the ‘stagnant fluid’ effect causes under-estimation of the heat transfer and the ‘liquid metal’ effect causes over-estimation of the heat transfer. Both effects are caused and remedied by certain combinations of the modeling parameters. Eliminating one effect can cause the other effect to appear. Therefore it is suggested that an optimal combination of the modeling parameters can be found which balances the errors caused by the two effects. This shows the benefit of using Meta-optimization, as the number of modeling parameters is large and the balance between them is too intricate to fine-tune manually. Furthermore, a certain predictability is found in the behavior of boundary layers with various modeling parameter settings, being that they all show a peak in shear stress at the same ‘Darcy force threshold’ and show relatively similar behavior in regions with Darcy forces above this threshold. This threshold forms a good measure of how the blurred density-based wall can be converted back to a crisp wall, and therefore can be used to post-process the density variables of a TO. Finally, the RRSE data obtained by the analysis framework shows good correlation with the qualitative analysis and is deemed functional for further quantitative analysis of the density-based fluid models.

5

Mitigation methods for weaknesses in the density method

In this chapter the mitigation methods for turbulent flow, introduced in the research questions (Section 2.6), are further detailed and their effectiveness is analyzed based on the modeling accuracy. For an analysis of the mitigation methods applied to a laminar flow, refer to Appendix B. In Section 5.1 the method of tuning the modeling parameters to reduce errors caused by blurred walls is discussed, and Section 5.2 discusses the method of altering the thermal conductivity to decrease the effects caused by porous solid material.

5.1. Mitigating blurred walls by tuning modeling parameters

As concluded in Section 4.2, the modeling parameters have a large effect on the boundary layer behavior of a turbulent density-based simulation and it is noted that a balance can be found between the ‘stagnant fluid’ and ‘liquid metal’ effects by tuning the modeling parameters. In Appendix C a ‘black-box’ analysis of the influence of the modeling parameters on the average RRSE is performed, and local optima were observed.

Because the optimization of modeling parameters to decrease the average RRSE seems feasible, a meta-optimization strategy is created in Section 3.7 to perform this optimization. Section 5.1.1 will conceptualize a mitigation method based on a meta-optimization of the turbulent density-based fluid model with selected modeling parameters. This mitigation method will be referred to as the ‘turbulent blurred walls mitigation method’ and abbreviated to ‘Mitigation 1’. After this, the results of this mitigation method will be discussed in Section 5.1.2.

5.1.1. Turbulent optimization problem

The following modeling parameters are considered as design variables for the meta-optimization strategy when applying mitigation 1 (refer to Table 3.2 for their definition and range):

- SIMP interpolation penalty p
- Darcy interpolation penalty q
- Projection filter slope β
- Penalty multiplier eddy viscosity n_v
- Penalty multiplier wall distance n_G
- Vertical position of density gradient y_{grad} .

The Darcy number is not considered as a design variable because from literature it is observed to show monotonically increasing influence on the accuracy with decreasing Darcy numbers and this observation is confirmed in Appendix C. Including a monotonic parameter in the meta-optimization would force it to continue infinitely, as the Darcy number would tend to 0 (in this case to $10^{-\infty}$ because $\log_{10}(\text{Da})$ would be used as the design variable). The solid-fluid transition length L_{trans} is also not considered for this same reason and is fixed at 20 % of the channel height ($L_{trans} = 1 \text{ mm}$).

The objective for this optimization is the average RRSE of the global velocity, global temperature and global turbulent viscosity (as discussed in Section 3.6). A sweep is performed for the Darcy number values of 10^{-7} , 10^{-8} and 10^{-9} , optimizing the above modeling parameters for each Darcy number.

5.1.2. Turbulent results

The results obtained from mitigation 1 are discussed here. The results are displayed graphically based on relevant simulation outputs in Figs. 5.1a to 5.1d. The two cases used in Section 4.2 are also displayed. For clarity, their distinguishing features are:

- Case 1: a low Darcy interpolation penalty ($q = 0.01$) and a high Darcy number ($Da = 10^{-7}$), resulting in an over-estimation of the heat transfer from solid to fluid.
- Case 2: a high Darcy interpolation penalty ($q = 1$) and a low Darcy number ($Da = 10^{-9}$), resulting in an under-estimation of the heat transfer from solid to fluid.

The results of mitigation 1 are obtained after on average 200 iterations of the meta-optimizer with an optimality tolerance of 0.1 (meaning that a stationary point is reached if “no improvement over the current best estimate can be found with steps in the scaled control variables of relative size larger than or equal to the optimality tolerance” [12]).

Comparison of graphical accuracy: In Figs. 5.1a and 5.1b it is observed that the optimized results follow the reference much closer than the two cases, although an increased pseudo heat transfer coefficient at the leading edge of the channel causes reduced pseudo wall temperature differences in the first half of the channel. (Pseudo boundary results are obtained following methodology discussed in Section 3.5). A clear difference is visible between the optimized result with $Da = 10^{-7}$ and the results with $Da = 10^{-8}$ and $Da = 10^{-9}$, but the latter two are nearly indistinguishable.

From Fig. 5.1c one can observe that the optimized results follow the behavior of case 1 very closely but show significant error with respect to the reference. This either indicates that there is no combination of modeling parameters that can make the pseudo boundary of the density-based model resemble the reference, or it indicates that doing so is possible but not advantageous for the overall accuracy of the model. Figure 5.1d confirms that the boundary layer resembles that of case 1. However as the Darcy number decreases, so does the velocity through the solid material.

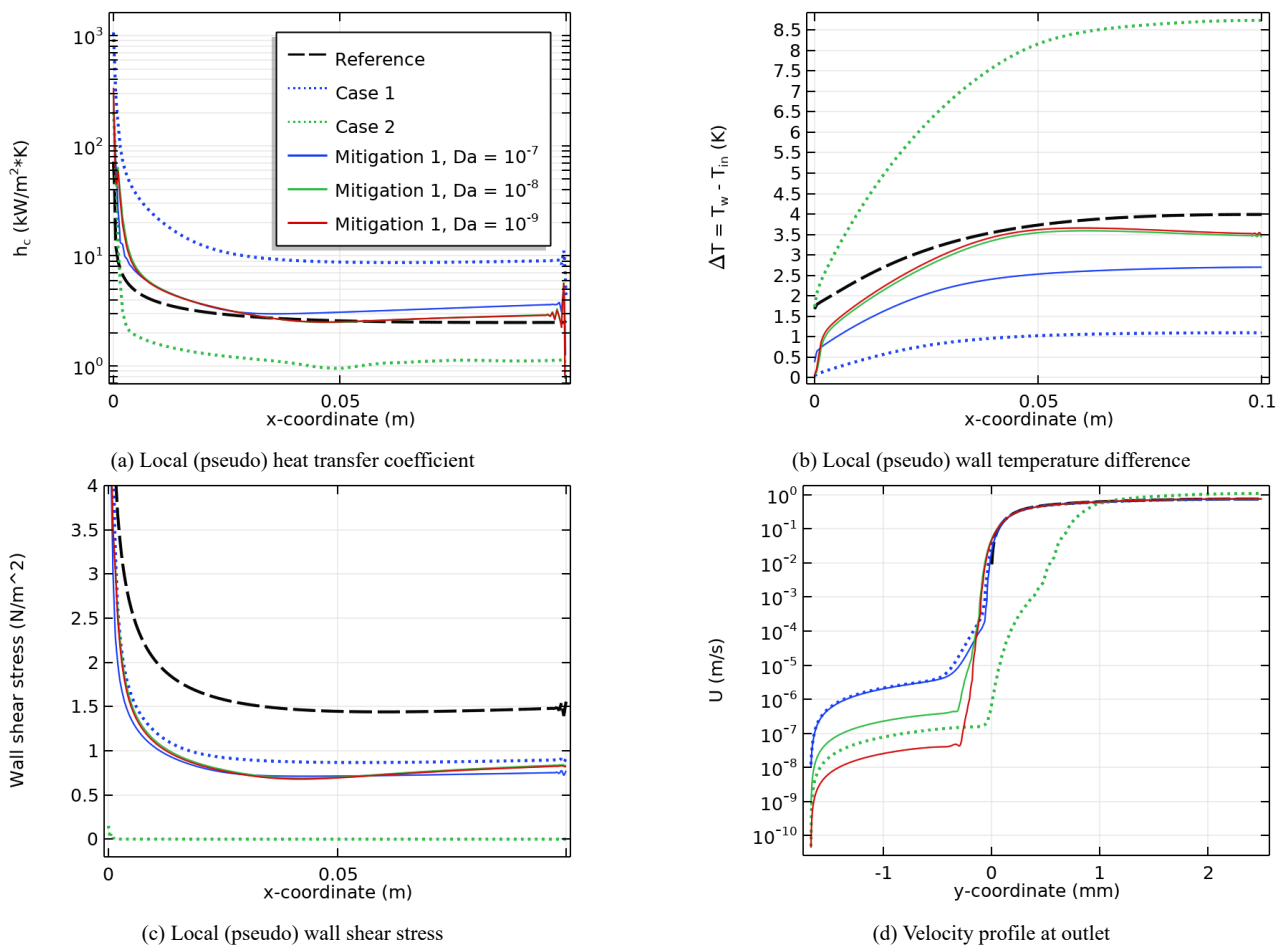


Figure 5.1: Comparison of mitigation 1 for three Darcy number values to the reference and the two cases which are used in Section 4.2. Pseudo boundary results are obtained following methodology discussed in Section 3.5

Comparison of numerical accuracy: The quantitative data, displayed in Table 5.1, confirm the trends described above. Compared to the two cases, the optimized results have significantly lower average RRSE values (a minimum improvement of 20 percentage point and a maximum improvement of 163 percentage point) and decreasing errors are observed with decreasing Darcy numbers. Between the three optimized results the lowest two Darcy number results are nearly indistinguishable. Interestingly, the velocity and turbulent viscosity RRSE values *increase* going from $Da = 10^{-7}$ to 10^{-8} although the temperature RRSE decreases significantly. This observation is consistent with the observations made in Sections 4.1 and 4.2 that the convection through porous solid material and the ‘liquid metal’ effect are both reduced by decreasing the Darcy number (both predominantly cause thermal errors). In contrast to this, the ‘stagnant fluid’ effect (which causes both thermal errors and velocity errors) is increased by lowering the Darcy number. This indicates that for high Darcy numbers, the optimization struggles to mitigate the convection through porous solid material and the ‘liquid metal’ effects by only tuning the set of modeling parameters introduced in Section 5.1.1. At lower Darcy numbers, the physical error caused by the ‘stagnant fluid’ effect dominates. The strong shift in balance between the two effects caused by the change in Darcy number indicates that there may be a Darcy number value between 10^{-7} and 10^{-8} which provides a better balance between these two effects. This could decrease the average RRSE further than either of the three Darcy numbers used for this analysis.

Table 5.1: Comparison of individual and average RRSE values of mitigation 1. The modeling parameter optimization is performed for three Darcy number values. The two cases used in Section 4.2 are also displayed

Case	Darcy number	RRSE Temperature	RRSE Velocity	RRSE Turbulent viscosity	Avg(RRSE)
Case 1	-	97.52 %	2.48 %	17.84 %	39.28 %
Case 2	-	192.96 %	82.01 %	243.62 %	172.86 %
Mitigation 1	10^{-7}	44.25 %	1.80 %	10.70 %	18.92 %
	10^{-8}	7.16 %	3.68 %	20.32 %	10.39 %
	10^{-9}	7.03 %	3.57 %	19.95 %	10.18 %

Resulting modeling parameters: The resulting modeling parameters of mitigation 1 (Table 5.2) confirm the observations that the ‘liquid metal’ effect dominates at a Darcy number of 10^{-7} and above, and the ‘stagnant fluid’ effect dominates at Darcy numbers of 10^{-8} , 10^{-9} and below. This is indicated by the decreasing trend (visible in Table 5.2) of the Darcy interpolation penalty q and the increasing trend of the SIMP interpolation penalty p when the Darcy number is decreased. At a Darcy number of 10^{-7} the Darcy force is low and the Darcy force thresholds are crossed close to the wall (Figs. 4.4a and 4.5a), giving rise to the ‘liquid metal’ effect. To minimize this, the meta-optimization increases the Darcy interpolation penalty such that the Darcy force thresholds lie further from the wall, reducing this effect. Increasing the Darcy interpolation penalty also brings the Darcy interpolation curve closer to that of the SIMP interpolation curve (Figs. 4.5b and 4.5c), which is strengthened by the low SIMP interpolation penalty. When the two interpolation curves match better, the ‘liquid metal’ effect is further minimized.

The result of the meta-optimization with Darcy number of 10^{-9} clearly shows how the ‘stagnant fluid’ effect is caused by the Darcy number and mitigated by the interpolation penalties. The lower Darcy number increases the Darcy force throughout the channel, which shifts the boundary layer away from the pseudo-boundary (Fig. 4.4a). The optimizer compensates for this by decreasing the Darcy interpolation penalty significantly. The resulting optimized boundary layer is placed such that the ‘liquid metal’ effect does not occur and the large mismatch between the interpolation functions (caused by a low Darcy interpolation penalty and high SIMP interpolation penalty, see Figs. 4.5b and 4.5c) does not affect the accuracy. The intermediate result with Darcy number of 10^{-8} shows a balance between the two effects.

The modeling parameters y_{grad} , $n_{\bar{v}}$, n_G and β do not show strong trends (in Table 5.2), and n_G as well as β remain nearly constant. The wall distance penalty multiplier n_G converges to the upper bound of 10^6 for each result, indicating that either the upper bound is placed too low to find an optimum, or that this modeling parameter has a monotonic effect on the accuracy. Strangely, the turbulent viscosity penalty multiplier $n_{\bar{v}}$ converges to a reasonably close range for each result although the RRSE of the turbulent viscosity is dominant at lower Darcy numbers. Either the scaling of the turbulent viscosity penalty with the Darcy number is perfect (the turbulent viscosity penalty is a product of the Darcy force and the turbulent viscosity penalty multiplier) which causes the meta-optimization to converge to the same multiplier, or the effect of the multiplier shows a very shallow gradient which causes the design variable to slowly move towards an optimum or one of the user-selected bounds. If the other modeling parameters converge to an optimum while the turbulent viscosity penalty multiplier only migrates slowly, the solution is still considered to be converged. The projection slope β is expected to have a monotonic effect on the accuracy, but the meta-optimization results cannot confirm this. Either a local optimum is found for β or the sensitivity to it is too low for proper convergence.

Finally, a conclusion can only be made on the vertical gradient position y_{grad} after analyzing the Darcy force profiles.

Table 5.2: Comparison of modeling parameter values resulting from the meta-optimization of mitigation 1 for three Darcy numbers

Darcy number	y_{grad}/L_{trans}	q	p	$n_{\bar{v}}$	n_G	β
10^{-7}	-5.97 %	4.05	2.49	3.39×10^{-2}	9.99×10^5	12.97
10^{-8}	-15.82 %	0.52	2.65	1.97×10^{-2}	1.00×10^6	11.61
10^{-9}	-14.98 %	0.05	3.22	2.09×10^{-2}	9.87×10^5	12.04

The combination of the interpolation penalty and gradient position results in a certain Darcy force occurring at the pseudo-boundary. This Darcy force can be used as a threshold value for post-processing results of a TO. Figure 5.2a displays that the optimized Darcy force profiles align closely at the Darcy force at which peak shear rate occurs. This results in the shear rate of the viscous region of the boundary layer aligning accurately with the shear rate of the reference (Fig. 5.2b). The vertical gradient position y_{grad} is used mainly to provide this alignment.

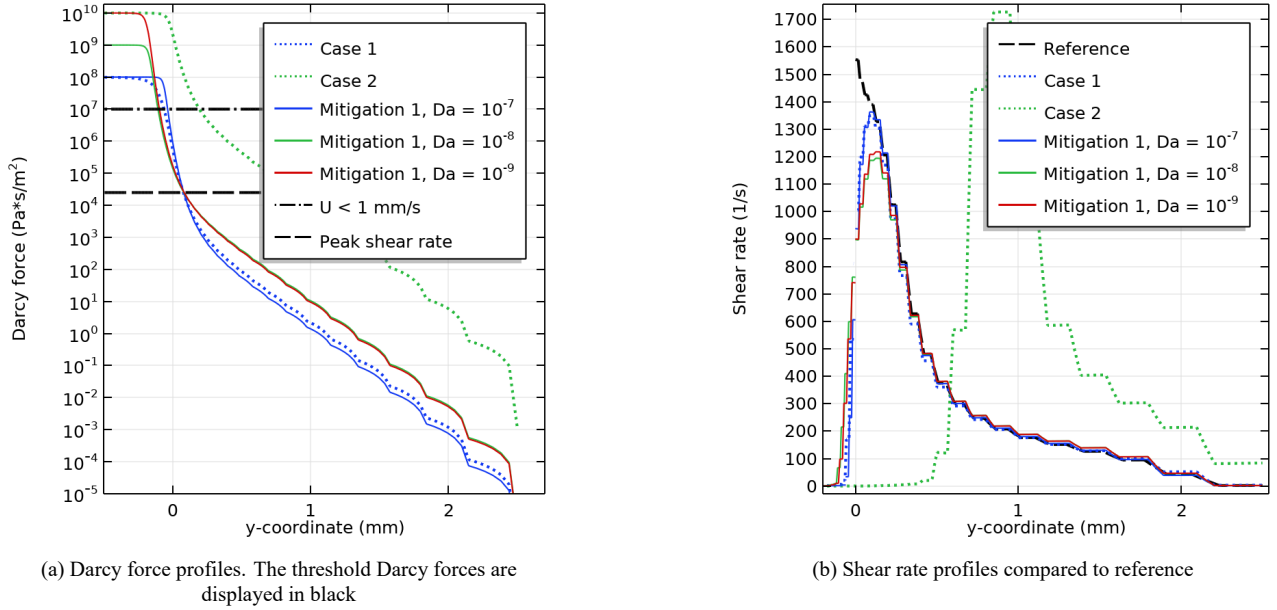


Figure 5.2: Comparison of mitigation 1, meta-optimized for three Darcy numbers. The two cases used in Section 4.2 are also displayed

5.2. Mitigating porosity of fully solid material by altering thermal conductivity

In a density-based fluid model even a fully solid material (defined as the filtered density variables approaching 0) has a finite porosity, caused by a finite maximum Darcy force. Aside from minimizing this porosity (Section 4.1) another method is sought to minimize the thermal errors. Conceptually, the error introduced by this porosity on the heat transfer behavior through the solid material can be reduced by altering the thermal conductivity of the solid material. In a true solid the heat flux q_d is purely conductive (Fig. 5.3a) and the temperature difference $T_1 - T_2$ over a finite length L is governed by Fourier's law (Eq. (5.1)):

$$T_1 - T_2 = \frac{q_d}{k_s} L. \quad (5.1)$$

The porous solid material of the density-based fluid models can accommodate a heat flux consisting of both a conductive and a convective component q_c (Fig. 5.3b). When the two heat flux components are summed and it is assumed that this combined total heat flux (q_t) is also governed by Fourier's law according to Eq. (5.2):

$$T_1 - T_3 = \frac{q_t}{k_s} L = \frac{q_d - q_c}{k_s} L, \quad (5.2)$$

the temperature difference $T_1 - T_3$ of this porous solid material is different than that of a true solid. If the convective heat flux opposes the conductive heat flux, as displayed, the temperature difference will be lower than that of a true solid. The opposite is true if the convective heat flux is in the same direction as the conductive heat flux.

Because the porosity of the fully solid material occurs at sub-grid level, heat transfer through the porous solid material is not modeled as a fluid flowing through solid material but rather like a single solid material which is in motion. Because of this single-material representation there is only one conductive heat flux instead of one for each material in a real porous

solid. Furthermore, this single conductive heat flux is carried by material with the thermal conductivity of a true solid, whereas in a real porous solid the two individual conductive heat fluxes would be carried by their respective material with individual thermal conductivity values. Therefore the single-material representation of a fully solid material with finite porosity shows behavior which is different than that of a true solid (the presence of a convective heat flux) as well as that of a true porous material (the lack of a conductive heat flux of the fluid and its respective thermal conductivity).

By altering the isotropic thermal conductivity of the porous fully solid material $k_s - \Delta k$ (Fig. 5.3c), the temperature difference $T_1 - T_4$ can be adjusted according to Eq. (5.3):

$$T_1 - T_4 = \frac{q_d - q_c}{k_s - \Delta k} L. \quad (5.3)$$

By equating the temperature difference of the true solid to that of the porous solid material with adjusted conductivity, the conductivity adjustment Δk can be computed with Eq. (5.4):

$$T_1 - T_2 = T_1 - T_4 \Rightarrow \Delta k = -\frac{q_c}{q_d} k_s. \quad (5.4)$$

This shows that a negative adjustment should be applied to the conductivity when the convective heat flux opposes the conductive heat flux. Because the convective heat flux is varying in direction and magnitude locally and the thermal conductivity is only adjusted globally, it is unlikely that this mitigation method solves the problem locally. However, a global improvement might already provide sufficiently increased accuracy. Furthermore, the assumption that the total heat flux is governed by Fourier's law is a simplification of reality, the single-material model of a porous solid departs from that of a true porous material and the convective heat flux is dependent on the local temperature of the porous medium (heat flux is proportional to enthalpy, which depends on temperature). These effects could alter the convective heat flux as the thermal conductivity is adjusted.

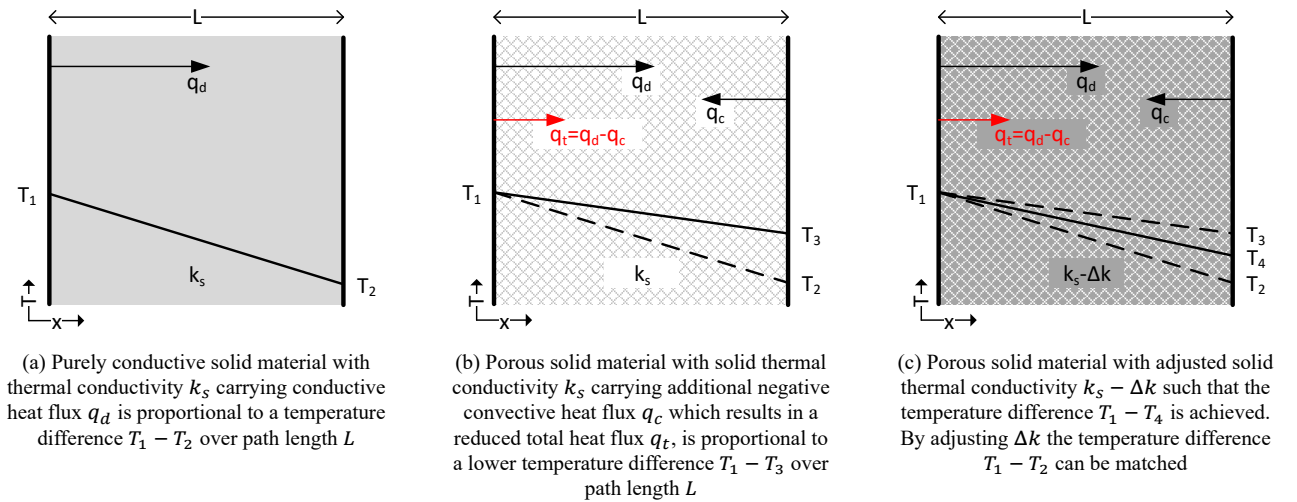


Figure 5.3: Three illustrations displaying the heat flux and temperature profiles through different material types

In the remainder of this section, the optimization problem for a turbulent density-based fluid model is introduced first (Section 5.2.1) after which the results of this optimization are discussed in Section 5.2.2.

5.2.1. Turbulent optimization problem

In addition to the design variables given in Section 5.1.1, the isotropic thermal conductivity of the solid material k_s is included as a design variable for this mitigation method. By including this modeling parameter, the behavior of a porous solid material can possibly be matched better with a true solid material. The same objective function is used as mitigation 1. Furthermore, the upper limit of the wall distance penalty multiplier is raised from 10^6 to 10^8 to eliminate the effect of reaching its upper bound which is visible for mitigation 1. The meta-optimization of a turbulent density-based fluid model with the modeling parameters given above will be referred to as the ‘turbulent porous solid material mitigation method’ and abbreviated to ‘Mitigation 2’.

5.2.2. Turbulent results

The results of mitigation 2 are discussed in this section. Visually, small differences are visible between the results of mitigation 1 and mitigation 2 but no clear increased accuracy is observed. A comparison of graphical accuracy is therefore omitted for mitigation 2.

Comparison of numerical accuracy: Looking at the quantitative data (Table 5.3), a counter-intuitive trend is visible: the result with the highest Darcy number (10^{-7}) shows the lowest error, and the error *increases with decreasing Darcy numbers*. Especially the turbulent viscosity RRSE of the highest Darcy number result is much lower than that of all other results. The turbulent viscosity RRSE increases strongly with decreasing Darcy numbers, while the other RRSE values remain relatively stationary for all Darcy numbers (a slight but insignificant increase of the velocity RRSE is visible with decreasing Darcy numbers). Only the mitigation 2 result of Darcy number 10^{-7} shows a lower average RRSE than mitigation 1. Except for the 10^{-7} Darcy number result, the meta-optimization does not succeed in decreasing the temperature RRSE of mitigation 2 compared to mitigation 1. Because the Nelder-Mead optimizer which is used for the meta-optimization does not contain any features to ensure global convergence, it is likely that the Meta-optimization has converged to a local minimum which performs worse at lower Darcy numbers. Ideally the convergence to an inferior local minimum is avoided by starting the optimization from multiple initial configurations. However, due to the high computational cost of a turbulent meta-optimization this approach is not feasible for this thesis. A more efficient approach might have been to start the meta-optimization for mitigation 2 from the optimum found by mitigation 1, but this did not fit in the time-frame of the research either because the two mitigation methods were developed in parallel.

Table 5.3: Comparison of individual and average RRSE values of the turbulent porous solid material mitigation method. The modeling parameter optimization is performed for three Darcy number values. The optimized results of mitigation 1 are also displayed

Mitigation method	Darcy number	RRSE Temperature	RRSE Velocity	RRSE Turbulent viscosity	Avg(RRSE)
Mitigation 1	10^{-7}	44.25 %	1.80 %	10.70 %	18.92 %
	10^{-8}	7.16 %	3.68 %	20.32 %	10.39 %
	10^{-9}	7.03 %	3.57 %	19.95 %	10.18 %
Mitigation 2	10^{-7}	15.48 %	1.78 %	1.77 %	6.34 %
	10^{-8}	12.55 %	1.86 %	20.68 %	11.70 %
	10^{-9}	15.02 %	3.61 %	37.84 %	18.82 %

Resulting modeling parameters: The resulting modeling parameters of the meta-optimizations performed for mitigation 2 are compared to mitigation 1 in Table 5.4. Since the RRSE values indicate that mitigation 2 has converged to an inferior optimum, only the most notable differences are discussed.

The wall distance penalty multiplier n_G shows a decreasing trend with decreasing Darcy numbers for mitigation 2, in contrast to mitigation 1 where the values consistently approached the upper bound (10^6). In Section 5.1.2 the conclusion was made that the wall distance penalty likely shows monotonic behavior, but this is not the case for mitigation 2. The cause for the decreasing trend of the wall distance penalty multiplier is unclear.

Where the thermal conductivity of the solid material is fixed for mitigation 1, it is a design variable for mitigation 2. An increasing trend is visible with decreasing Darcy numbers, starting at an extremely low conductivity at the result with Darcy number 10^{-7} and increasing to a value that is still significantly lower than the reference at a Darcy number of 10^{-9} . The effect of lowering the thermal conductivity to compensate for the convective heat flow through the solid material which is opposing the conductive heat flow is consistent with the concept of the mitigation method. The increasing trend of the thermal conductivity is related to the decreasing porosity of the solid material at lower Darcy numbers which lowers the convective heat flow through the solid material. Therefore less compensation of the thermal conductivity is needed at lower Darcy numbers.

Table 5.4: Comparison of modeling parameter values resulting from the meta-optimization of the turbulent porous solid material mitigation method for three Darcy numbers

Mitigation method	Darcy number	y_{grad}/L_{trans}	q	p	n_v	n_G	β	k_s [W/(m*K)]
Mitigation 1	10^{-7}	-5.97 %	4.05	2.49	3.39×10^{-2}	9.99×10^5	12.97	400
	10^{-8}	-15.82 %	0.52	2.65	1.97×10^{-2}	1.00×10^6	11.61	400
	10^{-9}	-14.98 %	0.05	3.22	2.09×10^{-2}	9.87×10^5	12.04	400
Mitigation 2	10^{-7}	-4.84 %	2.96	3.47	5.49×10^{-3}	2.15×10^6	13.32	15.42
	10^{-8}	-7.03 %	0.06	2.78	5.53×10^{-2}	3.41×10^5	12.18	46.20
	10^{-9}	-12.51 %	0.01	3.09	4.27×10^{-2}	2.71×10^4	11.00	156.76

Bottom boundary analysis: The same analysis of the heat flow and volume flow through the open bottom boundary is performed for the turbulent porous solid material mitigation method which was first introduced in Table 4.1. In Table 5.5 the results of both mitigation methods are displayed for comparison. Apart from the Darcy number result of 10^{-7} , the results of the two mitigation methods are extremely similar. The result of mitigation 2 with Darcy number of 10^{-7} shows a significantly *higher* convective outflow of heat through the bottom boundary compared to mitigation 1, although the volume outflow is exactly the same. This is a contradictory effect compared to the low RRSE observed with this result (Table 5.3), indicating that the extremely low resulting thermal conductivity of the solid material triggers the issues discussed in Section 5.2. Specifically the increased temperature in the solid material by reducing its thermal conductivity may have caused the convective heat flux to also increase. By reducing the range of the thermal conductivity for mitigation 2, the unwanted side-effects of extreme compensation could likely have been minimized.

Table 5.5: Heat- and volume flow through the open bottom boundary for the turbulent porous solid material mitigation method. The optimized results of mitigation 1 are also displayed. A positive sign represents flow into the domain. The volume flow is relative to the total flow rate

Mitigation method	Darcy number	Convective heat flow [W]	Conductive heat flow [W]	Total heat flow [W]	Volume flow [%]
Mitigation 1	10^{-7}	-95.81	1000.00	904.19	-1.36 %
	10^{-8}	-16.85	1000.00	983.15	-0.16 %
	10^{-9}	-1.84	1000.00	998.16	-0.02 %
Mitigation 2	10^{-7}	-165.06	1000.00	834.94	-1.36 %
	10^{-8}	-16.18	1000.00	983.82	-0.14 %
	10^{-9}	-1.57	1000.00	998.43	-0.02 %

To summarize, in this chapter the methods to mitigate the two main weaknesses of the density method for thermo-fluid TO are introduced and analyzed. First the errors caused by the blurred walls of the density method are mitigated by tuning the modeling parameters of the density-based turbulent fluid model. The tuning is performed by executing the meta-optimization procedure introduced in Section 3.7 for a subset of the modeling parameters. An optimization sweep is performed for three Darcy numbers. The results of the mitigation method are compared to the two cases introduced in Section 4.2, both graphically and based on their RRSE values. Mitigation 1 shows a significant decrease in error compared to the two reference cases (a minimum improvement of 52 % and a maximum improvement of 94 %) and appears to find a result which balances the main error effects ('liquid metal' and 'stagnant fluid') for each Darcy number. By reducing the Darcy number the error is reduced but a point of diminishing returns is visible. Furthermore, a shift is observed between the dominating error effects when decreasing Darcy numbers as indicated by the balance between the three individual RRSE values, which is consistent with the theory of these error effects but suggests the existence of an optimal Darcy number.

Mitigation 2 aims to reduce the thermal error caused by fluid flowing through the porous solid material by altering the thermal conductivity of the solid alongside the modeling parameters of the first mitigation method. The same meta-optimization procedure and optimization sweep as the first mitigation method is used to tune the thermal conductivity for three Darcy numbers. The results of this mitigation method are compared to the results of the first mitigation method. An analysis is performed on the erroneous loss of heat through the bottom boundary of the simulation domain. One of the results of the mitigation method (with Darcy number of 10^{-7}) shows a 66 % lower RRSE value than the first mitigation method but induces a much higher heat loss through the bottom boundary, indicating that the meta-optimization converged to an inferior optimum where unwanted side-effects such as temperature-dependency of the heat flux play a significant role. The other results of the turbulent porous solid material mitigation method show 13 % and 85 % higher error values than the first mitigation method.

Assessment of mitigation methods for weaknesses of the density method

In this chapter, the mitigation methods discussed in Chapter 5 will be tested and assessed to determine how well they apply to more complex geometries which are generated using a turbulent density-based fluid model. A test problem is introduced in Section 6.1 in which the accuracy of a density-based fluid model can be analyzed without the need for a converged TO. In Section 6.2 the average RRSE values as well as temperature error values obtained by the optimized modeling parameters are compared to a parameter sweep.

6.1. U-channel test case set-up

To assess the mitigation methods, a test case is devised which eliminates the dependency on a density distribution that is generated by a TO algorithm. This is achieved by disabling the TO algorithm and manually specifying a fixed density distribution. In contrast to the result of a TO, this test case has a reference geometry (with crisp boundaries) which can be used to analyze the modeling accuracy of the density-based solution of that same geometry.

To test the mitigation methods on a situation that is relevant to a TO, a more complex geometry than that of the meta-optimization (Section 3.7) is needed but the user should still be able to manually define the geometry. Therefore a U-channel geometry is selected with the same channel height as the meta-optimization geometry, as this meets the requirements and is commonly used as a verification scenario. Fig. 6.1 displays the U-channel scenario schematically. A fully developed turbulent velocity profile with $Re_{D_h} = 6000$ (same as meta-optimization) is specified at the inlet (arrows) and a heated solid region is specified at the inner region of the bend (red region). A more detailed set-up of the turbulent U-channel test problem is given in Appendix E.




Figure 6.1: Schematic of U-channel scenario. Fluid region is blue, solid regions are gray, heated solid region is red, arrows indicate inlet velocity distribution.

The flow of a conventional fluid through this geometry shows varying boundary layer thicknesses (Fig. 6.2a). The temperature of the heat source (Fig. 6.2b) is the simulation output which is most relevant to a thermo-fluid TO, since lowering this (and thus increasing heat transfer) will be the main objective of the TO optimizer. To evaluate whether the density-based fluid models over-estimate or under-estimate the heat transfer, their average heat source temperature will

be compared to that of the reference simulation. A ‘temperature error’ measure ϵ_T is constructed as:

$$\epsilon_T = \frac{T_d - T_r}{T_r - T_{in}} \quad (6.1)$$

where T_d is the average heat source temperature of the density-based simulation, T_r is the average heat source temperature of the reference simulation and T_{in} is the inlet temperature. This temperature error is thus normalized with respect to the $\Delta T = T_r - T_{in}$ of the reference simulation. A temperature error of 100 % indicates that the ΔT of the density-based simulation is twice as large as the reference and a temperature error of -100 % corresponds to a ΔT that is zero for the density-based simulation. Along with the ϵ_T , the average RRSE values defined in Section 3.6 are used to analyze the modeling accuracy. Although they do not show whether an over- or under-estimation of the heat transfer occurs, they cover more aspects of the modeling accuracy.

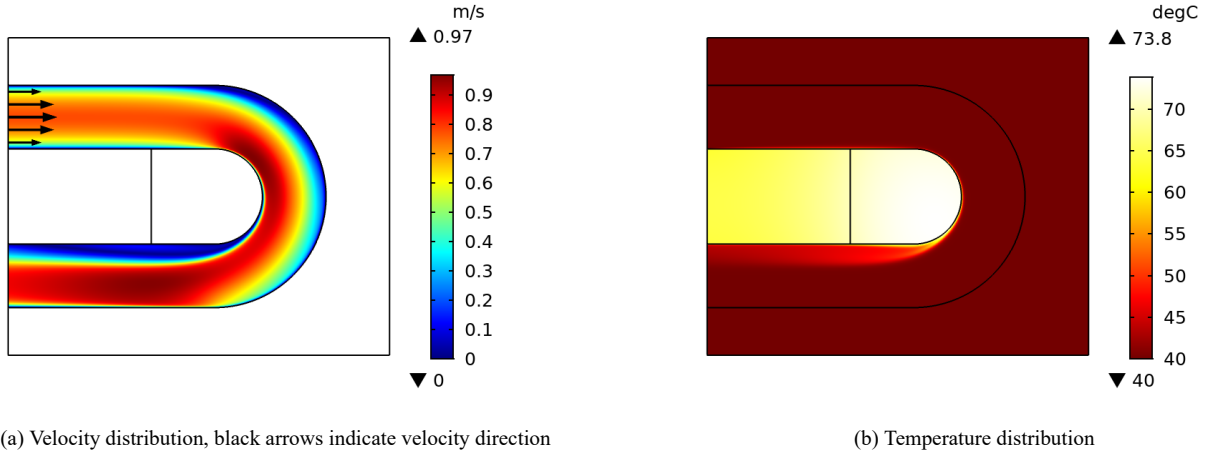


Figure 6.2: Results of the reference simulation of the U-channel geometry

6.2. U-channel test case results

A parameter sweep is performed for the turbulent density-based U-channel simulation to get a rough estimate of the range of average heat source temperatures achievable with the density-based fluid model. To avoid extremely long computation times, only the most sensitive modeling parameters are included in the parameter sweep. The selected parameters and their values are displayed in Table 6.1.

Table 6.1: Modeling parameters and their values included in the U-channel parameter sweep

Name	Symbol	Values
Darcy interpolation penalty	q	0.01 0.1 1 10
SIMP interpolation penalty	p	2 3 4
Turbulent viscosity penalty multiplier	$n_{\bar{\nu}}$	10^{-4} 10^{-2} $10^{-1.5}$
Darcy number	Da	10^{-9} 10^{-8} 10^{-7}
Projection slope	β	4 8 12

First the parameter sweep is analyzed, after which the behavior of the fluid with the meta-optimized modeling parameters from Table 5.2 is compared to the sweep. Fig. 6.3a displays box plots of the resulting temperature error values grouped by their Darcy number. Here the median of the data is displayed as a red line; The 25th and 75th percentiles are displayed as the lower and upper bounds of the box, respectively and the extremes of the dataset are displayed as whiskers. The spread of the temperature error becomes significantly larger with decreasing Darcy number, indicating that the density-based fluid model becomes more sensitive to the correct selection of modeling parameters when the Darcy number decreases. Furthermore a trend of decreasing temperatures of the density-based simulation is visible (sometimes resulting in large negative temperature errors) with increasing Darcy number. This indicates that with higher Darcy numbers the overestimation of the heat transfer becomes larger, which is caused by the increasing porosity of the solid material and the increasing effect of ‘liquid metal’ at higher Darcy numbers. At low Darcy numbers the low porosity allows for a more accurate heat transfer, but the ‘stagnant fluid’ effect causes an under-estimation of the heat transfer.

The combination of the lower sensitivity to the modeling parameters and the overestimation of the heat transfer at a Darcy number of 10^{-7} results in an impossible task to achieve zero temperature error for the turbulent U-channel scenario.

Box plots of the average RRSE of the parameter sweep are displayed in Fig. 6.3b. From these box plots it is observed that the same increasing spread is visible with decreasing Darcy numbers, indicating that not only the thermal behavior becomes more sensitive to the modeling parameter settings but the physical behavior does too. The high average RRSE values indicate that the situations where the temperature error is close to zero (see Fig. 6.3a) are not the result of a high overall accuracy, but rather of a situation that incidentally approaches the heat transfer of the reference while the physical flow might still be inaccurate.

The box plot in Fig. 6.3a also displays the temperature errors when using the modeling parameters resulting from mitigation 1 on the U-channel geometry and Fig. 6.3b also displays the average RRSE values for this scenario. The temperature errors (Fig. 6.3a) are all relatively close to 0 compared to the statistical results of the parameter sweep, although for the Darcy numbers of 10^{-8} and 10^{-9} better matching temperatures are observed in the sweep. In all cases, the mitigation results overestimate the heat transfer to some extent. Dependence on Darcy number is noticeably lower than for the parameter sweep results. When looking at the average RRSE, it is clear that the error achieved with mitigation 1 is significantly lower than those of the sweep for all Darcy numbers. The biggest difference in RRSE is found at a Darcy number of 10^{-9} where the lowest error found using the parameter sweep is 15 percentage point higher than the error resulting from the optimized parameters. This average RRSE is still 4 times higher than the average RRSE of the optimized modeling parameters applied to the straight channel geometry, indicating that the universality of optimal modeling parameters is limited. However, an argument can be made that the turbulent U-channel is an exceptionally challenging scenario for the density-based fluid model. Compared to a full parameter sweep which contains all modeling parameters and uses sufficiently small steps to capture the (possibly multiple) optimum/optima accurately, the meta-optimization achieves a reasonable accuracy with relatively few model evaluations (around 200 compared to thousands).

Fig. 6.3a also displays the comparison between the temperature errors resulting from mitigation 1 and mitigation 2 when their resulting modeling parameter values are applied to the turbulent U-channel geometry and Fig. 6.3b also displays the average RRSE values of the same comparison. Note that mitigation 2 cannot be directly compared to the parameter sweep because the latter does not include the thermal conductivity. It is observed that the temperature error of mitigation 2 grows significantly faster with increasing Darcy numbers than the error of mitigation 1. This is caused by the increasingly large compensation of the thermal conductivity made by mitigation 2 with increasing Darcy numbers. Where this compensation results in a relatively high accuracy in the straight channel geometry used by the meta-optimization, it causes an overshoot in the turbulent U-channel. Finally, the relatively low absolute temperature errors achieved by mitigation 2 at the Darcy numbers of 10^{-8} and 10^{-9} are the result of an inaccurate flow, as indicated by their higher average RRSE value compared to mitigation 1. These observations indicate that altering the thermal conductivity of the porous solid material (mitigation 2) is not a robust method to increase the thermal accuracy of the turbulent density-based fluid model, as the sensitivity to a different geometry is too high.

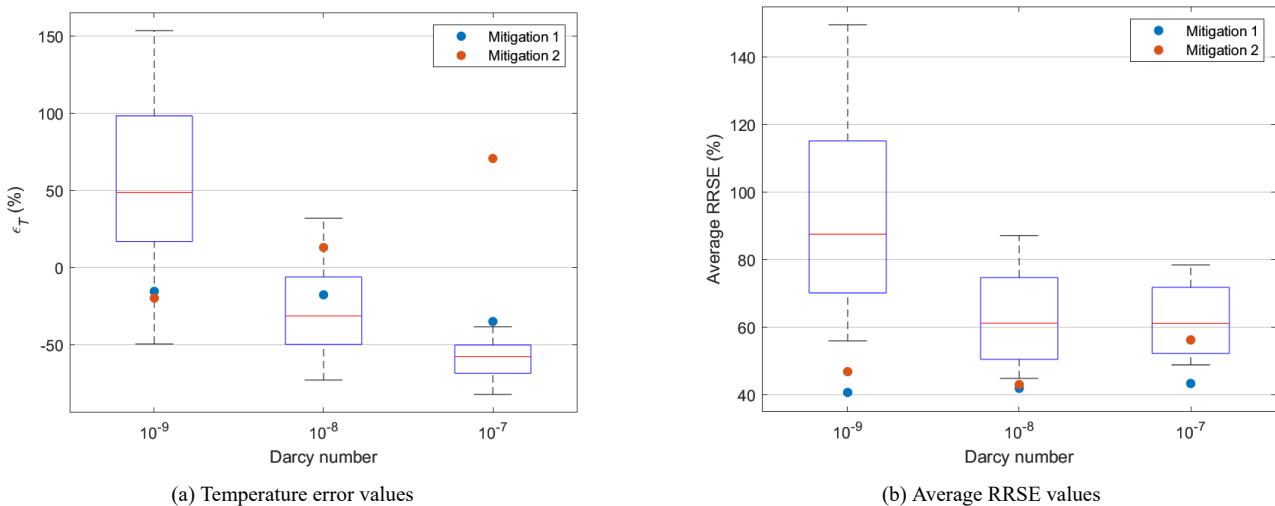


Figure 6.3: Box plots of temperature error values and average RRSE values resulting from turbulent U-channel parameter sweep, compared to the resulting modeling parameters of mitigation 1 and mitigation 2 applied to turbulent U-channel. Results are grouped by their Darcy number

Thermo-fluid Topology Optimization of a Prodrive Power Module cooling interface with meta-optimized modeling parameters

This chapter focuses on two topics: Applying thermo-fluid TO to optimize a cooling interface for a simplified Prodrive Power Module and comparing the performance of a TO with manually tuned modeling parameters to that with meta-optimized modeling parameters. This is first discussed for a laminar TO in Section 7.1 because no representative results are obtained from a turbulent TO due to instabilities. Therefore, Section 7.2 focuses on the issues found in the process of implementation of the turbulent TO and displays intermediate results. Both TO's serve to demonstrate the capability of the method in a representative setting and therefore some liberties are taken for improved visualization. A 2D TO is performed to keep computational costs relatively low compared to a 3D TO of the same scenario.

7.1. Laminar Topology Optimization

This section will discuss a laminar thermo-fluid TO applied to a cooling interface for a Prodrive Power Module. First, the problem definition is introduced in Section 7.1.1, after which the results of a manually tuned TO are discussed in section Section 7.1.2. Finally, in Section 7.1.3 the results of a 'meta-optimized' TO are discussed and compared to those of the manually tuned TO.

7.1.1. Problem definition

The problem definition of the Power Module cooling interface optimization encapsulates the following aspects:

- Fluid model
- Geometry of the optimization domain
- Fluid boundary conditions
- Thermal boundary conditions
- Materials
- Filtering functions
- Optimization objective
- Optimization constraints

Each aspect will be discussed in the following paragraphs.

The fluid model used with the laminar thermo-fluid TO is as described in Section 2.4.1 and the conjugate heat transfer model is implemented as discussed in Section 2.4.3. The geometry of the optimization is based on the dimensions of a Prodrive Power Module. A geometry with symmetry condition at the top horizontal boundary is selected to minimize computation time, see Fig. 7.1. The optimization domain (light blue) connects to inlet and outlet channels (dark blue) on the left and right side, respectively. The density variables in these channels are set to 100 % fluid. A fixed solid border domain (gray, density variables fixed at 0 % fluid) surrounds the optimization domain, to provide the correct boundary condition to the density variables. A gap is left between the solid border and the inlet/outlet channels to avoid disturbing the flow in these channels due to the exponential decay of the Darcy force (Section 4.2.3. The characteristic feature of the simplified Prodrive Power Module is the small local heat load (red) with dimensions of a single power semiconductor,

which is placed off-center from the fluid path. This heat source region is fixed at 0 % fluid, such that it cannot transfer heat directly into the fluid but rather has to conduct it through a heat spreader first. In a real Prodrive Power Module the local heat load of a power semiconductor cannot be directly applied to the cooling interface but is attached to a stack consisting of multiple layers before contacting the cooling interface. This stack spreads the heat flux non-uniformly over a larger surface at the contact face. The simplification of the heat load results in a more extreme heat flux for the TO than real-life. Initially, the density variables in the optimization domain (light blue) are uniformly set to $\gamma = 0.5$.




Figure 7.1: Schematic (not to scale) of the geometry used with the laminar Power Module TO. The optimization domain is indicated with light blue, the fluid inlet (left) and outlet (right) channels are dark blue, the solid heat source is red and the solid border is gray. Boundary conditions and dimensions of relevant geometries are displayed.

A fully developed velocity profile U_{in} is specified at the fluid inlet which accounts for the symmetry condition at the top of the inlet channel (see Fig. 7.1). The Reynolds number of this velocity profile is set to $Re_{D_h} = 100$ and is based on the hydraulic diameter of the inlet channel. Numerical experiments were initially performed with a Reynolds number of 1000 (selected based on applicability to the Power Module cooling system, Section 3.2) but these showed instabilities due to signs of turbulence. A zero-pressure condition is specified at the outlet boundary. Apart from the symmetry boundary and inlet and outlet boundaries, all other external boundaries are specified as no-slip. The uniform inlet temperature is set to $T_{in} = 40^\circ\text{C}$ and a total heat load of $Q = 10\text{ W}$ is applied uniformly to the 6*6 mm heat source. Except for the inlet, outlet and symmetry boundaries, the other external boundaries are thermally insulated.

The fluid used in the laminar TO is water (Table 3.1) and a solid material with $\rho = 2700\text{ kg/m}^3$, $Cp_s = 910\text{ J/(kg * K)}$, $k_s = 50\text{ W/(m * K)}$ is used. Compared to the material properties of the turbulent simulations in previous chapters, the thermal conductivity is lowered in this simulation. This is done because conduction increasingly dominates convective heat transfer when the inlet velocity is reduced while keeping material properties constant. Performing a laminar TO without altering material properties often results in a design which is predominantly solid. Lowering the thermal conductivity of the solid material balances these heat transfer modes better such that the optimizer will generate more detailed fluid channels, which is desired for a good comparison between a manually tuned and a meta-optimized TO. Note that this does not necessarily result in improved thermal performance. The value of 50 W/(m * K) is selected after a sweep of different values for k_s and was found to generate sufficiently detailed channels.

The density variables are filtered with a blurring filter (Eq. (2.17)) with filter radius $R_{min} = 0.3\text{ mm}$ to enforce a minimum length scale and a hyperbolic projection filter (Eq. (2.19)) to generate high contrast designs without porous material. The slope β of the projection filter is manually tuned for the manually tuned TO and obtained from the meta-optimization for the meta-optimized TO.

The objective of the optimization is twofold. Firstly it minimizes the average temperature difference

$$\frac{1}{A} \int_A \Delta T dA = \frac{1}{A} \int_A T - T_{in} dA \quad (7.1)$$

between the solid heat source T and the inlet temperature T_{in} , by minimizing the thermal resistance from the source to the fluid. This form of the temperature is used instead of the absolute temperature, because the heat source temperature approaches the inlet temperature as the heat transfer grows to infinity. This allows for easier comparison of heat transfer performance between results. Secondly, the volume factor of fluid in the optimization domain is to be minimized, such that redundant branching channels are removed. The power loss P_{loss} (Eq. (2.38)) of the fluid is constrained to be below a multitude w of a reference loss P_{ref} in a fully open channel:

$$\frac{P_{loss}}{wP_{ref}} - 1 < 0. \quad (7.2)$$

Constraining the power loss allows the user to control the level of detail in the resulting geometry, along with the blurring and projection filters. A high power loss constraint allows the optimizer to generate narrower channels (down to the limiting size controlled by the blurring and projection filters) which clears space for more branching channels. A high

power loss generally allows for increased heat transfer, but with decreasing channel width the flow through them becomes increasingly diffuse due to the exponential decay of the Darcy force. This could pose issues with post-processing of the density variables.

7.1.2. Manually tuned Topology Optimization

The manually tuned laminar thermo-fluid TO is performed with the values of the modeling parameters which are displayed in Table 7.1. These parameters are the result of 8 attempts of this specific TO and the insight gained by many more simulations before that. The weight factor $w = 5$ of the power loss constraint is chosen fairly low (values up to 30 have been tested) to generate wider channels with few obstructions. A Darcy number value of 10^{-6} is chosen based on insights of a laminar meta-optimization (Appendix B) and an experimental TO to check the stability of the optimizer. The selected Darcy number shows good stability and satisfactory accuracy.

Table 7.1: Manually tuned modeling parameters for laminar thermo-fluid TO

Parameter	Da	p	q	β	R_{min}	w
Value	10^{-6}	3	1	6	0.3 mm	5

The other modeling parameters are chosen such that the results meet the following qualitative requirements:

1. The optimizer generates no or few regions of intermediate density.
2. The optimizer generates smooth walls and no branching channels which have a dead end.
3. The Darcy force in an open channel is low enough to not obstruct the flow. Based on an analysis of the laminar boundary layer and laminar meta-optimizations for this specific problem (Appendix B), a Darcy force threshold of $4 \times 10^4 \text{ Pa} \cdot \text{s}/\text{m}^2$ is determined to indicate the transition from ‘free flow’ to flow which is influenced by the Darcy force.

After 8 iterations of manual tuning the resulting density distribution in Fig. 7.2a displays that for the largest part of the domain the transitions from solid to fluid are sharp, that the majority of the channel walls are smooth, and that no dead end branching channels are generated. In Fig. 7.2b, the Darcy force is plotted on a logarithmic scale, and a black contour line is plotted at the threshold Darcy force. It is observed that sufficiently large parts of the channel have Darcy force values below the threshold. The velocity distribution in Fig. 7.2c shows that no large regions of diffuse flow are present. The resulting temperature distribution displayed in Fig. 7.2d shows that the heat is spread fairly well throughout the solid heat spreader, which explains why the optimizer attempts to maximize the interface area between the heat spreader and the channels by making the heat spreader as large as possible and by adding ‘meanders’ near the heat source.

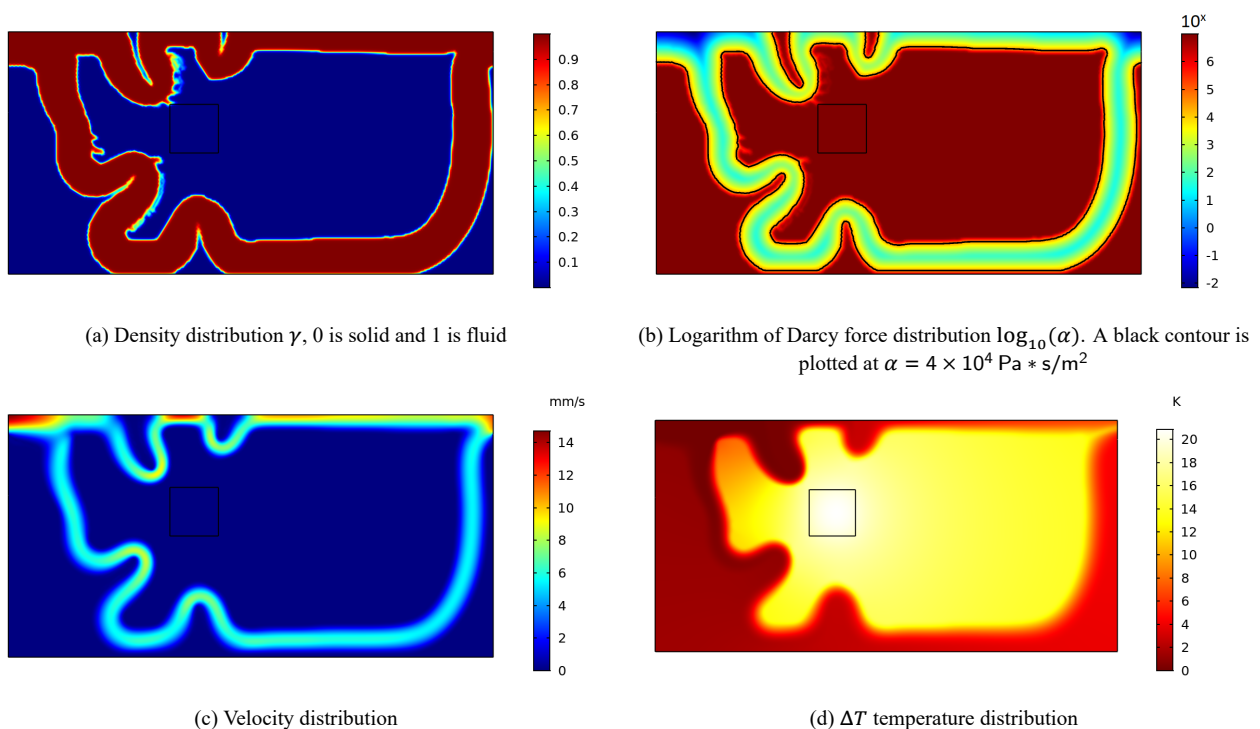


Figure 7.2: Resulting simulation outputs from manually tuned TO, with modeling parameters given in Table 7.1

Post-processing of the density distribution is done in two ways: with a threshold of the density variables of $\gamma = 0.5$ and with the threshold Darcy force $\alpha = 4 \times 10^4 \text{ Pa} \cdot \text{s}/\text{m}^2$ which is determined in Appendix B. Figures 7.3a and 7.3b display the post-processed geometries and their resulting temperature distributions. The geometry which is post-processed with $\gamma = 0.5$ (Fig. 7.3a) displays the same rough walls and sharp inner corners of the meanders as is visible in the density distribution (Fig. 7.2a), whereas the geometry which is post-processed with $\alpha = 4 \times 10^4 \text{ Pa} \cdot \text{s}/\text{m}^2$ is much more smooth, has wider inner corners of the meanders and more narrow channels.




Figure 7.3: Post-processed geometry and resulting ΔT temperature distribution using two threshold values for post-processing. Both geometries are post-processed from the manually tuned laminar thermo-fluid TO

Both post-processed results display higher temperatures and a lower power loss than the density-based result, as displayed in Table 7.2. This indicates that the heat transfer is over-estimated in the density-based simulations compared to the post-processed simulations. Because both post-processed have some power loss margin, their cooling performance can be increased when the velocity is increased such that the flow reaches the power loss limit, which is displayed in the last column of Table 7.2. The differences in ΔT between the density-based result and the post-processed results are decreased significantly by matching the power losses, with the best match being the geometry which is post-processed from $\alpha = 4 \times 10^4 \text{ Pa} \cdot \text{s}/\text{m}^2$ which has a difference of 0.34 K (1.7 %) with the density-based result.

Table 7.2: Power loss and average heat source ΔT values of the manually tuned density-based result and two of its post-processed results with different threshold values. The maximum allowed power loss is 2.33×10^{-5} . The last column displays the average ΔT value when the velocity of the post-processed simulation is increased to match the power loss constraint

Model	Threshold	Power loss	% of constraint power loss	Avg heat source ΔT	Avg heat source ΔT at 100 % constraint power loss
Density-based	n/a	2.33×10^{-5}	100	20.00 K	n/a
Post-processed	$\gamma = 0.5$	9.26×10^{-6}	40	25.63 K	22.77 K
	$\alpha = 4 \times 10^4 \text{ Pa} \cdot \text{s}/\text{m}^2$	1.53×10^{-5}	65	21.73 K	20.34 K

7.1.3. Meta-optimized Topology Optimization

A laminar meta-optimization is performed (in Appendix B, based on mitigation 1 with adaptation for laminar flow) which outputs a set of modeling parameters that can be used for the laminar thermo-fluid TO. The resulting modeling parameters are displayed in Table 7.3. The most notable differences between the manually tuned parameters and the meta-optimized parameters are the low SIMP interpolation penalty p (with a penalty of 1 the interpolation is linear) and the relatively low slope β of the hyperbolic projection filter.

Table 7.3: Meta-optimized modeling parameters for laminar thermo-fluid TO

Parameter	Da	p	q	β	R_{min}	w
Value	10^{-6}	1.0	0.8	2.0	0.3 mm	5

The resulting geometry (see Figs. 7.4a to 7.4d) meets all the qualitative requirements introduced in Section 7.1.2. The shape of the heat spreader generated by this TO is fairly similar to that of the manually tuned version. It displays a large heat spreader (Fig. 7.4d) with maximized interface area by meandering channels. Compared to the manually tuned TO, the meta-optimized TO has slightly more porous material at the inner corners of the meandering channels (Fig. 7.4a), but the Darcy force threshold $\alpha = 4 \times 10^4 \text{ Pa} \cdot \text{s}/\text{m}^2$ lies at a higher density value such that the resulting channel has smooth corners (Fig. 7.4b). In terms of the smoothness of the channel walls the meta-optimized TO performs better than the manually tuned TO.

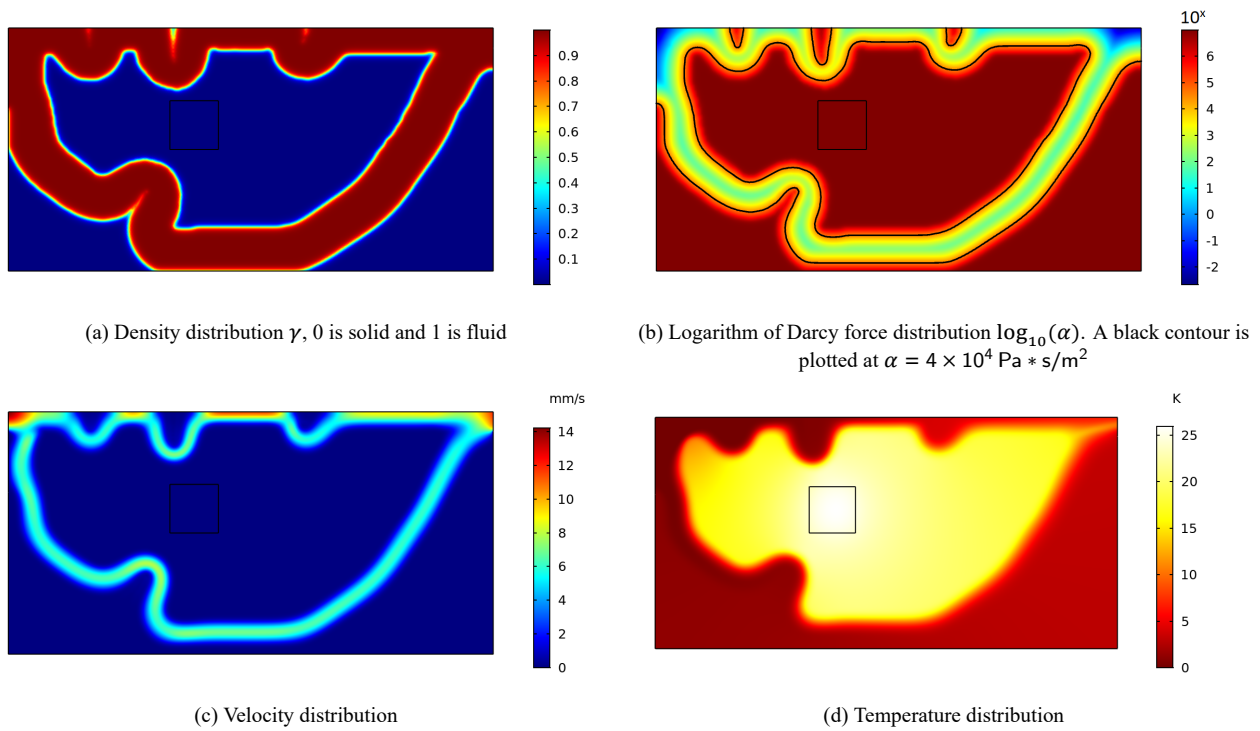


Figure 7.4: Resulting simulation outputs from meta-optimized TO, with modeling parameters given in Table 7.3

The differences between the two post-processing methods are quite extreme. Figure 7.5a displays that the meanders generated by the density distribution are largely filtered out when post-processing with $\gamma = 0.5$, where in Fig. 7.5b the meanders are represented well by post-processing with $\alpha = 4 \times 10^4 \text{ Pa} \cdot \text{s}/\text{m}^2$.




Figure 7.5: Post-processed geometry and resulting temperature distribution using two threshold values for post-processing. Both geometries are post-processed from the meta-optimized laminar thermo-fluid TO

Table 7.4 shows that the average heat source ΔT of the meta-optimized TO (density-based result) is 25 % higher than that of the manually tuned TO, indicating a worse cooling performance. However, when post-processing the density variables of the meta-optimized TO with the Darcy force threshold, a 7 % *lower* ΔT is obtained compared to the density-based result where the manually tuned TO results in a 9 % *higher* post-processed ΔT . The former effect (temperature decreases with post-processing) is desired over the latter, as the post-processed geometry results in ‘free’ extra performance compared to what is obtained by the TO. The post-processed geometries also have significant power loss margin, which gives room for improved cooling performance. After matching the power losses, the temperatures of the post-processed geometries decrease significantly (see last column of Table 7.4), but this increases the difference between the Darcy force-based post-processing and the density-based result (now 3.66 K, was 1.75 K before matching power losses).

Table 7.4: Power loss and average heat source ΔT values of the meta-optimized density-based result and two of its post-processed results with different threshold values. The maximum allowed power loss is 2.33×10^{-5} . The last column displays the average ΔT value when the velocity of the post-processed simulation is increased to match the power loss constraint

Model	Threshold	Power loss	% of constraint power loss	Avg heat source ΔT	Avg heat source ΔT at 100 % constraint power loss
Density-based	n/a	2.33×10^{-5}	100	25.09 K	n/a
Post-processed	$\gamma = 0.5$	5.72×10^{-6}	24	35.06 K	29.21 K
	$\alpha = 4 \times 10^4 \text{ Pa} \cdot \text{s/m}^2$	1.29×10^{-5}	55	23.34 K	21.43 K

Unfortunately, even after matching the power losses the ΔT of the highest performing post-processed geometry from the meta-optimized TO (21.43 K) is still 1.09 K (5 %) higher than that of the manually tuned TO (20.34 K, Table 7.2). Therefore it cannot be concluded that the current formulation of the meta-optimization improves the performance of a laminar thermo-fluid TO, but the margin between the two methods is fairly small. The meta-optimized TO does however have considerable qualitative benefits. Firstly, the resulting geometry is more favorable as it contains less artefacts (rough walls, porous material). Secondly, this TO provides satisfactory results at the first attempt, where a manual tuning procedure takes many more attempts. Finally, the Darcy force threshold which is determined from the meta-optimization has proven to provide satisfactory accuracy. Using such a threshold can eliminate the multiple post-processing cycles required to find the correct threshold manually.

7.2. Discussion of turbulent Topology Optimization

Turbulent thermo-fluid TO was attempted following the model description given in Section 2.4.2. Unfortunately no satisfactory results were obtained from many attempts due to various stability issues. This section will discuss these issues and proposes a path for future research.

The turbulent thermo-fluid TO is set-up in a similar fashion as the laminar case discussed in Section 7.1.1 with three main differences: the inlet velocity is set such that $Re_{D_h} = 6000$, the material properties are set to those given in Table 3.1 and the heat source is set to dissipate 60 W.

From literature it is expected to obtain accurate results from a turbulent thermo-fluid TO when interpolation penalties and the slope of a projection filter are increased and the Darcy number is decreased[15]. This behavior is partly confirmed by the mitigation methods analyzed in Chapter 5. However, when directly applying a set of modeling parameter values which is known to provide accurate results to a turbulent TO, the optimizer often does not converge beyond the initial geometry (a porous medium with 50 % porosity). For some combinations of modeling parameters the thermo-fluid solver cannot even find a solution for the initial geometry. The stability issues of the turbulent thermo-fluid TO can be divided into two categories: unstable optimizer and unstable thermo-fluid solver. Both categories will be discussed in their respective sub-sections (Sections 7.2.1 and 7.2.2). Following the assessment of the two categories of stability issues, a continuation strategy is implemented which changes the values of the modeling parameters while the optimization converges, such that instabilities are avoided on the path towards an accurate solution. The results obtained with various continuation strategies are discussed in Section 7.2.3.

7.2.1. Unstable optimizer with thermo-fluid Topology Optimization

An instability of the optimizer when performing a thermo-fluid TO implies that the optimization algorithm cannot compute a new iteration of the density variables and the optimization process is terminated. A precursor of a termination due to this instability when using COMSOL is a warning that the "Adjoint solution failed" and the optimizer requests a "reduced step". This indicates that the optimizer must reduce the magnitude of change in the update of the density variables to remain stable. Both the termination and the warning seem correlated to the smoothness and area of the wall (interface between solid and fluid). A rough wall with very little material having intermediate density causes more instability than a smooth wall with a large porous area. The underlying cause for this instability is the quality of the sensitivity information which is passed from the thermo-fluid solution to the optimizer. The rough and sharp walls generate sensitivity data which is concentrated on very small regions (sensitivities are highest at regions of intermediate density) and which is noisy (a rough wall causes erroneous local turbulence which increases the magnitude of the sensitivity). This hinders the optimizer in determining the update of density variables.

7.2.2. Unstable thermo-fluid solver with thermo-fluid Topology Optimization

An instability of the thermo-fluid solver when performing a thermo-fluid TO implies that the solution of the coupled thermo-fluid solver diverges for a given distribution of the density variables. This either results in a variable which locally diverges to infinity, or that the solver reaches its maximum number iterations without finding a converged solution. In both cases the optimization process is terminated. The cause for this instability is related to the mesh size, the density distribution generated by the optimizer and the turbulence model. Due to the type of turbulence model being 'low

Reynolds number', the mesh resolution must be fine enough throughout the domain to resolve the full turbulent boundary layer. In a thermo-fluid TO the shape and thickness of the boundary layer is controlled by the density distribution. On a fixed mesh, the density distribution can generate regions where the boundary layer becomes too thin to accurately resolve, resulting in divergence of the thermo-fluid solver. In other cases, the density distribution may force the turbulence model into a scenario where it cannot resolve the distribution of the eddy viscosity (see Section 2.2.2).

7.2.3. Results of continuation strategy

The instabilities discussed in Sections 7.2.1 and 7.2.2 are to some extent avoided by a carefully constructed continuation strategy. Each step of the continuation should be tuned to work towards the final accuracy without causing too much instability. Apart from choosing the continuation strategy, the settings of other modeling parameters also affect the result and stability. This process is attempted many times with different continuation strategies and modeling parameter values. Various parameters were included and removed from the continuation and different starting values and step sizes were tested. The manual trial-and-error approach is somewhat of a 'gambling game', as it cannot be predicted whether an attempt will converge without instabilities and the cause of a diverged result is not always clear. This section gives a few examples of results obtained from these attempts, with the goal to give insight to the sensitivity of the turbulent thermo-fluid TO to the modeling parameter settings. The selected modeling parameters and continuation strategies for these examples are displayed in Table 7.5. Most of the examples are the last converging step of the attempted continuation strategy.

Table 7.5: Modeling parameters used for four examples of different continuation strategies. Highlighted cells indicate that the parameter has changed in the continuation

Example	Continuation start									Continuation end	
	Da	p	q	β	$n_{\bar{v}}$	n_G	w	R_{min}	q	β	
1	10^{-7}	3	0.1	1.5	10^{-4}	10^3	30	0.3 mm	0.1	3	
2	10^{-7}	3	0.01	0	10^{-4}	10^3	10	0.5 mm	0.01	6	
3	10^{-7}	3	0.01	0	10^{-4}	10^3	10	0.5 mm	0.1	6	

Example 1, displayed in Figs. 7.6a and 7.6b, is the result after one continuation step of the projection filter's slope β . The blurring filter radius R_{min} is chosen to match the mesh size, which results in fine geometries. The power loss constraint w is set fairly high, to allow for narrow channels. The channels are erratic and the flow is forced to diffuse through solid material (bottom left channel).




Figure 7.6: Results of continuation method example 1, see Table 7.5 for modeling parameter values

Example 2, displayed in Figs. 7.7a and 7.7b, is the result after one step of the projection filter's slope. A slightly larger blurring filter radius and strongly restricting power losses are selected to enforce a longer length scale and create more open, smooth channels. The geometry is very basic but contains a large amount of porous material. The flow is predictable but still fairly diffuse. The combination of modeling parameters strongly triggers the liquid metal effect, resulting in an unrealistic over-estimation of heat transfer which causes the heat spreader area to be minimized.




Figure 7.7: Results of continuation method example 2, see Table 7.5 for modeling parameter values

Example 3, displayed in Figs. 7.8a and 7.8b, is the result of a continuation step following Example 2. The Darcy interpolation penalty is increased, which eliminates the liquid metal effect and causes the optimizer to design a larger heat spreader. The flow is only diffuse in some areas but the density variables show large regions of porous material.





Figure 7.8: Results of continuation method example 3, see Table 7.5 for modeling parameter values

To summarize, this chapter discusses the application of thermo-fluid TO to optimize a cooling interface for a Pro-drive Power Module. First the optimization problem of a laminar thermo-fluid TO has been defined in Section 7.1.1, after which the results of this optimization obtained with manually tuned modeling parameters as well as with meta-optimized parameters has been discussed in Sections 7.1.2 and 7.1.3. The thermal performance of the meta-optimized TO is lower than that of the manually tuned TO (a 5% higher ΔT is obtained after post-processing and matching power losses). However, the meta-optimized approach and result has clear qualitative advantages over the manually tuned TO, being mainly that no iterations are required to tune the modeling parameters and that a post-processing threshold can be determined from the meta-optimization results. The Darcy force-based post-processing threshold shows relatively high thermal accuracy compared to a conventional post-processing threshold of the density variables ($\gamma = 0.5$), although no post-processing sweep has been performed to substantiate its optimality.

The application of a turbulent thermo-fluid TO faces various stability issues, which have been discussed in Section 7.2. Because of these stability problems, no usable results were obtained for a turbulent cooling system. A number of intermediate results from different continuation strategies and modeling parameter settings have been discussed to indicate the sensitivity of the optimization to these settings.



Conclusions and outlook

8.1. Conclusions

This thesis is the result of an investigation into how to quantitatively analyze and improve the accuracy of (turbulent) density-based thermo-fluid Topology Optimization (TO). From the researched literature a number of issues regarding density-based thermo-fluid TO were identified. Firstly, a quantitative study of the modeling accuracy when using density-based thermo-fluid models was lacking. Comparative studies of modeling parameters were only performed qualitatively or based on the achieved objective of a resulting geometry, which promotes the use of parameter settings that over-estimate the cooling performance of the density-based model. Secondly, two main weaknesses of density-based thermo-fluid models were found: the finite porosity of solid material causes convection through solid regions and the blurred walls affect the shapes of adjacent boundary layers. The current heuristic countermeasures are not proven to be effective. Thirdly, the tuning of modeling parameters to increase accuracy (or rather, performance due to the lack of a quantified measure of accuracy) is currently a tedious manual process. Many authors state the high sensitivity of parameter settings, especially in turbulent flows [16][35]. With a large number of interacting modeling parameters, this manual tuning is guaranteed to be sub-optimal. Therefore the main research question of this thesis was: “How can the errors caused by the two main weaknesses of the density method when applied to thermo-fluids, being ‘porous solid material’ and ‘blurred walls’, be minimized?”

The first issue of lacking quantitative analysis of the modeling accuracy was solved by developing a simulation framework consisting of three components. Firstly, verified reference simulations of developing laminar ($Re_{D_h} = 100$) and turbulent ($Re_{D_h} = 6000$) boundary layers in a straight channel. Secondly, density-based simulations of the same scenario and thirdly, a quantitative measure of the modeling accuracy: the average of the RRSE values of the velocity, temperature and turbulent viscosity. The framework was shown to provide sufficiently detailed information on the overall accuracy of the density-based fluid model by comparing the RRSE data to a detailed qualitative analysis of the errors found in a density-based turbulent flow. This framework answers research question RQ.1a: “How can the modeling accuracy of the density method when applied to thermo-fluid TO be defined?”

A qualitative and quantitative analysis using the framework has provided insights into the influence of the different modeling parameters on the modeling accuracy, which answers research question RQ.1b: “How do the modeling parameters of the density-based fluid models affect the modeling accuracy?” A first conclusion to be drawn from this analysis is that by tuning the modeling parameters a balance can be found between different errors caused by the blurred walls. Secondly, the convection through solid material can be minimized by reducing the porosity of the solid material, but this causes large errors in adjacent boundary layers. Lastly, a ‘Darcy force threshold’ can be determined (for a given flow condition) that forms a good measure of how the blurred density-based wall can be converted back to a crisp wall, and therefore can be used to post-process the density variables of a TO.

To address the second issue (weaknesses found in density-based thermo-fluid TO), two mitigation methods were proposed. By optimizing the modeling parameters of a density-based thermo-fluid model (meta-optimization) the errors caused by the blurred walls were minimized and by altering the thermal conductivity of the solid material the errors caused by the porous solid material could theoretically be reduced. For a turbulent flow, the first mitigation method showed promising results, as it displayed between 52 % and 94 % reduced errors compared to two reference cases. The second mitigation method showed an error reduction of 66 % but also an increase of 85 % compared to the first method, as well as deviating behavior which made it less advantageous. The manually selected Darcy number was found to have a large influence on the achievable accuracy by both mitigation methods.

The generality of the meta-optimized modeling parameters generated for a turbulent flow was tested by applying them to a density-based thermo-fluid simulation of a turbulent U-channel. The first mitigation method showed promising results as it exhibited lower average RRSE values than those obtained by any combination of parameters in the sweep, although the mitigation results do still consistently over-estimate the heat transfer. The second mitigation method showed

a high sensitivity to the change of geometry, with results that over-compensate the errors caused by the porous solid material at high Darcy numbers, resulting in a large under-estimation of the heat transfer. The second mitigation method was therefore not deemed feasible at this point.

To test how the first mitigation method influences the results of a density-based thermo-fluid TO, a manually tuned TO was compared to a TO which uses meta-optimized modeling parameters. Unfortunately this comparison was only possible for a laminar TO, as attempts of a turbulent TO faced several stability issues caused by a high sensitivity to the modeling parameter settings. In comparison to a conventional post-processing threshold, the Darcy force threshold was found to provide geometries which consistently achieved higher thermal performance and a better reproduction of the density-based result. When using the Darcy force threshold, the mitigated TO shows similar accuracy compared to the manually tuned TO but with 5 % lower cooling performance. With the results of:

1. The meta-optimized mitigation methods,
2. The test for generality of the optimized modeling parameters,
3. The investigation into the effects of the most promising mitigation method on a laminar TO,

the research question RQ.2: “To what extent can mitigation strategies minimize this error?” is answered.

To address the third issue (tedious manual parameter tuning), the meta-optimized TO showed clear qualitative advantages over the manually tuned TO. By using meta-optimized modeling parameters, the TO provided satisfactory results from its first attempt, where the manually tuned TO required 8 attempts and significant understanding of density-based TO and its weaknesses. Furthermore, the use of the Darcy force threshold to post-process the density variables can eliminate the need to tune the post-processing threshold manually.

The framework for quantitative analysis is proven to be functional and can be used for benchmarking of future results as well as gaining further insight into the accuracy and stability of density-based fluid models. The concept of meta-optimization to mitigate the errors caused by blurred walls is promising especially for turbulent thermo-fluid TO. It is recommended to start using meta-optimization for laminar flows and to further develop the concept such that it can be applied to turbulent TO as well. Finally, post-processing the resulting density variables of a laminar TO by using the Darcy force threshold provides a better alternative to conventional post-processing thresholds.

To answer the main research question concisely: the errors caused by the two main weaknesses of density-based thermo-fluid TO can be minimized by meta-optimizing the modeling parameters, using a framework that also allows quantitative analysis of the density-based fluid models. The qualitative issue of manual parameter tuning is also minimized by this meta-optimization procedure. With this answer, the goal to quantitatively analyze and improve the accuracy of (turbulent) density-based thermo-fluid Topology Optimization has been reached.

8.2. Future work

The methods to quantitatively analyze and improve on the accuracy of density-based TO show many opportunities for future research which are categorized in the quantification of modeling accuracy, the meta-optimization and the thermo-fluid TO:

- Quantification of modeling accuracy
 - The accuracy of the reference scenario based on a developing boundary layer in a channel could not be directly verified because insufficient correlation data was found. Furthermore, any correlation data that was available either described a laminar flow or a fully turbulent flow whereas the ‘turbulent’ Reynolds number used for this thesis lies in the transitional regime. Any deviation of the reference from correlation data could therefore be attributed partly to this intermediate Reynolds number, convoluting the accuracy analysis. A different scenario could be found which provides sufficient correlation data for verification of the reference, while still being relevant to the TO.
 - The detailed analysis of the weaknesses of density-based TO was performed for a turbulent flow, but no specific turbulent simulation outputs were checked. If such measures are included in the analysis, possibly other error-inducing effects could be identified.
- Meta-optimization
 - The meta-optimization procedure was based on a gradient-free optimizer because implementation of a gradient-based optimizer faced issues. In theory this should be possible as long as the objective function, constraints and interaction between the optimizer and density-based fluid model are all smooth and differentiable. This is the case for all modeling parameters except for the vertical gradient position y_{grad} , which requires re-meshing of the domain and lies outside of the capabilities of the built-in gradient-based optimizers in COMSOL. Finding a method to enable gradient-based meta-optimization could result in a significantly more efficient and possibly more accurate procedure.

- The results of mitigation 1 indicate that there may be an optimal Darcy number instead of the expected monotonic effect on the modeling accuracy. A numerical experiment can be performed which includes the Darcy number as an optimization variable.
 - The large errors of mitigation 2 are partly accountable to the freedom to alter the thermal conductivity to extreme values. Decreasing the range of the thermal conductivity minimizes effects such as the dependency of convective heat flux on local temperature. Doing so enables a more realistic feasibility analysis of this mitigation method.
 - The meta-optimization used for mitigation 2 possibly converged to an inferior local optimum. It incorporated the same starting values as mitigation 1. By starting the meta-optimization for mitigation 2 at the optimized values resulting from mitigation 1, the optimizer can perhaps avoid this local optimum and converge to a better one.
 - The gradient-free optimizer converges to the nearest local optimum and therefore other optima might exist with better results. The current formulation of the meta-optimization can be started from multiple randomized positions to find these optima, or another optimizer with better global convergence performance can be implemented.
 - The generality of the meta-optimized modeling parameters was only investigated for one different geometry with the same flow conditions. The robustness to varying flow conditions could be investigated. It is expected that the flow conditions have a large influence on the accuracy of the optimized modeling parameters, as is indicated by the differences between the laminar and turbulent meta-optimizations. Perhaps an allowable range of deviating flow conditions can be determined that keeps errors below a certain threshold.
 - The turbulent U-channel was found to be challenging for the density-based fluid model, possibly due to the ‘sharp’ bend which was selected. It is expected that with a larger bend radius the geometry becomes more forgiving and the errors will likely be lower. The effect of geometry variations on the achieved accuracy with optimized modeling parameters could give insights on the generality of the optimum found by the meta-optimization.
 - The meta-optimization procedure was implemented for a 2D flow condition. The procedure can be extended to 3D and the differences between the 3D and 2D versions can be analyzed. For turbulent flow, switching to 3D enables the computation of streamline vorticity which could alter the behavior of a turbulent boundary layer for a given set of modeling parameters.
 - The optimizer stability issues faced when implementing turbulent thermo-fluid TO are largely accounted to erroneous sensitivity information. It is likely that the sensitivities become erratic for some combinations of modeling parameters. Possibly the ‘health’ of the sensitivity information can also be meta-optimized, such that the meta-optimization finds a set of modeling parameters which provides high accuracy but also ensures sufficient stability of the TO. A prerequisite to enable sensitivity-optimization with the meta-optimization procedure is the use of a gradient-based optimizer.
 - Other mitigation methods to reduce convective heat transfer through solid material, such as penalizing the convective term of the energy equation (used by Dilgen, Dilgen et al. [16]) could be investigated.
- Thermo-fluid Topology Optimization
 - The thermo-fluid solver stability issues faced when implementing turbulent thermo-fluid TO are possibly in part due to weaknesses of the Spalart-Allmaras turbulence model, causing it to diverge for certain geometries created by the TO. A more robust turbulence model such as the $k - \omega$ model could be implemented (which has been proven to work with TO [15]).
 - The thermal accuracy of the Darcy force-based post-processing was only compared to one conventional post-processing method. By performing a post-processing sweep the optimality of the proposed method can be determined.
 - In this work, all simulations and TO’s were performed in 2D. The problem of optimizing the cooling interface for a power module with high power density is inherently 3D, and a 3D TO will likely provide results with significantly increased performance.
 - If 3D TO lies within reach, the applicability of 2D meta-optimized modeling parameters to a 3D TO could be investigated.

Acronyms

CFD Computational Fluid Dynamics. 4–7, 19, 21, 68

DNS Direct Numerical Simulation. 5, 6, 14

FEM Finite Element Method. iii, 4, 5, 9, 11, 12, 21, 26, 28

FVM Finite Volume Method. 5

N-S Navier-Stokes. iii, 5, 6, 13–15, 83

PDE Partial Differential Equation. 7, 8

PT Prodrive Technologies. 1, 2


RANS Reynolds-averaged Navier-Stokes. ii, iii, 6–8, 13–15

RRSE Root Relative Squared Error. ii, 28, 29, 31, 32, 37, 38, 40, 43–47, 56, 74, 75, 77–79, 81, 84

S-A Spalart-Allmaras. ii, iii, 6–8, 13–15, 58, 82

SIMP Solid Isotropic Material with Penalization. 10, 12, 16, 29, 35, 36, 38, 40, 46, 51, 74, 76–78, 82

TO Topology Optimization. i–iv, 2–5, 7–29, 36, 37, 41, 44, 45, 48–58, 80, 83



Accuracy of reference simulations

A.1. Verifying accuracy of developing external flow along flat plate reference simulation

The correlation functions that apply to the external flow along a flat plate can be categorized in two main categories, being thermal correlation functions and fluid mechanics correlation functions. A version for a laminar flow and for a fully turbulent flow is available for all the selected correlation functions and both will be discussed. For local functions the x coordinate is defined as the streamwise distance from the leading edge of the plate to the local evaluation point, as is visualized in Fig. 3.1b.

A.1.1. Fluid mechanics correlation functions.

The correlation functions for fluid mechanics correlate a quantity of mechanical behavior of the fluid to a known boundary condition. All fluid mechanics correlation functions discussed in the following paragraphs are obtained from 'Fluid Mechanics' by White [33].

1. Local Reynolds number. The local Reynolds number at coordinate x is given in Eq. (A.1).

$$\text{Re}_x = \frac{\rho U x}{\mu} \quad (\text{A.1})$$

where ρ is the fluid's density, U is the streamwise free-stream velocity of the fluid at infinite y distance from the plate and μ is the fluid's dynamic viscosity.

2. Velocity profiles. For a laminar boundary layer, an exact analytical solution for the velocity profile is defined by Blasius. This solution satisfies a nonlinear PDE and nondimensionalises the velocity and coordinate from the wall by the conversions stated in Eq. (A.2). The PDE is stated in Eq. (A.3).

$$\frac{u}{U} = f'(\eta), \quad \eta = y \left(\frac{U}{\nu x} \right)^{1/2} \quad (\text{A.2})$$

$$f''' + \frac{1}{2} f f'' = 0 \quad (\text{A.3})$$

where u is the local velocity in the boundary layer, U is the free-stream velocity, x and y are the absolute coordinates with respect to the leading edge of the plate and ν is the kinematic viscosity of the fluid. The boundary conditions for the PDE are stated in Eqs. (A.4a) and (A.4b).

$$\text{At } y = 0 : \quad f(0) = f'(0) = 0 \quad (\text{A.4a})$$

$$\text{As } y \rightarrow \infty : \quad f'(\infty) \rightarrow 1.0. \quad (\text{A.4b})$$

For a fully turbulent boundary layer no exact velocity profile exists. This boundary layer does however show a number of regions which can be correlated separately (also displayed in Fig. A.1):

1. Viscous wall layer: Viscous shear dominates.
2. Outer layer: Turbulent shear dominates.




Figure A.1: Typical velocity and shear distributions in turbulent flow near a wall: (a) shear; (b) velocity. Source: White [33]

3. Overlap- or ‘log-law’ layer: Both types of shear are important.

To analyze the boundary layer of different geometries, a set of coordinate transformations is performed as displayed in Eqs. (A.5a) to (A.5c).

$$u^+ = \frac{u}{u_\tau} \quad (\text{A.5a})$$

$$y^+ = \frac{yu_\tau}{\nu} \quad (\text{A.5b})$$

$$u_\tau \text{ or } u^* = \sqrt{\tau_w/\rho} \quad (\text{A.5c})$$

where τ_w is the local wall shear stress.

In the viscous wall layer, the ‘law of the wall’ is used to define the linear relation between u^+ and y^+ which is given in Eq. (A.6). For the overlap layer, the velocity profiles show a logarithmic relation with the wall distance. For this region, a ‘logarithmic law’ is defined which is stated in Eq. (A.7), where $\kappa = 0.41$ and $B = 5.0$. The presence of this overlap region is a clear indicator of a turbulent boundary layer.

$$u^+ = \frac{yu_\tau}{\nu} = y^+ \quad (\text{A.6})$$

$$u^+ = \frac{1}{\kappa} \ln \frac{yu_\tau}{\nu} + B = \frac{1}{\kappa} \ln y^+ + B. \quad (\text{A.7})$$

When plotting u^+ with respect to y^+ on a logarithmic axis, the different regions of the boundary layer can clearly be identified as labeled in Fig. A.2.

3. Local boundary layer thickness. From the Blasius profile, the growth of the boundary layer can be determined. The velocity boundary layer thickness δ is defined as the distance from the surface of the plate where the velocity is 99% of the free-stream velocity. The definition for the local boundary layer thickness is stated in Eq. (A.8).

$$\frac{\delta}{x} \approx \frac{5.0}{\sqrt{\text{Re}_x}}. \quad (\text{A.8})$$

For the turbulent boundary layer a ‘one-seventh-power law’ approximation is suggested by Prandtl (obtained from White [33]), which results in a local boundary layer thickness as is displayed in Eq. (A.9).

$$\frac{\delta}{x} \approx \frac{0.16}{\text{Re}_x^{1/7}}. \quad (\text{A.9})$$

Another form of the same correlation is proposed by Schlichting [28], as displayed in Eq. (A.10):

$$\frac{\delta}{x} \approx \frac{0.37}{\text{Re}_x^{1/5}} \quad (\text{A.10})$$

4. Local wall shear stress. From the Blasius profile a dimensional form of the wall shear stress can be created for laminar flow, which is stated in Eq. (A.11).

$$\tau_w(x) = \frac{0.332\rho^{1/2}\mu^{1/2}v^{3/2}}{x^{1/2}}. \quad (\text{A.11})$$

Similarly, from the power-law approximation the wall shear stress can be approximated using Eq. (A.12).

$$\tau_{w,turb}(x) \approx \frac{0.0135\mu^{1/7}\rho^{6/7}v^{13/7}}{x^{1/7}}. \quad (\text{A.12})$$




Figure A.2: Experimental verification of the inner, outer, and overlap layer laws relating velocity profiles in turbulent wall flow. Adapted from White [33]

A.1.2. Thermal correlation functions.

The thermal correlation functions all relate a thermal aspect of the characteristic external flow along a flat plate to a known boundary condition. These correlation functions are obtained from ‘Basic heat and mass transfer’ by Mills [23] unless stated otherwise.

1. Local heat transfer coefficient, constant heat flux surface. The local Nusselt number (dimensionless heat transfer coefficient) at coordinate x for a laminar boundary layer and a constant heat flux surface is given in Eq. (A.13). For a fully turbulent boundary layer, the local Nusselt number is given in Eq. (A.14).

$$\text{Nu}_x = 0.453\text{Re}_x^{1/2}\text{Pr}^{1/3}; \quad \text{Pr} > 0.5 \quad (\text{A.13})$$

$$\text{Nu}_x = 0.030\text{Re}_x^{0.8}\text{Pr}^{0.4}; \quad 0.5 < \text{Pr} < 400, \quad 5 * 10^5 < \text{Re}_x < 5 * 10^6 \quad (\text{A.14})$$

where Pr is the fluid’s Prandtl number, defined as $\text{Pr} = \mu C_p / k$. To convert the local Nusselt number into the local heat transfer coefficient, the definition of the Nusselt number is rewritten to Eq. (A.15):

$$h_{c,x} = \frac{k\text{Nu}_x}{x}. \quad (\text{A.15})$$

From a simulation, the local heat transfer coefficient can be determined using Eq. (A.16):

$$h_{c,x} = \frac{q_{in}}{T_s(x) - T_{in}} \quad (\text{A.16})$$

where q_{in} is the constant heat flux at the surface of the plate, $T_s(x)$ is the local surface temperature of the plate and T_{in} is the fluid temperature at the inlet.

2. Average heat transfer coefficient, constant heat flux surface. The average heat transfer for a flat plate of length L with a constant heat flux surface is determined using the average Nusselt number, which is given by Eq. (A.17):

$$\overline{\text{Nu}}_L = 0.680\text{Re}_L^{1/2}\text{Pr}^{1/3}. \quad (\text{A.17})$$

With the average Nusselt number, the average heat transfer is determined similar to Eq. (A.15), with Eq. (A.18):

$$\overline{h}_{c,L} = \frac{k\overline{\text{Nu}}_L}{L}. \quad (\text{A.18})$$

To compare this correlation-based heat transfer to a simulated heat transfer, the average difference between the surface temperature and the inlet temperature must first be calculated using Eq. (A.19):

$$\overline{(T_s - T_{in})} = \frac{1}{L} \int_0^L (T_s(x) - T_{in}) dx \quad (\text{A.19})$$

and use that to compute the average heat transfer coefficient using Eq. (A.20):

$$\bar{h}_c = \frac{q_{in}}{\overline{(T_s - T_{in})}} \quad (\text{A.20})$$

A.1.3. Flat plate reference simulation results

The set-up of the reference simulation based on an external developing flow along a flat plate is introduced in Section 3.1. This section discusses the accuracy of the results obtained from this reference simulation, based on a comparison with the correlation data introduced in Appendices A.1.1 and A.1.2.

Boundary condition checks The simulation domain of the flat plate reference simulation should be tall enough to approach the fluid velocity and temperature at infinite vertical distance from the plate. Furthermore, the constant heat flux boundary condition is applied to the bottom of the heat spreader domain. The heat flux at the solid-fluid interface at the top of this domain should be checked, to determine how well the constant heat flux approximation holds (and therefore how well the correlation functions for this hold). Lastly, the buffer region before the leading edge of the flat plate should be large enough to not disturb the pressure peak at the leading edge of the plate. These boundary conditions are checked in the following paragraphs.

First, the assumption of an "infinite fluid" is checked by comparing the velocity along the top boundary of the simulated domain with the inlet velocity. If these are close enough, the fluid velocity at the inlet can be used as the velocity at infinite distance from the plate. As displayed in Fig. A.3, at the lowest Reynolds number ($Re_x = 5000$) this velocity has a maximum error of 0.07% which is acceptable.




Figure A.3: Relative top boundary velocity U_{top}/U_{in} along flat plate, evaluated for various Reynolds number values

The same assumption of an "infinite fluid" is made for thermal effects. The inlet temperature is assumed to be the temperature at infinite distance from the plate. We check this assumption by looking at the relative temperature of the top boundary with respect to the inlet. Fig. A.4 displays that at the lowest Reynolds number the temperature drops to an almost immeasurable amount below the target. This is also acceptable. Therefore the domain is high enough to approximate an infinite fluid and the inlet conditions can be used as the conditions at infinite distance from the plate.




Figure A.4: Relative top boundary temperature T_{top}/T_{in} along flat plate, evaluated for various Re_L values

Secondly, a constant heat flux boundary condition is applied to the bottom of the heat spreader block which is at the bottom of the domain. For many correlations, either an isothermal wall condition or a constant heat flux wall condition needs to be applied to the cooled surface. Neither of these is the case for the reference simulation, so the best-matching scenario needs to be selected. When looking at the local wall flux at the solid-fluid interface (Figs. A.5a and A.6a), it stabilizes to the flux applied to the bottom side of the heat spreader after a quite strong peak at the leading edge of the plate. The laminar simulation however builds a consistent error (minimum of 10%) towards the trailing edge of the plate (see Fig. A.5b), while the turbulent simulation converges to a very small error of 0.35% at the trailing edge with increasing Reynolds numbers (see Fig. A.6b). The maximum deviation from the applied heat flux is visible at the outlet at a Reynolds number of 10^4 , and is a maximum of 25% for both baseline simulations.

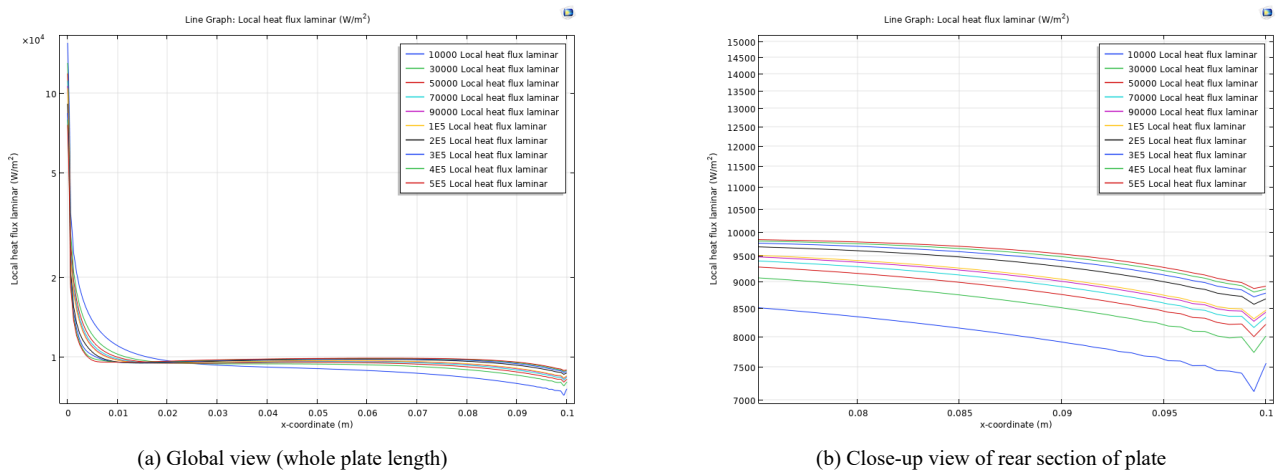


Figure A.5: Local wall heat flux at solid-fluid interface of reference simulation for varying Re_L values, laminar fluid model

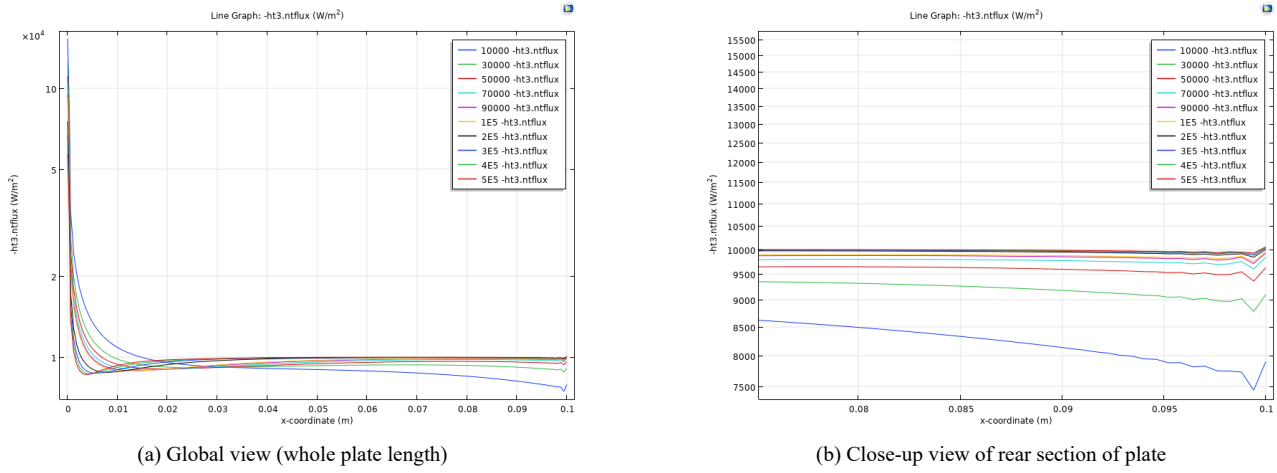


Figure A.6: Local wall heat flux at solid-fluid interface of reference simulation for varying Re_L values, laminar fluid model

The peak in local wall heat flux at the start of the plate looks extreme, as it has a magnitude of over 10 times higher than the average heat flux along the complete plate. However, the area where this peak occurs is fairly small. Therefore the influence of this peak on the average heat flux is analyzed. Fig. A.7a displays how determining the average heat flux along the rearward portion of the plate, starting at position x , the effect of including or excluding the peak can be analyzed. Figure A.7b displays the average starting at varying values of x . It is observed that at the end of the peak (manually chosen at 0.02 m), the average heat transfer drops by a maximum of 13%, meaning that the effect is less extreme than initially observed. At the highest Reynolds numbers, which are of most interest for both fluid models, the peak only alters the average by 3%.

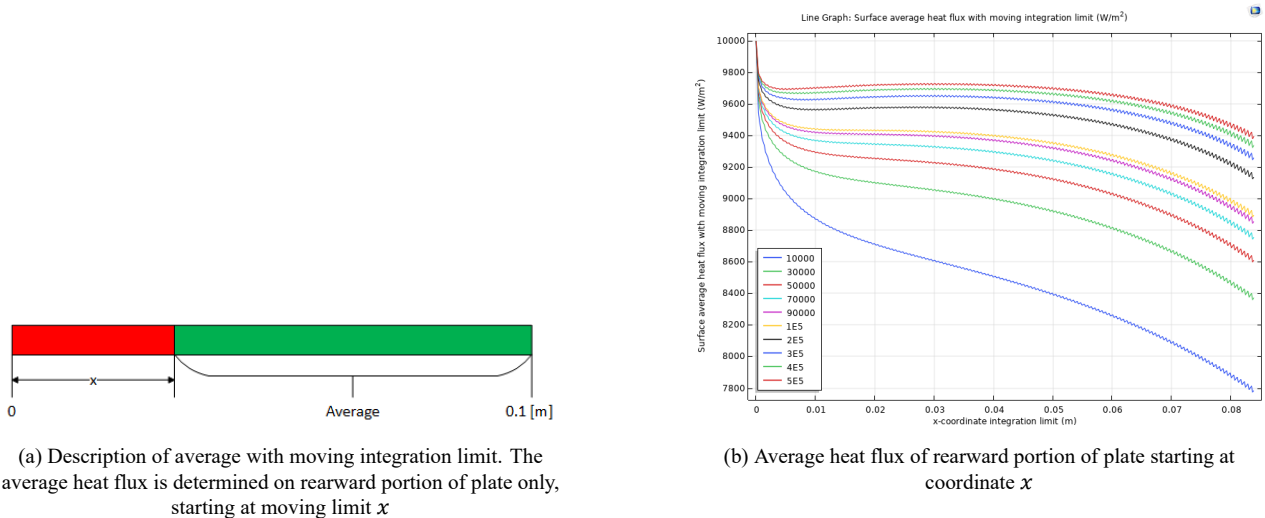


Figure A.7: Average heat flux with moving integration limit, plotted for varying Re_L values with turbulent fluid model

When, instead of the local heat flux, observing the local surface temperature at the solid-fluid interface it is observed that it continues to climb along the plate, and deviates with more than 300% when comparing the surface ΔT at the start and end of the plate. In conclusion, the wall boundary condition is much closer to a constant heat flux than a constant surface temperature and therefore the constant heat flux correlation functions will be used.




Figure A.8: Local wall ΔT with respect to inlet temperature for varying Re_L , laminar fluid model

Finally, the size of the buffer region should be checked. The leading edge of the plate creates a pressure peak which propagates upstream until it equalizes with the inlet pressure. If the inlet is placed too close to the pressure peak the gradient is forced to be unnaturally high, which suppresses the boundary layer development at the leading edge. A check is performed if the gradient propagates naturally by observing the shape of its contours in upstream direction. A natural gradient will remain oval (Fig. A.9a) where an unnatural gradient is stretched in the direction of the inlet, often touching the inlet boundary (Fig. A.9b). This effect can also be observed when checking the pressure distribution along the inlet. A simulation with proper buffer zone only exhibits a 0.5% pressure variation at the inlet at the highest Reynolds number (Fig. A.10a) where a buffer zone which is too short shows a variation of about 24% at the same Reynolds number (Fig. A.10b). We can conclude that the chosen buffer zone of 0.05 m as displayed in Fig. A.9a is large enough.

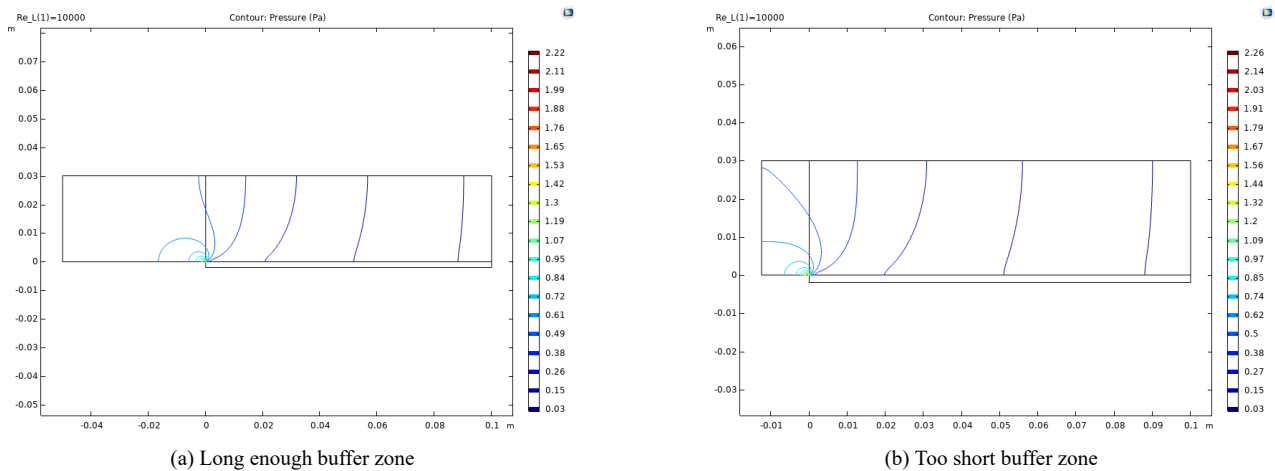


Figure A.9: Comparison of pressure contours between a long and short buffer zone, $Re_L = 10^4$

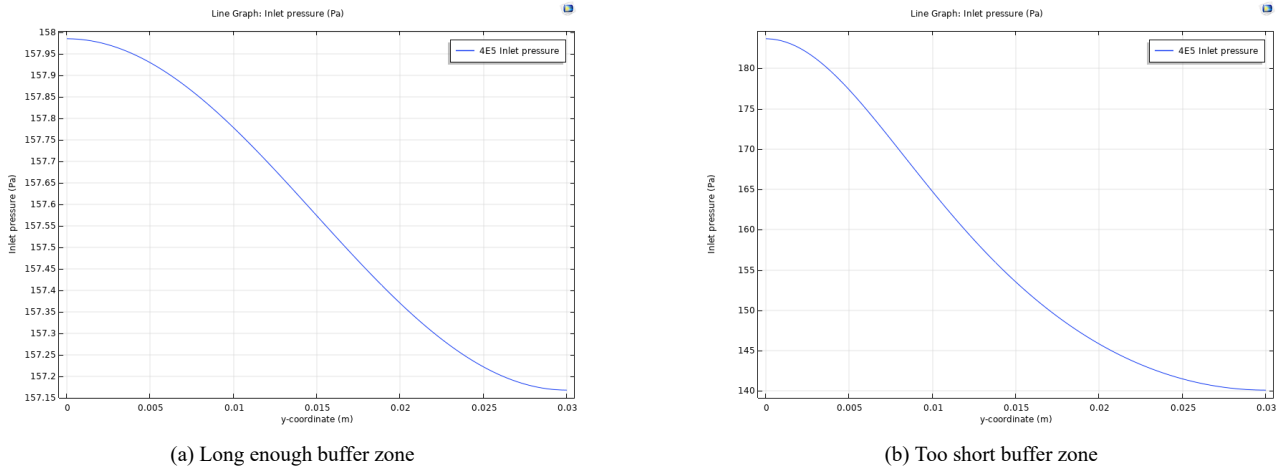


Figure A.10: Comparison of inlet pressure distribution between a long and short buffer zone, $Re_L = 4 \times 10^5$

Boundary layer thickness The development of the boundary layer of the baseline simulations (both a laminar and a turbulent fluid model) are compared to the correlations for both a laminar and turbulent boundary layer thickness at varying Reynolds numbers (Eqs. (A.9) and (A.10)). In the reference simulation, the boundary layer thickness δ is locally defined as the height of the region adjacent to the wall with velocity below 99% of the velocity at infinite distance from the plate (assumed equal to inlet velocity). Fig. A.11a displays the growth of the boundary layer along the channel at $Re_L = 10^4$. It is observed that a small region upstream from the leading edge of the plate also matches the criterion for ‘boundary layer’ (Fig. A.11b), which is an erroneous effect caused by the pressure peak at the leading edge of the plate. Furthermore the boundary layer grows in discrete steps, as it is limited in resolution by the mesh size.

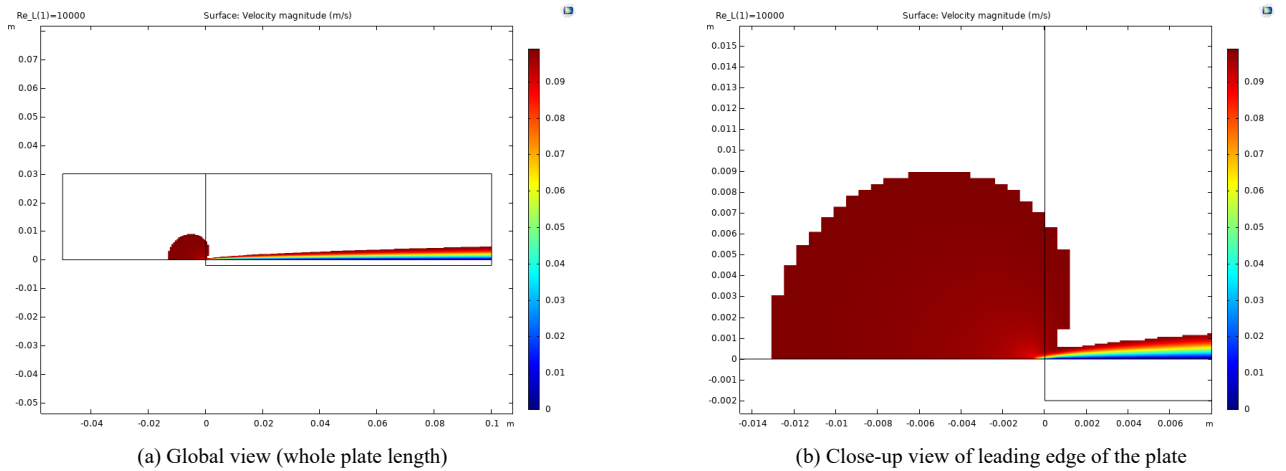


Figure A.11: Plot of velocity distribution in the boundary layer at $Re_L = 10^4$, filtered by boundary layer thickness δ definition

When comparing the boundary layer thicknesses at a Reynolds number of 5×10^3 (Fig. A.12a), the erroneous area of decreased velocity is again visible at the leading edge, and the boundary layer thickness therefore does not start at 0. It is observed that both fluid models behave very similarly and closely follow the laminar correlation. At a Reynolds number of 5×10^4 (Fig. A.12b), a clear difference is visible between the laminar and turbulent simulation, but both are still following the trend of the laminar boundary layer correlation. When further increasing the Reynolds number to 5×10^5 (Fig. A.12c), the turbulent model clearly does not follow the laminar correlation anymore while the laminar model does. The turbulent model does not follow the turbulent boundary layer correlations either. Therefore the flow is likely in a regime of transition to turbulence, which is captured by the turbulent fluid model but not by the laminar fluid model.

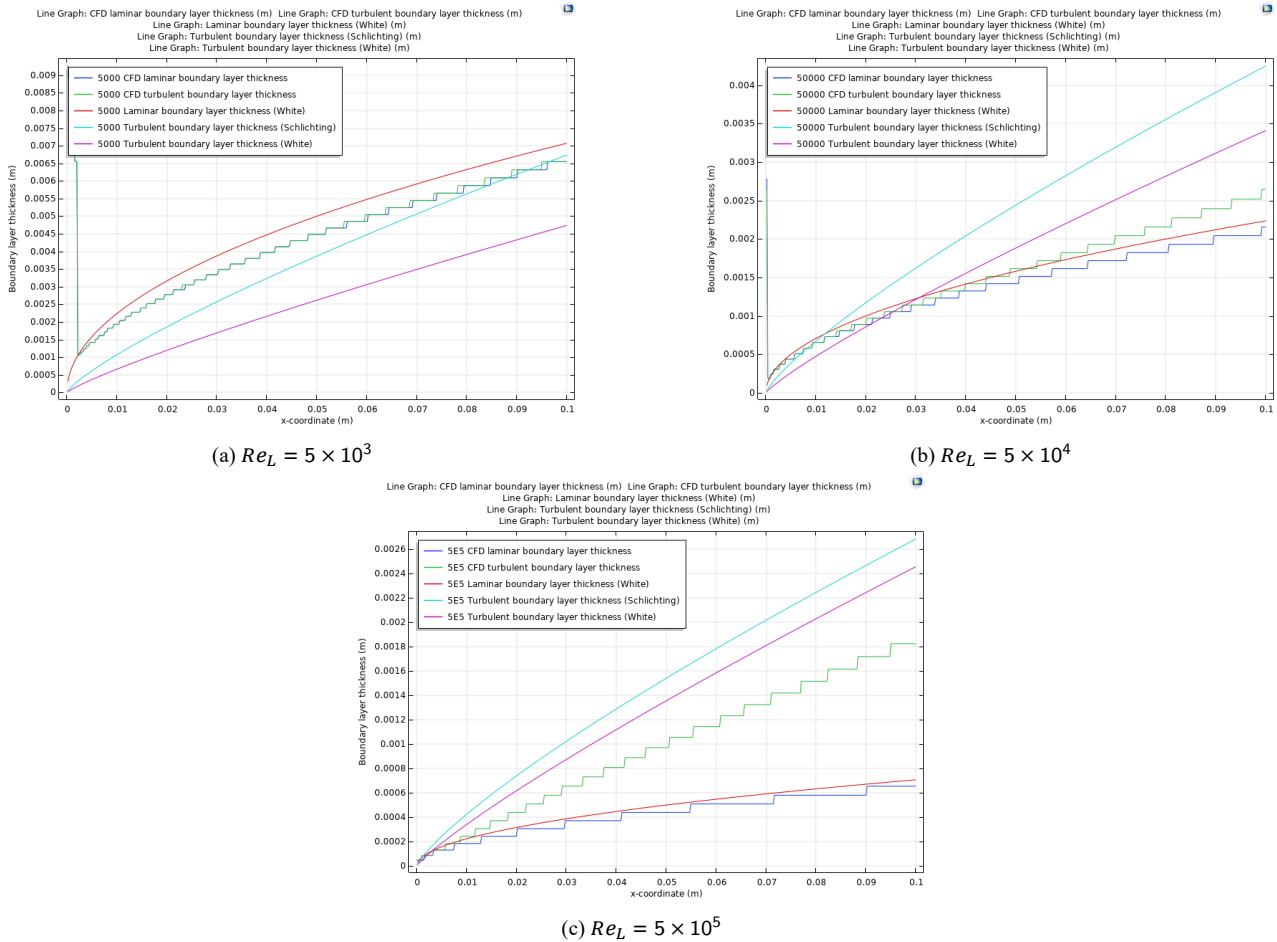


Figure A.12: Boundary layer thickness along the length of the plate at three Reynolds numbers. A comparison is made between a laminar CFD solution, a turbulent CFD solution, a correlation function for laminar flow and two correlation functions for fully turbulent flow

Wall shear stress The wall shear stress of both (laminar and turbulent) reference models can be compared to the correlations which are defined in Eqs. (A.11) and (A.12). The wall shear stress is proportional to the velocity gradient at the wall, which is a good measure for indicating turbulence (turbulent flow has a much higher velocity gradient).

When comparing the wall shear stresses obtained at a Reynolds number of 1×10^4 (Fig. A.13a), we see that both reference simulations are extremely similar and follow the trend of the laminar correlation reasonably well. A maximum error of about 6% between the turbulent simulation and the laminar correlation is visible near the trailing edge of the plate. Already at a Reynolds number of 5×10^4 transition to turbulent flow is visible with the turbulent reference (Fig. A.13b). The wall shear stress follows the laminar correlation quite well up to a certain point (≈ 0.01 m behind the leading edge) and transitions to follow the turbulent correlation with about 13% error at the trailing edge of the plate. The laminar reference simulation continues to closely follow the laminar correlation with around 1% error. At a Reynolds number of 2×10^5 the transition point of the turbulent reference to a turbulent flow is moved closer to the leading edge of the plate (Fig. A.13c). The error between the turbulent reference and the turbulent correlation is slightly reduced to around 9%. It is clear that the laminar reference will likely never deviate from the laminar correlation and therefore never show any signs of turbulence. At the highest Reynolds number (5×10^5), no transition point to turbulence can be distinguished for the turbulent reference (Fig. A.13d), as both the laminar and turbulent correlation functions show very similar values at the leading edge of the plate. The error between the turbulent reference and turbulent correlation is not significantly reduced compared to $Re_L = 2 \times 10^5$. This could indicate that the turbulent reference simulation never reaches a fully turbulent boundary layer within the simulated domain for this Reynolds number but remains within the transitional regime.

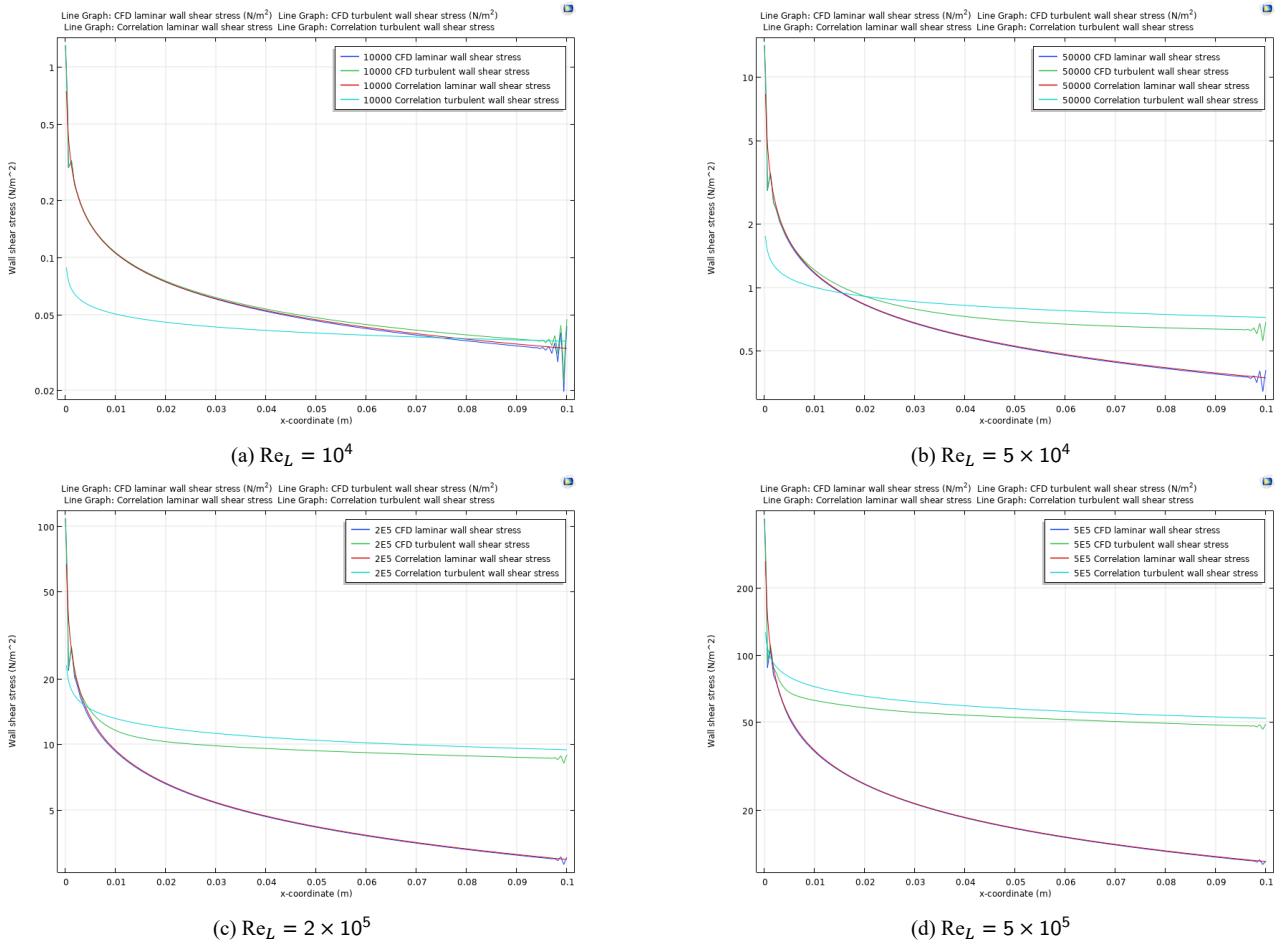



Figure A.13: Wall shear stresses of laminar and turbulent flat plate reference simulations compared to a laminar and a turbulent correlation for various Reynolds numbers

Velocity profile The Blasius velocity profile (Eq. (A.8)) can be used to verify the accuracy of the velocity profile of a laminar boundary layer. The accuracy of the turbulent boundary layer can be qualitatively checked by comparing its dimensionless velocity profile (in scale of u^+ vs y^+ , see Eqs. (A.5a) and (A.5b)) to the correlations for the viscous wall layer (Eq. (A.6)) and log-law layer (Eq. (A.7)). The velocity profiles obtained from the reference simulations are measured at the outlet of the simulated domain.

Comparing the laminar reference results to the Blasius velocity profile (Fig. A.14a), they visibly match perfectly for all Reynolds numbers simulated. This is good proof that the simulation is properly set-up (the boundary conditions are correct and a fine enough mesh is used). Figure A.14b displays that for low Reynolds numbers the turbulent fluid model is capable of reproducing the Blasius velocity profile.



(a) Dimensionless velocity profiles of laminar reference simulation compared to Blasius profile for various Reynolds numbers

(b) Dimensionless velocity profiles of laminar and turbulent reference simulations compared to Blasius profile for $Re_L = 5 \times 10^3$

When the dimensionless velocity profile (u^+ vs y^+) of the turbulent fluid model is compared to the viscous wall layer and log-law functions, it is observed that at a Reynolds number of 10^4 the velocity profile follows the viscous wall layer function quite well up to $y^+ = 1$, after which it deviates (Fig. A.15a). This shows that the Blasius velocity profile of the laminar boundary layer is not necessarily the same as the velocity profile of the viscous wall layer. At $Re_L = 7 \times 10^4$ a clear 'kink' can be seen where the reference simulation deviates further from the viscous wall layer function (Fig. A.15b). At $Re_L = 2 \times 10^5$ the reference simulation starts to follow the log-law function for a very short length (Fig. A.15c) and at $Re_L = 5 \times 10^5$ it clearly follows the log-law function for a reasonable distance (Fig. A.15d). The fact that the reference simulation does not follow the log-law layer exactly again indicates that the boundary layer is not fully turbulent at the outlet for the given Reynolds number.

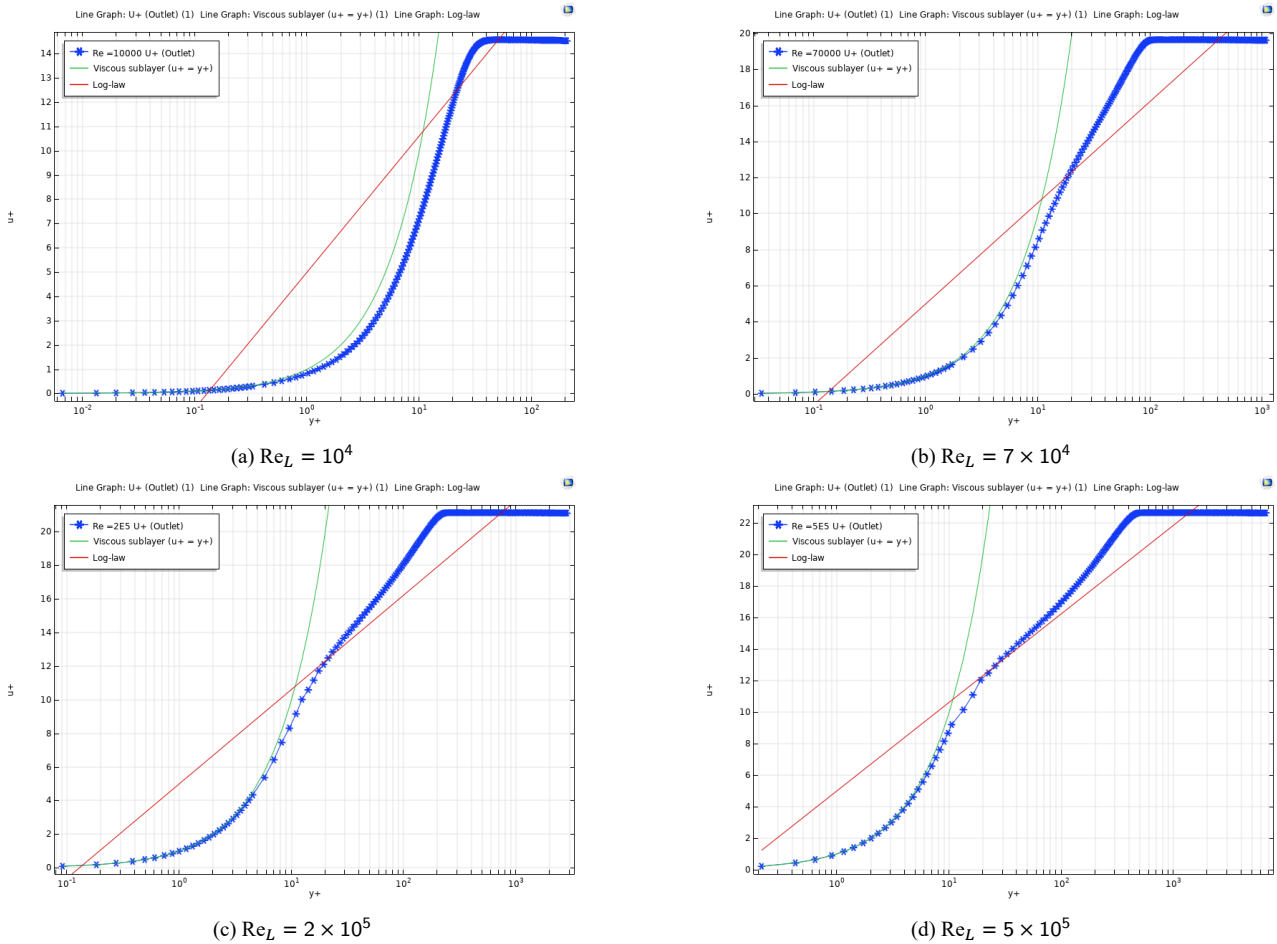




Figure A.15: Dimensionless turbulent velocity profiles compared to correlations for the viscous wall layer and log-law layer for four different Reynolds numbers

Local & average heat transfer coefficient The local and average heat transfer coefficients of the laminar and turbulent reference simulations can be compared to the correlations for a constant heat flux given in Eqs. (A.13) to (A.20). However, when observing the average heat transfer coefficient for varying Reynolds numbers (Fig. A.16a), the turbulent reference simulation diverges from the correlation with increasing Reynolds numbers. The correlations for local and average heat transfer built into COMSOL Multiphysics use a transitional Reynolds number of 6×10^5 by default, where the turbulent reference already shows transition at a Reynolds number between 5×10^4 and 10^5 . The only method of bringing this transition point up to 6×10^5 is by re-tuning the coefficients that define the turbulence model, which is beyond the scope of this project.

Another method which may be less scientifically acceptable is to alter the transitional Reynolds number of the correlation function to match that of the turbulent baseline. The best matching transitional Reynolds number when looking at the average heat transfer behavior is approximately 2×10^4 . With this solution, at least the behavior of the turbulent baseline with respect to increasing Reynolds number can be checked. Figure A.16b displays that the average heat transfer correlation continuously has a slightly higher value than the turbulent baseline, but they follow the same trend very well. The maximum error between the aforementioned models is about 22% at a Reynolds number of 5×10^4 .




(a) $Re_{tr} = 6 \times 10^5$ (default)




(b) $Re_{tr} = 2 \times 10^4$ (adapted)

Figure A.16: Average heat transfer coefficient of reference simulations compared to correlation data for varying Reynolds numbers and two different values of the correlation's transitional Reynolds number Re_{tr}


Figs. A.17a to A.17d display the local heat transfer coefficients along the plate of the laminar and turbulent reference simulations compared to the correlation data for various Reynolds numbers. It is observed in Figs. A.17c and A.17d that the transition between the laminar and turbulent correlation is quite abrupt. Before switching to the turbulent definition, the turbulent reference follows the correlation closely, and after switching it has a maximum overshoot of approximately 26%. The overshoot is caused by the correlation assuming that the boundary layer is fully turbulent after switching, where the turbulent baseline transitions over a large range of Reynolds numbers. At a Reynolds number of 5×10^5 the turbulent reference and correlation data match relatively well at the outlet. This local overshoot explains the increased average heat transfer coefficient of the correlation data compared to the turbulent reference.




(a) $Re_L = 5 \times 10^3$



(b) $Re_L = 1.99 \times 10^4$



(c) $Re_L = 1.26 \times 10^5$



(d) $Re_L = 5 \times 10^5$

Figure A.17: Local heat transfer coefficients of laminar and turbulent reference simulations compared to correlation data for various Reynolds numbers

A.2. Verifying accuracy of developing internal flow through channel reference simulation

After adapting the reference simulations of the flat plate to the channel geometry (see Section 3.1), only the correlations of a fully developed channel flow can be used to determine the accuracy of the final reference simulation.

A.2.1. Correlation functions for fully developed channel

The following correlation functions are all obtained from Mills [23]. In a tube, the Reynolds number based on the tube diameter Re_D is given by Eq. (A.21):

$$Re_D = \frac{GD}{\mu} \quad (\text{A.21})$$

where G is the mass velocity ($G = \dot{m}/A_c$, with \dot{m} the mass flow and A_c the cross-section area) and D is the tube diameter. If a channel with non-circular cross-section is used, the Reynolds number should be based on the hydraulic diameter D_h :

$$Re_{D_h} = \frac{GD_h}{\mu} = \frac{G(4A_c/\wp)}{\mu} \quad (\text{A.22})$$

where \wp is the wetted perimeter.

For a hydrodynamically fully developed laminar flow in a pipe with a smooth wall, the friction factor has a constant value given in Eq. (A.23):











$$f = \frac{64}{Re_D} \quad (\text{A.23})$$

and for a hydrodynamically fully developed turbulent flow in a pipe with a smooth wall, the friction factor f is given by Petukhov's formula in Eq. (A.24):

$$f = (0.790 \ln Re_D - 1.64)^{-2}; \quad 10^4 < Re_D < 5 \times 10^6. \quad (\text{A.24})$$

If the channel is not tubular and the flow is laminar, the product of friction factor and Reynolds number based on the hydraulic diameter fRe_{D_h} is obtained from Table A.1. In the selected scenario of a channel flow, this factor is 96. If a turbulent flow through a non-circular channel is used, the Reynolds number based on the hydraulic diameter can be inserted into Eq. (A.24) without conversion. This can be done "because the viscous sublayer around the perimeter of the duct is very thin, and the velocity and temperature are nearly uniform across the core fluid. Since the viscous sublayer is the major resistance to momentum and heat transfer, the precise shape of the core fluid is not critical" [23].

Table A.1: "Nusselt numbers and the product of friction factor times Reynolds number for fully developed laminar flow in ducts of various cross-sections." [23]

Cross Section	Nu_{D_h}		fRe_{D_h}
	Constant Axial Wall Heat Flux	Constant Axial Wall Temperature	
 Equilateral triangle	3.1	2.4	53
 Circle	4.364	3.657	64
 Square	3.6	2.976	57
 Rectangle 1 1.4	3.8	3.1	59
 Rectangle 1 2	4.1	3.4	62
 Rectangle 1 3	4.8	4.0	69
 Rectangle 1 4	5.3	4.4	73
 Rectangle 1 8	6.5	5.6	82
 Rectangle 1 ∞	8.235	7.541	96
 Heated ∞ Insulated	5.385	4.861	96

From a simulation, the skin friction coefficient can be easily obtained by using its definition given in Eq. (A.25):

$$C_f = \frac{\tau_s}{(1/2)\rho V^2} \quad (\text{A.25})$$

where τ_s is the wall shear stress, ρ is the fluid density and V is the bulk velocity (here equal to the inlet velocity). To compare C_f from the simulation to f from the correlation, the two can be related by:

$$f = 4C_f \quad (\text{A.26})$$

when the flow is hydrodynamically fully developed.

Correlation functions are also available for the heat transfer of a *thermally* fully developed flow. However, these are not included in the analysis, as the channel required to fully develop a flow thermally from a uniform inlet velocity is too long to simulate within realistic times. Furthermore, the correlation functions for this are often given only for constant wall temperatures.

A.2.2. Channel reference simulation results

The friction factors f given for hydrodynamically fully developed flow by Eqs. (A.23) and (A.24) are converted to skin friction coefficients C_f such that they can be compared to the skin friction coefficients obtained by the laminar and turbulent reference simulations. Fig. A.18a displays that the laminar reference stabilizes to a skin friction coefficient that is 0.4% lower than that of the laminar correlation. Figure A.18b displays that the turbulent reference simulation shows some oscillations along the channel length, before stabilizing to a value that is 2.5% higher than its respective correlation. The boundary layer is not fully turbulent as indicated in Section 3.2 and the Reynolds number of 6000 lies below the applicable range of Eq. (A.24). These two factors can explain the larger error of the turbulent simulation compared to the laminar simulation. Both cases are accepted to be sufficiently accurate for further use in this research.




Figure A.18: Skin friction coefficients of reference simulations compared to the correlations for fully developed channel flow

B

Blurred wall mitigation method applied to laminar flow

This appendix discusses the application of mitigation 1 (blurred walls mitigation method) to a laminar flow with $Re_{D_h} = 100$. First, in Appendix B.1 the optimization problem is introduced and after that, in Appendix B.2 the results of this mitigation method are discussed.

B.1. Laminar optimization problem

For the laminar blurred walls mitigation method, the following modeling parameters are considered as design variables for the meta-optimization procedure (refer to Table 3.2 for their definition and range):

- SIMP interpolation penalty p
- Darcy interpolation penalty q
- Projection filter slope β
- Vertical position of density gradient y_{grad} .

Similar to the turbulent mitigation method (Section 5.1), the Darcy number and the solid-fluid transition length L_{trans} are excluded from the design variables as they are expected to have a monotonic effect on the modeling accuracy. The objective of this meta-optimization is the average RRSE of the global velocity and temperature. A sweep is performed for the Darcy numbers of 10^{-4} , 10^{-5} , 10^{-6} and 10^{-7} . This range spans higher Darcy numbers than those used for the turbulent mitigation methods, because from previous numerical experiments it is observed that a higher Darcy number can be used for a laminar mitigation method to achieve similar results as a turbulent mitigation method. This might be caused by the laminar flow carrying less momentum than a turbulent flow and therefore requiring a lower Darcy force / Brinkman penalty to slow it down. However, the Brinkman penalty already includes the fluid's velocity (Section 2.4.1) and therefore automatically scales with the flow velocity. Therefore it is not completely clear why these increased Darcy numbers can be used. Furthermore, a laminar density-based fluid model seems to become unstable at higher Darcy numbers compared to the turbulent density-based fluid model. This may be caused by increased disturbances with lower Darcy numbers, which can cause local turbulent eddies to initiate. These eddies cannot be fully resolved if the mesh is too coarse (it is practically impossible to use a mesh that is fine enough to fully resolve these eddies) and therefore they cause instabilities in the fluid solver.

Compared to the turbulent mitigation methods, the channel for this laminar mitigation method is shortened to 10 mm (from 100 mm, defined in Section 3.2) to reduce computational cost and to preserve the detailed view of the developing section of the boundary layer (the laminar flow develops faster than the turbulent flow). Unfortunately, when reducing the domain, the streamwise element size remained constant which resulted in few elements in this direction. Therefore the obtained results may be coarser than desired. The blurring filter radius R_{min} (Eq. (2.17)) is set to 0.3 mm and with that the solid-fluid transition length L_{trans} (Eq. (2.18)) becomes 2.08 mm.

B.2. Laminar results

In this section, the results of the blurred walls mitigation method (laminar mitigation 1) applied to a laminar flow are discussed. First, a graphical comparison will be made based on a number of simulation outputs. After that the resulting RRSE values are discussed. Finally, the resulting modeling parameter values are discussed.

Comparison of graphical accuracy: Figs. B.1a to B.1d compare the results of the mitigation method to the laminar reference simulation. In Fig. B.1a the heat transfer coefficients of all mitigation method results show an inverted region near the leading edge of the channel wall, but the result of Darcy number 10^{-4} shows inverted heat transfer nearly along the whole channel length. In Fig. B.1b the temperature difference is seen to match better with the reference as the Darcy number decreases. The wall shear stresses displayed in Fig. B.1c show data that counteract the previous observation, because the mitigation results match the reference better with *increasing* Darcy numbers. Figure B.1d displays that the velocity through the solid region decreases significantly with decreasing Darcy numbers. The results with Darcy numbers of 10^{-6} and 10^{-7} achieve nearly identical heat transfer coefficients, temperature differences and wall shear stresses, indicating that a point of diminishing returns is present when decreasing the Darcy number.

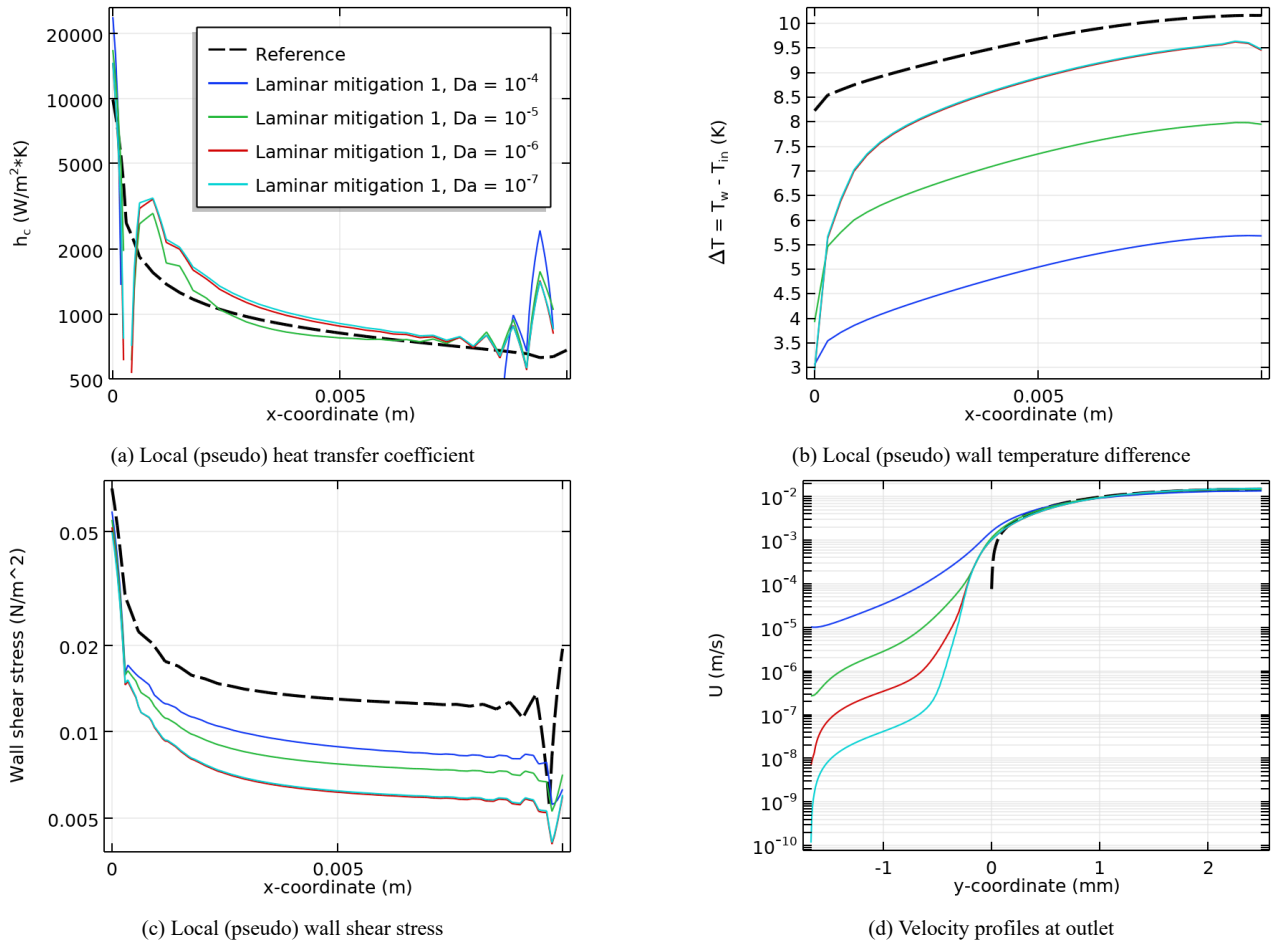


Figure B.1: Comparison of laminar mitigation 1 for four Darcy number values to the reference simulation. Pseudo boundary results are obtained following methodology discussed in Section 3.5

Comparison of numerical accuracy: The RRSE values obtained by laminar mitigation 1 are displayed in Table B.1. Here the observation of increasing thermal accuracy with decreasing Darcy numbers is reinforced because the temperature RRSE shows this trend as well. The velocity RRSE shows a minimum at the Darcy number of 10^{-5} , which is somewhat consistent with the wall shear stresses displayed in Fig. B.1c. This again indicates a shift from dominating errors caused by the liquid metal and porous solid material effects at higher Darcy numbers to the error caused by the stagnant fluid effect being dominant at lower Darcy numbers. The average RRSE is shown to decrease with decreasing Darcy numbers (Table B.1), but a point of diminishing returns is again visible.


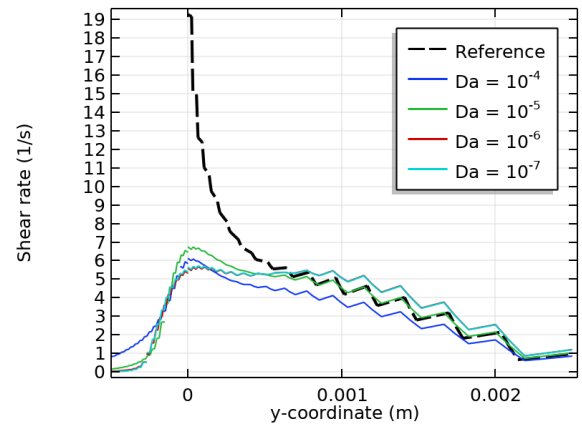
Table B.1: RRSE data of laminar blurred walls mitigation method

Darcy number	RRSE Temperature	RRSE Velocity	Average RRSE
10^{-4}	87.23%	19.55%	53.39%
10^{-5}	57.25%	4.94%	31.09%
10^{-6}	17.92%	5.52%	11.72%
10^{-7}	7.67%	7.22%	7.44%

Resulting modeling parameters: The resulting modeling parameter values from laminar mitigation 1 are displayed in Table B.2. Interestingly, for Darcy numbers of 10^{-4} down to 10^{-6} , the optimizer converges to the upper limit for the Darcy interpolation penalty q and to the lower limit for the SIMP interpolation penalty p . This is a strong indication that the liquid metal effect is highly dominant at these Darcy numbers and can be partly mitigated by matching the interpolation functions as good as possible (both functions are linear at these penalty values). Switching the interpolation of the thermal material properties from the SIMP interpolation to the Darcy interpolation can possibly improve the overall accuracy of the laminar density-based fluid model. It might still be advantageous to use separate SIMP interpolation penalties for the Darcy force interpolation and the interpolation of material properties. Furthermore, a trend is visible in Table B.2 of the relative vertical gradient position y_{grad}/L_{trans} . When observing Fig. B.2a, it is visible that this gradient position is used to align the Darcy force profiles of the different Darcy number settings. The alignment of the Darcy force profiles matches relatively closely with the Darcy force threshold (determined to lie at 4×10^4 Pas/m² for this flow condition) at which the shear rate peaks. The Darcy force profiles of the results with Darcy numbers of 10^{-6} and 10^{-7} are nearly identical from the pseudo-boundary ($y = 0$) upwards, even though they have completely different modeling parameter values. The shear rate profiles displayed in Fig. B.2b are also identical between the Darcy numbers of 10^{-6} and 10^{-7} . This reinforces the observation that a point of diminishing returns is present when reducing the Darcy number.


Table B.2: Resulting modeling parameters of laminar blurred walls mitigation method

Darcy number	y_{grad}/L_{trans}	q	p	β
10^{-4}	-2.42%	10.00	1	4.65
10^{-5}	-7.94%	10.00	1	1.67
10^{-6}	-22.06%	10.00	1	5.84
10^{-7}	-24.35%	0.21	1	4.93

(a) Darcy force profiles. The Darcy force threshold of 4×10^4 Pas/m² is displayed in black

(b) Shear rate profiles compared to reference

Figure B.2: Comparison of laminar mitigation 1 results for selected Darcy numbers




Analyzing effect of modeling parameters on modeling accuracy

This appendix discusses a ‘black-box’ approach to analyzing the effects of the modeling parameters on the accuracy of the laminar density-based thermo-fluid model. An older version of the analysis framework is used to measure the modeling accuracy of the density-based fluid model for randomly generated combinations of a subset of the modeling parameters for a laminar flow, as given in Table C.1. Because this method requires solving the verification simulation many times, this analysis is performed on a laminar flow model to limit computation times. The older version of the analysis framework uses a simplified, linear gradient of the density variables instead of the realistic gradient which is filtered by the blurring and projection filters. Therefore the solid-fluid transition length L_{trans} and the projection filter slope β are also excluded from this analysis. The outdated version of the framework may not provide results that are as accurate as the final version, but it is believed that the conclusions drawn from this black-box analysis are still valid. The goal of this analysis is to determine for each modeling parameter whether or not it has a monotonic or convex effect on the accuracy. This behavior gives an indication which parameters can be used for optimization of the modeling accuracy and which parameters should be specified manually.


Table C.1: Modeling parameters used for laminar Monte-Carlo simulation

Variable	Symbol	Definition	Relevant range	Linear/logarithmic parameter
SIMP interpolation penalty	p	Eq. (2.12)	2 - 5	linear
Darcy interpolation penalty	q	Eq. (2.13)	0.01 - 10	logarithmic
Darcy number	Da	Eq. (2.30)	10^{-10} - 10^{-5}	logarithmic
Vertical position of density gradient	y_{grad}	Section 3.5	-	linear

The random generation of parameter combinations is performed by a Monte-Carlo simulation which ensures a uniform distribution of samples throughout the user-selected domain. The Monte-Carlo simulation outputs a 4-dimensional point cloud of samples, with each sample having an average RRSE value. Because of the high dimensionality of the dataset, it cannot be visualized at once. When plotting the RRSE with respect to one parameter, some parameters clearly show a certain behavior type where others do not. Figure C.1a shows that the Darcy number generally has a monotonically increasing effect on the RRSE and Fig. C.1b shows convex behavior of the gradient position. Furthermore, it is known that the transition length has a monotonic effect on the RRSE. The SIMP interpolation penalty and Darcy interpolation penalty do not show any distinctive global behavior and therefore monotonicity and global convexity can be eliminated for them. Because of the monotonicity of the Darcy number and the transition length, the ‘modeling parameter optimization’ problem is not bounded when these parameters are included as a variable. Therefore the Darcy number and the transition length are considered as user-selected parameters.




(a) Average RRSE with respect to Darcy number, Da . A monotonically increasing error with increasing Darcy numbers is identified




(b) Average RRSE with respect to relative vertical gradient position, y_{grad}/L_{trans} . L_{trans} is fixed. Global convexity of the error with respect to the gradient position is identified

Figure C.1: Unfiltered scatter plots of data points generated with a Monte-Carlo simulation, visualizing global behavior of the average RRSE with respect to selected modeling parameters

By applying a narrow range to the Darcy number and removing data points that lie outside of this range, the behavior of the remaining ‘free’ parameters can be analyzed locally at the selected Darcy number. The same can be applied to any combination of the modeling parameters, further filtering the dataset. In some cases a higher resolution is required locally and an extra Monte-Carlo simulation is performed with one or multiple fixed modeling parameters. From such an extra simulation at a fixed Darcy number of 10^{-7} and by filtering the relative vertical gradient offset to lie within a range of 0 to 30%, a locally monotonic behavior is visible of the Darcy interpolation penalty (Fig. C.2a). However, a clear point of ‘diminishing returns’ is visible where the decrease of the RRSE levels off. This is understandable when looking at the interpolation curve of the Darcy interpolation function (Fig. 2.4b), as the function becomes nearly linear for penalty values above 10. Interestingly, the point of diminishing returns shifts with both the Darcy number and the gradient position. In contrast, the SIMP interpolation penalty does not show a clear behavior (Fig. C.2b).




(a) Average RRSE with respect to Darcy interpolation penalty, q . Local monotonicity is identified.




(b) Average RRSE with respect to SIMP interpolation penalty, p . No behavior type is distinguished

Figure C.2: Scatter plots generated using fixed Darcy number (10^{-7}) and filtered gradient offset (0 to 30%) to visualize local behavior of the average RRSE with respect to the modeling parameters

Furthermore, it is observed that the optimal gradient position depends on the Darcy number selection and appears to converge to 0 as the Darcy number decreases (Figs. C.3a and C.3b). At higher Darcy numbers the convexity of the RRSE with respect to the gradient position changes into a behavior with a local minimum around 0 offset and a decreasing error with increasingly positive gradient positions. The decrease in error levels off at relative gradient positions y_{grad}/L_{trans} of approximately 75%. This suggests that for these high Darcy numbers the Darcy force is not high enough, even at low density values (nearly solid), to generate an accurate boundary layer.




(a) Dataset filtered to include samples with $Da = 1 \times 10^{-5}$ to $1 \times 10^{-5.4}$. A local minimum is identified around 0 offset, and a decreasing error which levels off towards 75% relative gradient position



(b) Dataset filtered to include samples with $Da = 1 \times 10^{-7.6}$ to 1×10^{-8} . A global optimum is identified

Figure C.3: Scatter plot of the average RRSE with respect to the vertical gradient position y_{grad}/L_{trans} filtered at different ranges of the Darcy number

These observations partly support the hypothesis that a (local) optimum can be found within the set of modeling parameter combinations. The problem is deemed a feasible candidate for meta-optimization, but it is expected to encounter issues with local minima and drift of parameters for which the accuracy has low sensitivity.



Overview of density method variables for thermo-fluid Topology Optimization

80

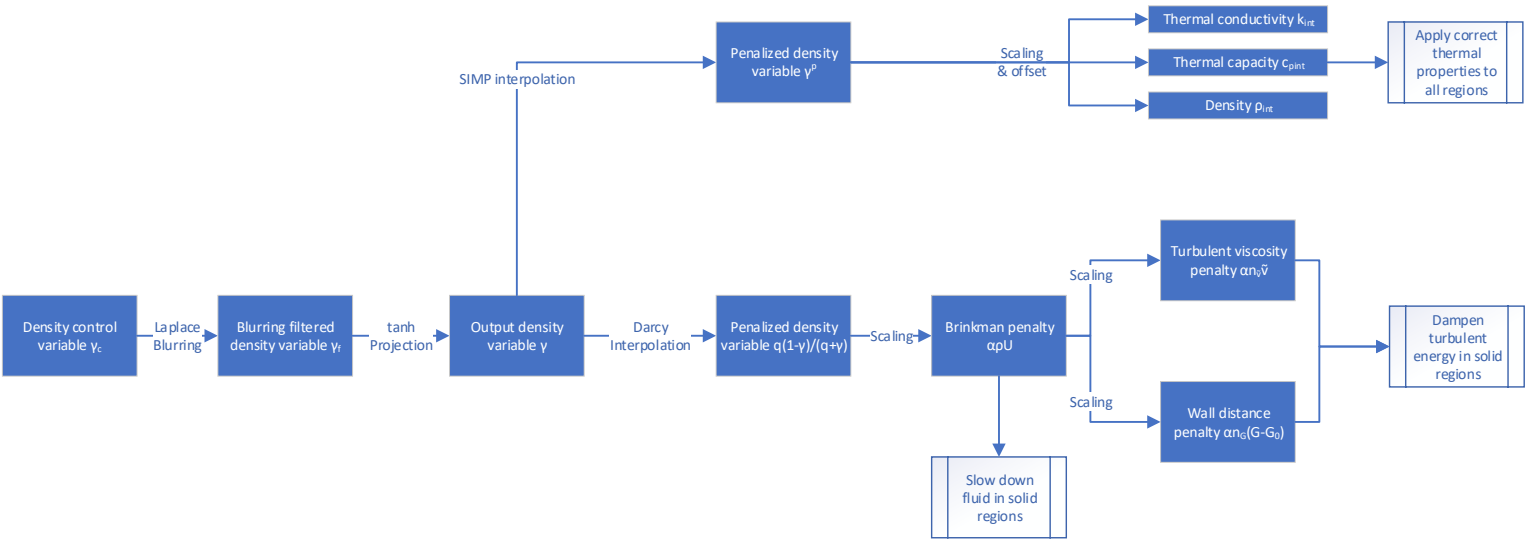



Figure D.1: Flowchart of density method variables



Set-up of turbulent U-channel test case

This appendix will discuss the set-up of the turbulent U-channel test case in detail. A reference simulation with a conventional turbulent fluid model and sharp walls is first executed. After that, a density-based simulation is performed of the same geometry, by manually specifying the density variables to form the shape of the U-channel. The modeling error is here defined as the RRSE values of the density-based result compared to the reference. To set up the reference and density-based simulations, the following aspects are relevant and will be discussed in the following sections:

1. Dimensions of geometry (Appendix E.1)
2. Boundary conditions (Appendix E.2)
3. Turbulence model (Appendix E.3)
4. Material properties and interpolations (Appendix E.4)
5. Mesh (Appendix E.5)
6. Stabilization (Appendix E.6)
7. Measure of accuracy (Appendix E.7)

E.1. Dimensions

The shape of the U-channel geometry and the parametric dimensions are displayed in Fig. E.1 and the values of the dimensions used for the reference simulation as well as the density-based simulation are given in Table E.1.

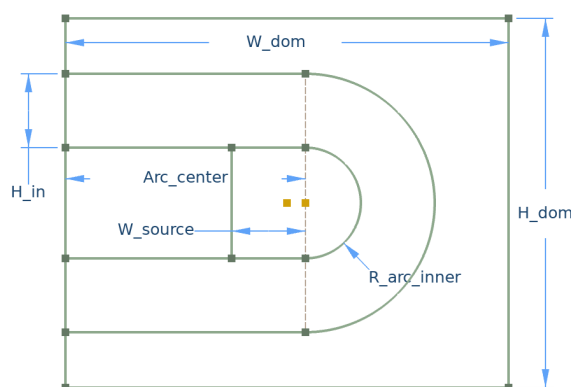


Figure E.1: Geometry and dimensions of U-channel simulation. Dimensions are given in Table E.1

For the density-based simulation in COMSOL, the density variables can be manually specified by drawing the ‘pseudo-boundaries’ where they are fixed at $\gamma = 0.5$. In the quantitative analysis framework (straight channel geometry), the density variables are defined similarly, with a variable y_{grad} to specify the vertical position of the density variable gradient with respect to the reference. Because the simulations in the framework only contained half of the channel (with a symmetry condition at the top boundary), the full channel height (H_{in}) has to be offset by $-2y_{grad}$ to match the alignment of the mitigation method results. Because the simulation is 2D, the fluid channel can be considered as an infinite out-of-plane extrusion, and therefore the hydraulic diameter of this channel is $2H_{in}$. For the energy equation an out-of-plane thickness of 10 mm is selected, such that total heat flows can be used instead of 2D heat fluxes.

Table E.1: Dimensions of turbulent U-channel simulations. The reference dimensions and the altered dimensions for the density-based model are listed

	W_dom	H_dom	H_in	R_arc_inner	Arc_center	W_source
Reference	30 mm	25 mm	5 mm	3.75 mm	16.25 mm	5 mm
Density-based	30 mm	25 mm	5 mm $- 2y_{grad}$	3.75 mm $+ y_{grad}$	16.25 mm	5 mm

E.2. Boundary conditions

Figure E.2 displays the boundary conditions for the reference simulation of the turbulent U-channel. The blue domain is fluid, the gray domains are conductive solid and the red domain is conductive solid with a volumetric heat source. A fully developed turbulent velocity profile with average velocity U_{in} is specified at the inlet (upper left) such that $Re_{D_h} = 6000$. At the outlet (lower left), the pressure is forced to 0. Along the walls of the channel, a no-slip condition is applied. A uniform inlet temperature $T_{in} = 40^\circ\text{C}$ is specified and thermal outflow is applied to the outlet. A heat flow of 60 W is applied uniformly to the heated (red) domain. The external boundaries of the domain (apart from the inlet and outlet) are adiabatic. In the density-based simulation, the no-slip boundary condition is removed as there are no sharp walls with this model. The boundary behavior is reproduced by interpolating the density variables from solid to fluid.




Figure E.2: Boundary conditions of U-channel reference simulation

E.3. Turbulence model

The density-based fluid model uses the S-A turbulence model with penalization for the eddy viscosity and wall distance, as specified in Section 2.4.2. However, it is expected that the U-channel geometry triggers some inaccuracies with this turbulence model (mostly because of detached flow after the U-bend). Therefore a different turbulence model is selected for the reference simulation which does not show these inaccuracies for the U-channel. The Wilcox revised two-equation $k - \omega$ model with realizability constraints is selected for its robustness in internal flows, geometries that exhibit strong curvature and detached flows [3]. The $k - \omega$ turbulence model is a low-Reynolds model, similar to the S-A model, and therefore requires a high mesh resolution and a good initial guess of the solution.

E.4. Material properties and interpolations

The same material properties are used as with the other turbulent simulations in this research, which are stated in Table 3.1. The reference simulation applies these properties to their respective domains. The density-based simulation applies material properties which are interpolated between solid and fluid based on the density variables. The thermal material properties (thermal conductivity, thermal capacity and density) are all interpolated with the same SIMP interpolation function (Eq. (2.12)). This is essentially the same as interpolating the thermal diffusivity of the material. The turbulent fluid is slowed down in solid domains by interpolating the Darcy force / Brinkman penalty (Eq. (2.32)) with a Darcy interpolation (Eq. (2.13)). The turbulent kinetic energy is dampened in solid domains by applying a penalty to the eddy viscosity (Eq. (2.33)) and the wall distance (Eq. (2.34)) which are based on the Darcy force.

E.5. Mesh

The reference simulation and the density-based simulation use separate mesh strategies. First the reference mesh is discussed, after which the density-based mesh is introduced. The mesh of the reference simulation uses a pre-defined ‘physics controlled mesh’ generation by COMSOL, set to the ‘Extremely fine’ setting. This generates a free triangular mesh with rectangular boundary layer elements near the wall, of which a close-up is displayed in Fig. E.3a. The resolution is checked by determining the y^+ value of the first velocity evaluation point adjacent to the wall. This value is found to be $y^+ = 0.04$ which is well within the acceptable range (the rule-of-thumb for a low-Reynolds number turbulence model is $y^+ < 1$). However, the boundary layer thickness varies greatly along the length of the U-channel and with that the y^+ of the first evaluation point may become significantly higher. The y^+ values along the reference channel walls are displayed in Fig. E.3b. From this it is observed that the highest y^+ of the first evaluation point is 0.2, which is still within acceptable range. The mesh inside solid domains is significantly coarser than that in the fluid domain, as the heat conduction is linear and can therefore be easily resolved.

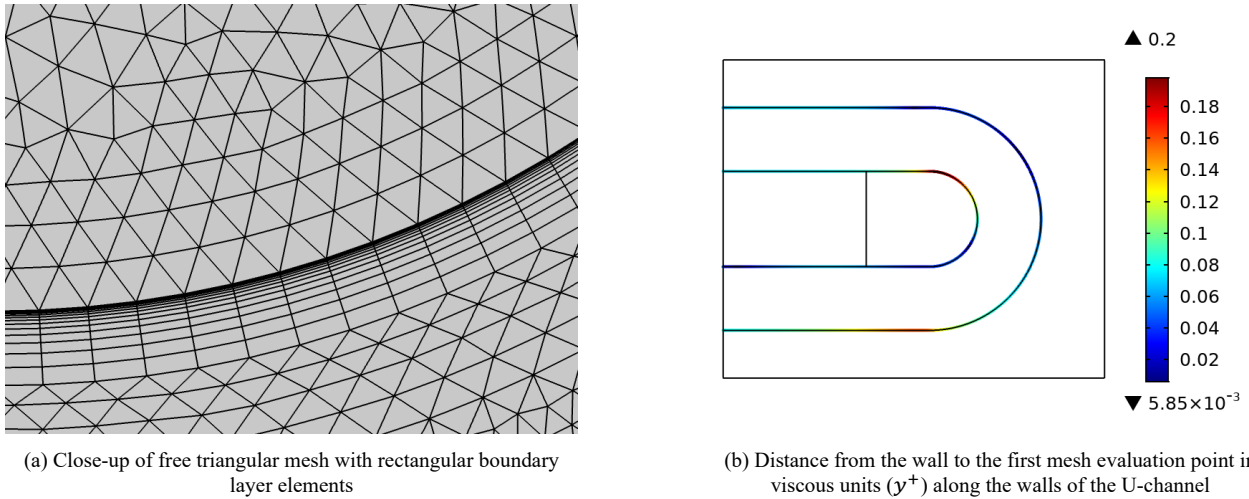



Figure E.3: Definition and viscous resolution of the mesh used for the turbulent U-channel reference simulation


The mesh of the density-based simulation is also based on a free triangular mesh with rectangular boundary layer elements, as is displayed in Fig. E.4a. However, it needs to meet different resolution requirements. Firstly, evaluating the mesh resolution of the density-based simulation in viscous units (y^+) is not possible, as there is no wall to compute the wall distance in viscous units with. By basing the mesh resolution on the resolution of the reference, a good estimation can be made on the achieved density-based resolution. Secondly, the mesh of the density-based simulation should also meet the requirement set by the blurring filter that $h < R_{min}$, where h is the element height and R_{min} the blurring filter radius. In this case, to match the blurring filter radius used for the mitigation methods ($L_{trans} = 1$ mm and therefore $R_{min} = 0.14$ mm), the mesh resolution is selected to be 0.1 mm which is lower than the base size of the reference mesh. This mesh resolution is equal both in the fluid and in the solid domains, such that the blurring filter requirement is met throughout the whole domain. Furthermore, mesh refinements are used near the pseudo-boundary to achieve similar resolution at the pseudo-boundary as the reference. Note that this would not be possible in a density-based TO, as the walls migrate through the domain and the mesh remains stationary. These refinements are added to reduce the computational time compared to a uniform mesh resolution equal to the resolution at the pseudo-boundary. Lastly, in contrast to the reference simulation, the refinements are placed on both sides of the pseudo-boundary for the density-based simulation, because the fluid can penetrate into the solid domain which is known to cause instabilities if not properly resolved. Figure E.4b shows the element size along the left vertical external boundary of the simulation domain. The dips indicate areas of refinement. Here, the constant mesh size of 0.1 mm of the density-based model in both the fluid and solid regions is visible, whereas the reference simulation shows varying mesh sizes in these domains. The density-based mesh shows refinement on both sides of the pseudo-boundary to a similar resolution as the reference, whereas the reference only shows the refinement inside the fluid domain. By meeting all three mesh resolution requirements (filter radius, y^+ in the fluid, resolving penetrating fluid inside solid region) ensures a reasonable convergence of the density-based fluid model.

E.6. Stabilization

Both the reference and density-based simulations use consistent stabilization (streamline, crosswind) for the N-S, the turbulence equations and the energy equation. For the fluid equations this stabilization is required because linear discretization is used for all fields.



(a) Close-up of free triangular mesh with rectangular boundary layer elements both in the fluid and the solid domain



(b) Mesh size of density-based and reference models along the left vertical external boundary. The dips indicate mesh refinement regions

Figure E.4: Mesh of density-based simulation

E.7. Measure of accuracy

The same average RRSE measure as with quantitative analysis framework is used to gauge the modeling accuracy of the density-based fluid model in the turbulent U-channel. The RRSE values of the velocity, temperature and eddy viscosity are computed by comparing the density-based data to the reference data. The average RRSE is defined as the unweighted average of the three aforementioned RRSE values.

Bibliography

- [1] FEM vs. FVM | COMSOL Blog. URL <https://www.comsol.com/blogs/fem-vs-fvm/>.
- [2] Understanding Stabilization Methods | COMSOL Blog. URL <https://www.comsol.com/blogs/understanding-stabilization-methods/>.
- [3] Which Turbulence Model Should I Choose for My CFD Application? | COMSOL Blog. URL <https://www.comsol.com/blogs/which-turbulence-model-should-choose-cfd-application/>.
- [4] Joe Alexandersen and Casper S. Andreasen. A Review of Topology Optimisation for Fluid-Based Problems. *Fluids*, 5(1):29, 3 2020. ISSN 2311-5521. doi: 10.3390/fluids5010029. URL <https://www.mdpi.com/2311-5521/5/1/29>.
- [5] Steven R. Allmaras, Forrester T. Johnson, and Philippe R. Spalart. Modifications and Clarifications for the Implementation of the Spalart-Allmaras Turbulence Model. In *Seventh International Conference on Computational Fluid Dynamics (ICCFD7)*, pages 1–11, Big Island, Hawaii, 7 2012. URL <http://ae.metu.edu.tr/tuncer/ae546/docs/ICCFD7-1902.pdf>.
- [6] Martin P. Bendsøe and Noboru Kikuchi. Generating optimal topologies in structural design using a homogenization method. *Computer methods in applied mechanics and engineering*, 71:197–224, 4 1988. doi: 10.1016/0045-7825(88)90086-2.
- [7] Martin P. Bendsøe and Ole Sigmund. *Topology optimization. Theory, methods, and applications*. Springer-Verlag Berlin Heidelberg GmbH, 2 edition, 2004. ISBN 978-3-642-07698-5. doi: 10.1007/978-3-662-05086-6.
- [8] Rajesh Bhaskaran and Lance Collins. Introduction to CFD Basics. Technical report, Cornell University-Sibley School of Mechanical and Aerospace Engineering, 2002. URL <https://dragonfly.tam.cornell.edu/teaching/mae5230-cfd-intro-notes.pdf>.
- [9] Thomas Borrvall and Joakim Petersson. Topology optimization of fluids in Stokes flow. *International Journal for Numerical Methods in Fluids*, 41(1):77–107, 1 2003. ISSN 02712091. doi: 10.1002/flid.426.
- [10] Hendrik C. Brinkman. A calculation of the viscous force exerted by a flowing fluid on a dense swarm of particles. *Applied Scientific Research*, 1(1):27–34, 12 1949. ISSN 03657132. doi: 10.1007/BF02120313.
- [11] COMSOL. CFD Module User’s Guide. Technical report, 2018. URL <https://doc.comsol.com/5.4/doc/com.comsol.help.cfd/CFDModuleUsersGuide.pdf>.
- [12] COMSOL. Optimization Module User’s Guide. Technical report, 2018. URL www.comsol.com/blogs.
- [13] Ercan M. Dede. Multiphysics Topology Optimization of Heat Transfer and Fluid Flow Systems. In *Proceedings of the COMSOL Users Conference*, Boston, 2009. URL <https://www.comsol.pt/paper/download/44388/Dede.pdf>.
- [14] Ercan M. Dede. Optimization and design of a multipass branching microchannel heat sink for electronics cooling. *Journal of Electronic Packaging, Transactions of the ASME*, 134(4):041001–1, 12 2012. ISSN 10437398. doi: 10.1115/1.4007159.
- [15] Cetin B. Dilgen, Sumer B. Dilgen, David R. Fuhrman, Ole Sigmund, and Boyan S. Lazarov. Topology optimization of turbulent flows. *Computer Methods in Applied Mechanics and Engineering*, 331:363–393, 4 2018. ISSN 00457825. doi: 10.1016/j.cma.2017.11.029.
- [16] Sumer B. Dilgen, Cetin B. Dilgen, David R. Fuhrman, Ole Sigmund, and Boyan S. Lazarov. Density based topology optimization of turbulent flow heat transfer systems. *Structural and Multidisciplinary Optimization*, 57(5):1905–1918, 5 2018. ISSN 16151488. doi: 10.1007/s00158-018-1967-6.
- [17] Alberto Donoso and Ole Sigmund. Topology optimization of multiple physics problems modelled by Poisson’s equation. *Latin American Journal of Solids and Structures*, 1(2):169–184, 3 2004. ISSN 1679-7825. URL <https://www.lajss.org/index.php/LAJSS/article/view/45>.

- [18] Adriano A. Koga, Edson C. C. Lopes, Helcio F. Villa Nova, Cícero R. de Lima, and Emílio C. N. Silva. Development of heat sink device by using topology optimization. *International Journal of Heat and Mass Transfer*, 64:759–772, 2013. ISSN 00179310. doi: 10.1016/j.ijheatmasstransfer.2013.05.007.
- [19] Andrej N. Kolmogorov. Local structure of turbulence in an incompressible viscous fluid at very large Reynolds numbers. *Doklady Akademii Nauk SSSR*, 30(4):299–301, 1941. URL <https://cds.cern.ch/record/739746>.
- [20] Boyan S. Lazarov and Ole Sigmund. Filters in topology optimization based on Helmholtz-type differential equations. *International Journal for Numerical Methods in Engineering*, 86(6):765–781, 5 2011. ISSN 00295981. doi: 10.1002/nme.3072.
- [21] Hao Li, Xiaohong Ding, Fanzhen Meng, Dalei Jing, and Min Xiong. Optimal design and thermal modelling for liquid-cooled heat sink based on multi-objective topology optimization: An experimental and numerical study. *International Journal of Heat and Mass Transfer*, 144, 12 2019. doi: 10.1016/j.ijheatmasstransfer.2019.118638.
- [22] Danny J. Lohan, Ercan M. Dede, and James Allison. Topology Optimization Formulations for Circuit Board Heat Spreader Design. In *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Reston, Virginia, 6 2016. AIAA. ISBN 978-1-62410-439-8. doi: 10.2514/6.2016-3669.
- [23] Anthony F. Mills. *Basic heat and mass transfer, Pearson new international edition*. Pearson Education Limited, second edition, 11 2013. ISBN 9781292042480.
- [24] Laurits H. Olesen, Fridolin Okkels, and Henrik Bruus. A high-level programming-language implementation of topology optimization applied to steady-state Navier-Stokes flow. *International Journal for Numerical Methods in Engineering*, 65(7):975–1001, 2 2006. ISSN 00295981. doi: 10.1002/nme.1468.
- [25] Carsten Othmer. A continuous adjoint formulation for the computation of topological and surface sensitivities of ducted flows. *International Journal for Numerical Methods in Fluids*, 58(8):861–877, 11 2008. ISSN 02712091. doi: 10.1002/flid.1770.
- [26] Osborne Reynolds. IV. On the dynamical theory of incompressible viscous fluids and the determination of the criterion. *Philosophical Transactions of the Royal Society of London. (A.)*, 186:123–164, 12 1895. ISSN 0264-3820. doi: 10.1098/rsta.1895.0004.
- [27] Christopher Rumsey. Turbulence Modeling Resource, 2019. URL <https://turbmodels.larc.nasa.gov/>.
- [28] Hermann Schlichting and Klaus Gersten. *Boundary-layer theory*. Springer-Verlag Berlin, 9th edition, 2016. ISBN 978-3-662-52917-1. doi: 10.1007/978-3-662-52919-5.
- [29] Philippe R. Spalart and Steven R. Allmaras. One-equation turbulence model for aerodynamic flows. In *30th Aerospace Sciences Meeting & Exhibit*, number 1, pages 5–21, Reno, NV, 1 1992. AIAA. doi: 10.2514/6.1992-439.
- [30] Thomas B. van der Hout. Topology Optimization methods for cooling of a SiC transistor Power Module. Technical report, TU Delft, Eindhoven, 11 2019.
- [31] Theodore von Kármán. Mechanical similitude and turbulence. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen - Fachgruppe I (Mathematik)*, 5:58–76, 1930. URL <https://resolver.caltech.edu/CaltechAUTHORS:20140804-121813080>.
- [32] Fengwen Wang, Boyan S. Lazarov, and Ole Sigmund. On projection methods, convergence and robust formulations in topology optimization. *Structural and Multidisciplinary Optimization*, 43(6):767–784, 6 2011. ISSN 1615147X. doi: 10.1007/s00158-010-0602-y.
- [33] Frank M. White. *Fluid Mechanics*. McGraw-Hill, 7th edition, 2011. ISBN 978-0-07-352934-9.
- [34] Gil H. Yoon. Topology optimization for turbulent flow with Spalart-Allmaras model. *Computer Methods in Applied Mechanics and Engineering*, 303:288–311, 5 2016. ISSN 00457825. doi: 10.1016/j.cma.2016.01.014.
- [35] Gil H. Yoon. Topology optimization method with finite elements based on the k- ϵ turbulence model. *Computer Methods in Applied Mechanics and Engineering*, 361, 4 2020. ISSN 00457825. doi: 10.1016/j.cma.2019.112784.
- [36] Radyadour Kh. Zeytounian. Joseph Boussinesq and his approximation: A contemporary view. *Comptes Rendus - Mécanique*, 331(8):575–586, 8 2003. ISSN 16310721. doi: 10.1016/S1631-0721(03)00120-7.