Multi-Objective Design Optimization of Electric Vehicle Battery Cooling Plates Considering Thermal and Pressure Objective Functions

by

Anthony Jarrett

A thesis submitted to the
Department of Mechanical and Materials Engineering
in conformity with the requirements for
the degree of Master of Applied Science

Queen's University
Kingston, Ontario, Canada
September 2011

Abstract

The current stimuli of climate change and rising oil prices have spurred the development of hybrid electric (HEV), and battery electric vehicles (BEV): collectively termed EVs. However, the battery technology needs much development: at the time of writing, the range of a BEV is too low to be practical in many situations. A critical limitation is the sensitivity of batteries to temperature: the heat generated during operation affects their performance and reduces the lifetime.

This study investigates battery cooling using cooling plates: thin rectangular fabrications inserted between battery cells. A coolant pumped through internal channels absorbs heat and transports it away from the battery. Previous studies of liquid heat exchangers have indicated that the geometry of the channels plays a significant role in the performance; however, there is a lack of rigorous numerical optimization applied to EV cooling plates.

By developing a numerical optimization framework utilizing parametric geometry generation and computational fluid dynamics, this research has investigated the characteristics of optimum cooling plate geometry with respect to three objectives: average temperature, temperature uniformity, and coolant pressure drop. By applying each objective separately, improvements of up to 70% have been made compared

to a reference design. The influence of boundary conditions on performance and optimum design has been assessed, and multi-objective optimization has investigated the trade-off between competing objective functions.

Although care should be taken when extrapolating the results beyond the geometry and conditions in the study, some general design principles can be proposed. Objectives of average temperature and pressure drop can both be satisfied by a common design with wide cooling channels, but different characteristics are needed for temperature uniformity. Additional assessments have revealed that optimizations of temperature uniformity are especially sensitive to the boundary conditions, whereas the other objective functions are largely insensitive.

The optimization process developed in this work can be applied to any potential cooling plate design and will lead to gains in the targeted performance measure. In doing so, the performance of the EV will be incrementally improved, thereby advancing the day when an EV is not only an environmental choice, but also a practical choice.

Acknowledgments

First and foremost, I would like to thank my supervisor, Dr Il Yong Kim, for his insightful advice and steady guidance. Not only have I learned much from him, I have also found my research to be interesting and enjoyable.

Secondly, my thanks go to all the faculty and staff whom I have encountered in the Department of Mechanical and Materials Engineering. Whether teaching, research advice, or the labour of scholarship applications, they have always been there at the right time.

The other students in the Structural and Multidisciplinary Systems Design Laboratory have been a constant source of advice, support, entertainment and friendship. All of them—past and present—have my warmest thanks.

My family and friends—who now span two continents—have been always encouraging and enthusiastic; reassuring and practical; loving and supportive—in whatever measure was required. Whatever the next chapter will bring, I look forward to sharing it with them.

This research was funded by AUTO21, a member of the Network of Centres of Excellence of Canada Program. From our industrial partner, General Motors of Canada, I would like to thank Carlton Fuerst, Derrick Chow, Justin Gammage, Brian Tossan and colleagues for their input and information.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

The recent combination of rising oil prices and environmental concerns has led to unprecedented pressure from consumers and government for the automotive industry to develop viable alternatives to oil-based combustion engines. While hybrid electric vehicles (HEVs) have found a niche in the mainstream auto market, battery electric vehicles (BEVs) still have several hurdles to overcome before they find widespread appeal. One of the primary technical obstacles is the battery capacity and performance: batteries of a practical size and cost cannot currently compete with the vehicle range provided by gasoline and an internal combustion engine. The heat generated during operation has a significant effect on the performance and durability of the battery, and as technological development increases the energy density of battery cells for both HEVs and BEVs (collectively EVs), this becomes a more critical factor.

Although low temperatures also reduce battery performance, the more pressing

problem is cooling the battery: high temperatures increase the rate of corrosive reactions so limiting the battery lifetime. In addition, temperature variations within a cell lead to non-uniform rates of battery discharge meaning that the available energy cannot all be readily accessed thereby reducing the driving range. The most effective cooling systems require the active circulation of a liquid coolant, and this consumes power. As an EV is a closed system, any power needed to operate a cooling system directly reduces the power available to the other vehicle systems, including the motive power. Therefore, a further feature of the cooling system is that its operation should require the minimum amount of power possible.

The implementation of an improved battery cooling system has the potential to increase battery life, and increase the useful energy capacity of the battery. These are currently two of the principal drawbacks of a BEV, and principal performance indicators of a HEV; addressing them will go some way to making all EVs more practical and desirable.

1.2 Objective

For the most efficient use of space, EV battery packs are often configured as a stack of thin rectangular battery cells. In this arrangement, one method of regulating the temperature of the pack is to use cooling plates inserted between battery cells. These are thin metal fabrications with internal channels through which a coolant flows; heat from the cells is transferred into the coolant and transported away from the battery. Studies of similar devices used for high-powered electronics and fuel cells have indicated that the geometry of the cooling channels can have a significant influence on the cooling performance. However, there is a lack of rigorous design

optimization work that can be readily applied to the design of a battery cooling plate.

The problem defined in this study is the numerical optimization of a single EV battery cooling plate. Using a combination of previous academic studies, and technical information from General Motors (GM) of Canada, a reference cooling plate is defined and its geometry characterized parametrically. Using the analysis technique of computational fluid dynamics (CFD), the performance of the plate with respect to heat transfer, temperature uniformity, and power consumption is assessed. The parametric geometry generation and CFD analysis is incorporated into a gradient-based optimization algorithm thereby allowing the optimum plate geometry to be determined for any given combination of objective function and boundary conditions.

Using this optimization routine, the optimum cooling plate designs are derived for objective functions of minimum average temperature, temperature uniformity, and power consumption. Variation in the optimum designs at different boundary conditions is investigated, and multi-objective optimization is used to define the trade-off between competing objectives and define optimum designs at intermediate objective functions. A feature of battery cells is that their heat generation is greatest in the vicinity of the battery connections—this boundary condition is investigated by assessing different cooling plate orientations and deriving optimum designs with different heat generation distributions.

The objectives of the study are as follows:

1. Develop a numerical model that can predict the thermal behaviour of a cooling plate.

2. Develop a verified procedure for the optimization of an individual cooling plate.

3. Demonstrate that single objective optimization can improve the design of a cooling plate.

4. Perform multi-objective optimization of a cooling plate to determine trade-offs between single objective designs.

5. Determine the effect of changing boundary conditions—on plate performance, and the optimum design.

6. Determine general design principles that can be used to improve the performance of a cooling plate.

1.3 Contributions

Previous studies have assessed fuel cell cooling plates and applied numerical optimization to the design of electronics cooling plates, but this study represents the first numerical optimization of a cooling plate with the dimensions and boundary conditions corresponding to an EV battery pack. It demonstrates that numerical optimization can be applied to a relatively complex geometry to produce significant improvements in performance. The results indicate that for the geometry utilized in this study—a single serpentine channel—the objective functions of cooling performance and power consumption are closely aligned; a cooling plate can be designed that satisfies both performance criteria, albeit at the expense of temperature deviation.

The guidelines developed in this study will contribute to the intelligent design of cooling plates by allowing the desired behaviour to be targeted: either overall heat transfer, or temperature uniformity. In addition, the assessment and optimization of

the cooling plate at a range of boundary conditions, coupled with a multi-objective optimization study further expands the scope of applicable optimum designs.

This is the first study to assess the influence of a non-uniform heat source on the performance and optimum design of a cooling plate. Guidelines have been developed indicating the preferable orientation of cooling plates with respect to the heat source depending on the thermal requirements. Furthermore, the demonstration that optimization must be carried out with the appropriate heat source distribution in order to achieve the optimum designs is a finding that can be applied to multiple disciplines.

Finally, there is a significant contribution for future researchers within the Structural and Multidisciplinary Systems Design Lab in the form of a verified and versatile optimization algorithm. Future studies can use this as a starting point for investigations with more elaborate modelling, different boundary conditions, or alternative geometries.

Chapter 2

Background

2.1 Global Trends in Automotive Vehicles and Energy Use

Global production of road vehicles in 2010 was approximately 77.6 million (which was a 26% increase on the previous year) [1]. In total, there were 840 million cars in operation in 2010, forecast to rise to 1.4 billion by 2030 [2]. This increase is almost entirely attributed to developing countries.

At the time of writing, road vehicles are almost exclusively powered by oil-based fuels, which are a significant contributor to greenhouse gas emissions. Transport (of which road vehicles make up around 77%) accounts for 26% of global energy use, and 23% of global greenhouse gas emissions [3]. Although oil consumption per vehicle is forecast to decline annually by around 1.6% [2], in the absence of major technological shifts, the total consumption of oil by road vehicles will rise significantly in the period to 2030.

The price of oil in recent years has been extremely volatile, starting the 21st century at around $20 (US$ per barrel) and showing a general increase throughout most of the decade to a highpoint of $147 in July of 2008. However, by the end of the year, the price had collapsed to less than $35, before once more beginning a steady rise to exceed $100 in the first quarter of 2011 [4]. The oil market is influenced by a unique combination of economic, political and environmental factors [5], and the reasons behind the price fluctuations of 2008–2011 seem to be a combination of supply/demand issues coupled with political unrest and financial speculation [6].

Despite having relatively stable periods in 2009 and 2010, the price of oil should be expected to remain unpredictable. Furthermore, with the exception of Canada, all of the largest estimated reserves of oil occur in countries with some degree of political volatility. The Canadian reserves (and similar deposits in Venezuela) are concentrated in non-conventional reserves such as oil sands that are expensive and environmentally damaging to extract [2]. Much of the conventional oil reserves outside of the middle-east are offshore reserves in deep water, and the 2010 accident at the Deepwater Horizon rig in the Gulf of Mexico [7] highlights the potential risks involved with deep water oil extraction.

The increasing price volatility and move towards less accessible reserves are to some extent indicative of the finite nature of the resource. Estimations of remaining oil supplies are difficult to make, as new reserves are frequently being discovered, but the onset of 'Peak Oil'—the point at which oil production begins to decline—has been forecast for many years [8].

Even before the discussion of climate change, there are clearly compelling reasons to reduce the dependence on oil.

2.2 Anthropogenic Climate Change

Since the beginning of the 20th century, global surface temperatures have steadily increased, by between 0.56 and 0.92 degrees in the 100 years to 2005. The rate of increase in the last 50 of these years has been almost twice that for the 100 year period. Sea level rises of 2–3 mm per year have been observed—consistent with the recorded temperature rise—due to thermal expansion and increased meltwater. Significant decreases in the extent of arctic and glacial ice have been observed, and seasonal snow cover has decreased in both hemispheres. Globally, patterns of rainfall have changed, with some regions seeing increases, while others have seen decreases. Extreme weather such as hurricanes, flooding and drought has become more common. [9]

These changes in climate and temperature correlate with an increase in emissions of gases, collectively known as greenhouse gases (GHGs), following the industrial revolution. These gases collectively have an effect of absorbing heat radiated from the earth's surface and re-emitting it, thus reducing the amount of heat lost into space. Of the naturally occurring gases, those with the most significant warming effect are water vapour, carbon dioxide (CO_2), methane (CH_4), nitrous oxide (N_2O), and ozone (O_3) [10]. With the exception of water vapour, these have all seen significant atmospheric increases due to human activity: CO_2 has increased from a pre-industrial concentration of 280 ppm to 379 ppm in 2005; CH_4 has more than doubled in concentration to 1774 ppb in 2005; N_2O has increased by 18% to 319 ppb in 2005; ozone is a short-lived gas that displays regional variation, but the increase in warming effect is estimated to be almost as great as from CH_4. [9]

In addition to the effect of naturally occurring gases, there are many artificial gases

that have a significant warming effect, the largest group being the chlorofluorocarbons (CFCs). Despite occurring in very low concentrations, their warming effect can be thousands of times greater than CO_2, and this leads to a warming contribution of around 12% of the total. Most of these gases have been banned under international agreements, but they remain in the atmosphere for many years before breaking down.

Climate modelling is extremely difficult, and future behaviour can be influenced by a myriad of factors. Nonetheless, the IPCC (Intergovernmental Panel on Climate Change) has produced a set of forecasts corresponding to different scenarios. The estimates of global temperature rise over the course of the 21st century range from 1.1 to 6.4 degrees, with the best estimates falling in the 1.8–4.0 range. Corresponding sea level rises are generally in the range of 0.2 m to 0.5 m. The changes brought about by such a major change in the climate are considered by many to be extremely damaging.

The principal factor accounting for future climate change is the continued emission of carbon dioxide. With the rapid industrialization of developing nations, the increase of atmospheric CO_2 shows no sign of slowing, and indeed, measurements indicate that it is accelerating [11]. Despite scientific consensus that CO_2 emissions should be reduced, governments are reluctant to impose stringent measures that could prove economically damaging, and so the technologies and infrastructure to enable emissions reduction on a large scale have yet to be put into practice.

A worldwide shift in automotive technology away from gasoline fuels has the potential to significantly reduce global CO_2 emissions. The idea of an electric car is not new, but the current need for change coupled with technical advances may mean that its time has come.

2.3 Electric Vehicles

2.3.1 History

Despite its long-time image as car-of-the-future, electric vehicles were actually sold in their thousands before internal combustion (IC) engine vehicles came to dominate the market in the early 20th century. Larminie [12] gives a succinct description of the early days of the electric car:

> The first electric vehicles of the 1830s used non-rechargeable batteries. By the end of the 19th century, with mass production of rechargeable batteries, electric vehicles became fairly widely used. Private cars, though rare, were quite likely to be electric, as were other vehicles such as taxis. Indeed if performance was required, the electric cars were preferred to their internal combustion or steam powered rivals. [In 1899] the electric vehicle known as 'La Jamais Contente' set a new land speed record of 106 kph.
>
> At the start of the 20th century electric vehicles must have looked a strong contender for future road transport. The electric vehicle was relatively reliable and started instantly, whereas internal combustion engine vehicles were at the time unreliable, smelly and needed to be manually cranked to start. The other main contender, the steam engine vehicle, needed lighting and the thermal efficiency of the engines was relatively low.
>
> By the 1920s several hundred thousand electric vehicles had been produced for use as cars, vans, taxis, delivery vehicles and buses. However,

> despite the promise of the early electric vehicles, once cheap oil was widely available and the self starter for the internal combustion engine had arrived, the IC engine proved a more attractive option for powering vehicles. Ironically, the main market for rechargeable batteries has since been for starting IC engines.

For the remainder of the century, IC engine vehicles made up the overwhelming majority, due to their superior performance and cost. Electric vehicles remained in use for some niche applications for which low speed and limited range were not critical shortcomings; their quiet operation, zero emissions, and painless start-stopping has made them useful in some situations such as the British milk float and golf buggies. However, up until the 1990s, their re-adoption into the mainstream car market was unsuccessful, primarily because their performance and price was no match for the IC engine.

However, the growing environmental, economic and political movement to reduce the reliance on gasoline powered vehicles has prompted a shift in policy makers, manufacturers, and consumers to address some of the shortfalls in electric vehicle technology and produce a viable market and product. This can most clearly be seen in the success of the Toyota Prius: the integration of an electric battery with an IC engine has created a vehicle with a driving experience comparable to that of a regular IC engine vehicle, but with improved fuel economy. Local government subsidies have offset the higher cost of the car, and among environmentally-minded consumers it has become a fashionable product. Over 1.6 million have been sold as of 2010 [13].

2.3.2 Vehicle Technology

Although there are several categories of electric vehicle including fuel cell powered, and vehicles using external power sources, this discussion is limited to vehicles utilizing an electric battery. These vehicles can be further subdivided into those that are solely powered from the battery: BEVs (battery electric vehicles), and those that also utilize an additional energy source: HEVs (hybrid electric vehicles).

BEVs operate with a very simple architecture: a battery, charged from an external source, supplies electricity to one or more motors that drive the wheels. Because the torque generated by an electric motor is independent of speed, BEVs can operate without gear shifting, and it is often advantageous to connect individual motors directly to the wheels. Furthermore, energy can be recovered during braking by having the motors act as generators and charging the battery. 2010 saw the launch of the first modern BEV from a major manufacturer: the Nissan Leaf [14]. It is advertised with a top speed of 145 km/hour and a range of 117 km; adequate for urban driving, but the limited battery capacity and lengthy recharging make it impractical for long journeys.

Until the capacity of EV batteries can be significantly increased, HEVs will represent a practical compromise: by adding an IC engine and gasoline tank the range can be extended while retaining some of the efficiency benefits. Different methods of linking the conventional and electric drivetrain, and charging the battery give rise to further categories of HEV. For example, a parallel hybrid can be driven either by electric motor or IC engine; a serial hybrid is driven exclusively by the electric motor, with the IC engine used to charge the battery via a generator. In either configuration, there is the option of charging the battery from an external source, i.e.

a plug-in vehicle; or restricting it to electricity generated on-board by the generator and regenerative braking. In the latter, as the energy is still coming exclusively from gasoline, improvements in fuel economy and emissions reductions are limited; only when plugged in to an external source of electricity is there the possibility to dramatically reduce gasoline usage and tail-pipe emissions. The method of generating the electricity in the first place will dictate the total emissions reduction that can be attributed to replacing gasoline with battery power.

2010 saw the US launch of the Chevrolet Volt [15], a plug-in series hybrid that has been designed to be powered solely from the battery for the first portion of a journey before starting an IC engine to replenish the battery. From fully charged, it can travel 56 km on battery power, at which point the IC engine will be started. Only 10% of Canadians commute more than 25 km to get to work [16], and a study in the UK [17] found that 64% of vehicle CO_2 emissions arise from journeys of less than 40 km. So despite the limited electric range, the Volt still has the capability to significantly influence vehicle oil consumption and emissions. As battery technology is improved, its electric range can be increased, so that for an increasing number of journeys it will operate essentially as a BEV.

2.4 Batteries

A battery is a device that stores chemical energy and releases it as electrical energy. It is comprised of two electrodes joined by an electrolyte, and it is the chemical reactions that take place between these components that generates electricity. In the case of a 'primary cell', the battery components are consumed during use; however, in a 'secondary cell' the reaction can be reversed by applying a current to the cell,

thereby recharging the battery.

The most familiar rechargeable battery is possibly the lead-acid battery that is already found in most vehicles to supply electricity for the starter motor and ancillary devices. However, while this battery technology is cheap and robust, it does not provide an efficient method of storing energy. Advances in battery technology in the last 20 years, in part driven by the demand for portable consumer electronics, have made available a variety of cell chemistries of increasing performance.

The 2000 Toyota Prius HEV used a 53 kg Nickel-Metal-Hydride (NiMH) battery pack with a capacity of 1.8 kWh (kilowatt-hours) [18]. At the time of writing, the current state of the art is the 2010 Chevrolet Volt HEV that uses a 190 kg lithium-ion battery pack with a capacity of 16 kWh [19].

Although to some extent different battery types have different behaviour, this study is not intended to be specific to any one battery chemistry or configuration: the technical challenges associated with one type of battery are often equally applicable to other batteries. It is not the purpose of this study to assess the particular mechanisms and chemical processes at work within battery cells, but it is necessary to understand some of the basic performance measures and characteristics.

2.4.1 Battery Characteristics

One of the basic characteristics of a battery cell is the potential difference (voltage) developed across its terminals. Each cell has a nominal voltage that describes the voltage that can be expected when the cell is delivering power. However, this voltage is influenced by many factors, not least the internal resistance of the cell; i.e. as the current is increased, more voltage is expended internally, and so the terminal voltage

decreases. The terminal voltage is also affected by the state of charge of the battery; i.e. the voltage decreases as the battery charge decreases.

Conventional terminology is to define the 'capacity' of a battery by the duration for which it can supply an electric current; i.e. a measure of charge. Although the SI unit for this is the coulomb (the charge transported by one amp flowing for one second), it is more convenient to use a larger unit, the amp-hour, defined as one amp flowing for one hour. Thus, a battery may be said to have a capacity of 50 amp-hours; able to supply a current of 1 amp for 50 hours. This implies that it can supply a current of 2 amps for 25 hours, however the capacity is dependent on the discharge rate, and at higher currents the capacity is decreased.

The total electrical energy supplied by a battery depends on the voltage of the battery and its capacity. Using the unit of charge introduced above, the energy stored in a battery in watt-hours (Wh) is calculated by multiplying the voltage by the capacity. Given that both voltage and charge vary depending on the condition and operation of the battery, the electrical energy delivered by a battery is also subject to variation.

One of the critical limitations of battery cells for use in electric vehicles has been that in order to hold an amount of energy comparable to a tank of gasoline, the mass and volume of the batteries would be prohibitively large. Therefore, two additional measures of energy storage are used: specific energy (energy per unit mass), and energy density (energy per unit volume).

Figure 2.1 shows the specific energy and the energy density of several available battery technologies. The lead acid cell has an energy density of less than 100 Wh l^{-1} and a specific energy of less than 50 Wh kg^{-1}. The established battery technologies of

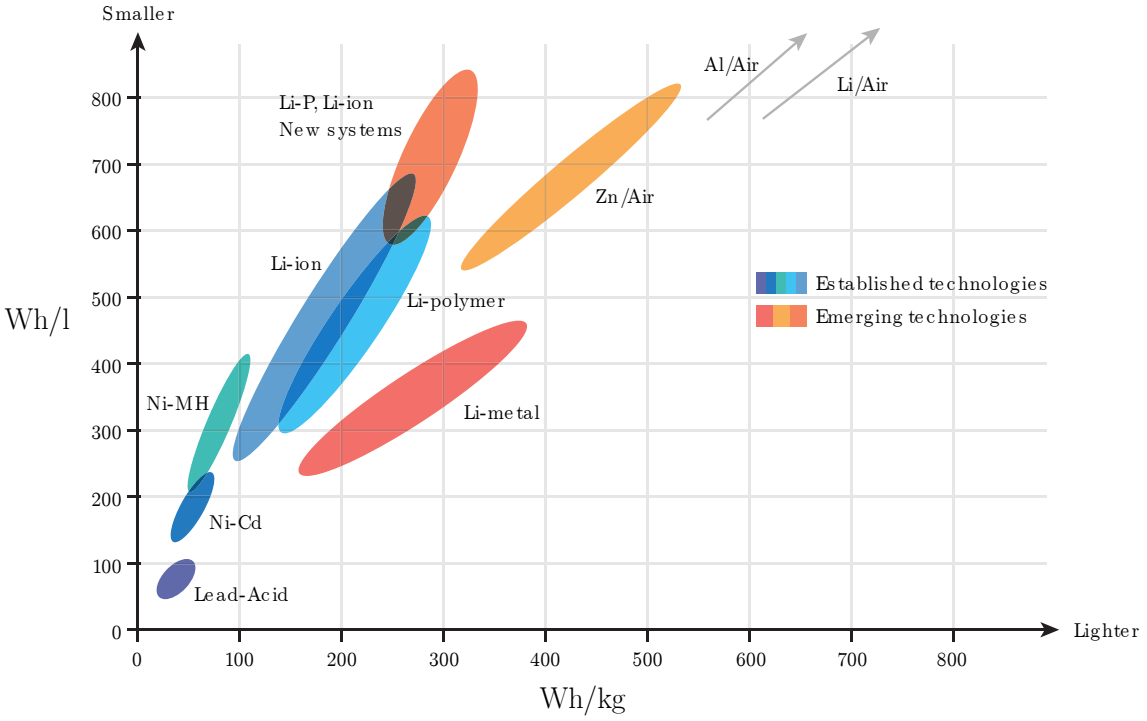

Figure 2.1 – Comparison of energy density for various battery cell chemistries. Note that one Watt-hour is equivalent to 3600 Joules. Modified from [22].

lithium-ion and lithium-polymer deliver specific energy of greater than 200 Wh kg^{-1} and energy density of up to 600 Wh l^{-1}. Continuing research into alternative cell chemistries promises to increase this further. By comparison, the specific energy of gasoline is approximately 13000 Wh kg^{-1} [20]. Even allowing for the low efficiency of the IC engine, typically between 20 and 30% [21], it is clear that for a BEV to compete with a conventional vehicle, the battery will be of considerable volume and mass.

The thermal characteristics of battery cells are of particular interest in this study, and these can be divided into cause and effect; i.e. factors that influence the generation of heat, and parameters that are dependent on the temperature of the cell. Heat generated within a battery cell is a product of electrical resistance, and so many of

the same factors apply: drawing a higher current and operating at low temperatures generates more heat. In addition, cells have a temperature range over which they perform optimally: operation at temperatures above this cause corrosive reactions to proceed at a higher rate, reducing the life of the battery. Operation at low temperature reduces the rate of necessary electrochemical reactions, and raises the internal resistance; these effects decrease the instantaneous power developed by the battery. Therefore, thermal management of the battery pack is essential to extract the best performance from the battery.

An additional, critical, consideration for some battery chemistries, including lithium-ion, is that at excessively high temperatures there is risk of thermal runaway. In this situation, there is a positive feedback loop whereby an increase in battery output causes an increase in temperature; and this in turn increases the battery output. If allowed to continue, this may lead to catastrophic failure of the battery and combustion. Studies [23, 24] have shown that for a lithium-ion cell, this occurs at temperatures of approximately 170–180°C.

Previous Work

Previous studies have measured the heat generated by lead-acid cells [25, 26, 27], nickel-zinc cells [27], and lithium-ion cells [27, 28] used for powering electric vehicles. In all studies, there was significant variation depending on temperature and discharge rate; most heat was generated at high discharge rates, and at low temperatures. Typical values for lithium-ion cells during representative driving profiles were measured at an average of between 2 and 20 W per cell, with peaks up to 50 W [28].

The spatial heat generation distribution in an EV battery cell has been assessed

both empirically and analytically. Pesaran *et al.* [25, 27] utilized thermal imaging to measure the temperature of various battery cells under different conditions. The temperature distribution was highly dependent on the battery chemistry, and geometry: cylindrical lead-acid cells showed a strong temperature concentration near the terminals, but cuboid nickel-zinc cells had the highest temperature in the body away from the terminals. Kim *et al.* used numerical modelling to simulate the electro-chemical reactions in a rectangular laminate lithium-ion cell, and predicted the temperature rise, validated with thermal imaging. They found a significant temperature concentration around the terminals, resulting in a marked temperature gradient across the cell from 60°C near the terminals to 40°C at the other end.

The effect on battery efficiency can be determined from the heat generation rates: by the principle of energy conservation, any energy released as heat reduces the energy available to be converted to electricity. Tests on lead-acid and nickel-zinc cells [27] measured the electrical efficiency over discharge cycles at different conditions. Slow discharge rates at high temperatures gave the most efficient operation; at high discharge rates and low temperature the cells supplied significantly less energy.

Although at high temperatures more electrical energy is available, there is a significant degradation effect. Shim *et al.* [29] cycled lithium-ion cells at 25°C and 60°C: at 25°C the cell retained 95% of its capacity after 140 cycles, but at 60°C only 35% capacity remained at the end of the test.

2.4.2 Thermal Management

In a conventional IC engine powered vehicle, systems are required that both heat the engine prior to ignition, and/or cool it during operation, depending on the environmental conditions. In cold climates, the cylinder block of an engine can be installed with an electric block heater that when supplied with mains electricity will keep the engine warm enough to start easily and operate at acceptable efficiency. In all vehicles, a cooling system is needed to limit the temperature of the IC engine during operation, to avoid damage to the engine material and lubricants. This is usually accomplished by a combination of liquid and air cooling.

In electric vehicles, there are similar constraints on the temperature, although it is the battery and not the engine that needs attention. The requirements of a battery thermal management system are to keep the battery temperature within the desired range, whilst keeping the temperature across the entire pack and within each individual cell as uniform as possible. This should be achieved at the lowest possible power consumption, as all power for the thermal management system directly reduces the power available to the other vehicle systems.

For operation in mild climates, and for parallel HEVs, it might be possible to use passive air cooled batteries, i.e. directing ambient or cabin air over the battery pack whilst the vehicle is in motion. However, for battery cells of higher energy density and greater heat generation, as found in BEVs, it is probable that active cooling systems with a liquid coolant will be required [26]. Another factor that affects the configuration of the thermal management system is the geometry and location of the battery pack: different solutions will be required for cylindrical and laminar cells; if the battery is in an inaccessible location, then passive air cooling will be more difficult.

It may also be desirable to insulate the battery pack to reduce its susceptibility to the ambient temperature.

Previous Work

The design of a battery thermal management system must consider numerous factors, and an attempt to define a design approach was made by Pesaran *et al.* [26]. The seven stages progress from definition of performance objectives and thermal parameters through system modelling to testing and optimization. This research group, at the National Renewable Energy Laboratory (NREL), USA, has subsequently published work discussing the characterization of battery thermal behaviour [27]; the best methods of pre-heating batteries in cold environments [30]; an assessment of the Prius battery [18]; and the thermal behaviour of a cylindrical jacket cooling configuration [31].

Chen *et al.* [24] created thermal models of a battery pack and used them to predict the effect on the battery performance of modifying parameters such as thermal conductivity, pack insulation, and warm-up/driving time. On the issue of battery pack cooling, the importance of temperature uniformity, and providing every cell with identical thermal boundary conditions was emphasized.

2.5 Cooling Plates

The focus of this study is on battery packs utilizing rectangular laminate cells, for which the desired cooling system is to use cooling plates inserted between battery cells. This is shown conceptually in Figure 2.2. The cooling plates are thin metal

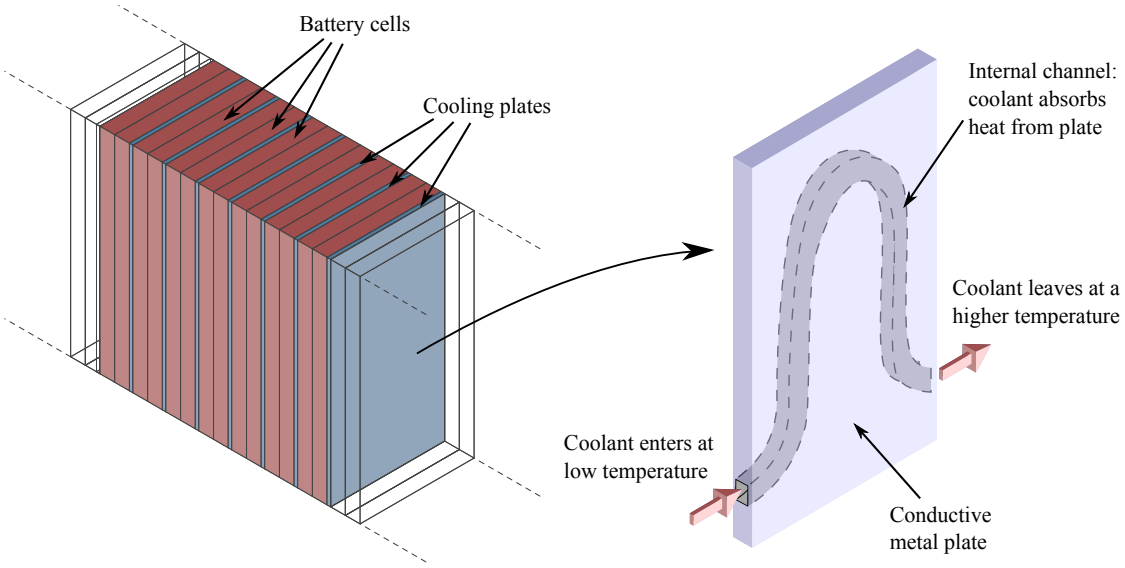

Figure 2.2 – Representation of a battery pack and cooling plate arrangement.

fabrications, with an internal channel through which a coolant flows. Heat is conducted from the battery cells into the plate, and then into the coolant which is pumped away to a heat exchanger. Similar arrangements have been used to cool microelectronics and fuel cells. Design requirements for a cooling plate are to have a high rate of heat transfer, whilst operating with a low coolant pressure drop and flow rate. Additionally, cooling plates that promote a uniform temperature within the battery cells are desirable, and as with all automotive applications, low mass and cost are also objectives.

As well as considering the performance of an individual cooling plate, there are design considerations that influence the entire battery pack. For example, the number of cooling plates is a design parameter, as well as the rate of coolant delivered to each one. (Is it preferable to deliver a high rate of coolant to few plates, or a low rate to many?) The piping to deliver coolant to the plates can be designed in a serial or parallel configuration, or some hybrid of the two.

Previous Work

Literature on the subject of cooling plates for the specific application of EV batteries is sparse, but in the fields of fuel cell and electronics cooling, previous studies provide some useful background.

The size and shape of fuel cell cooling plates is similar to that required for battery cooling, and their performance criteria are the same. Chen *et al.* [32], Choi *et al.* [33], Yu *et al.* [34] have all used computational fluid dynamics to assess the performance of a number of different plate geometries. They found that even with identical heat generation and coolant flow rate, the geometry of the channel had a significant influence on the temperature mean, maximum, and uniformity. Serpentine channel configurations (in which a single channel covers the whole plate) led to more uniform temperature distributions than parallel channels; this was a result of poor coolant distribution between the channels. These studies, however, were all limited to a single inlet and outlet point, and chose the optimum plate from a predetermined set of designs with uniform channel width; there was no further optimization.

One of the most important non-thermal characteristics of a cooling plate is the drop in pressure of the coolant as it flows through the channels. Maharudrayya *et al.* [35] considered the pressure lost at 90 degree bends of a PEMFC coolant channel and found that they can contribute significantly to the total pressure loss, especially at high Reynolds Numbers. They recommended further optimization to calculate the ideal channel configuration.

For the active cooling of high-powered electronics, liquid-cooled heat sinks have been in use for many years. They operate on a smaller scale than an EV battery, but have similar thermal requirements and have been the subject of numerous analysis

and optimization studies. A micro-channel heat sink typically has many parallel coolant channels running through a heat sink attached to the heat source. Given the small size of many electronics components, the heat sink and channels extend perpendicularly away from the component in order to increase the available surface area for heat transfer (this is one of the main differences from a typical battery cooling plate, which has a small dimension in the perpendicular direction).

Tuckerman and Pease [36] first analyzed the performance of a micro-channel heat sink using idealized heat transfer to characterize thermal resistances in the system. Following this work, many efforts have been made to optimize the design. Kim [37] reviewed three analysis and optimization methods: the fin approach, the porous medium approach, and numerical optimization. The first two of these use approximations to address the problem, but for a simple system with uniform parallel channels, the fin method can be applied with good results. Numerical optimization is considered the most accurate, but also the most time consuming.

Knight *et al.* [38] [39] used the fin method to optimize channel number, thickness and ratio for both laminar and turbulent flow. Bau [40] also used a fin method optimization, first for a uniform channel, and then allowing the channel width to vary along its length thereby improving temperature uniformity. Wei and Joshi [41] introduced multiple layers of channels to increase the area for heat transfer and optimized both the channel geometry and number of layers. Ryu *et al.* [42] used the numerical optimization approach with a three-dimensional finite volume model of the cooling channel: the channel number and geometry were selected as design variables and thermal resistance as the objective function. With a constraint on pumping power, they found the channel width to be the most significant design variable.

Kurnia *et al.* [43], also working in the field of microelectronics, applied the same methodology as the fuel cell studies: assessing eight fixed channel geometries at a range of flow rates for the objectives of temperature uniformity, average temperature and pressure drop. As in the studies mentioned above, the channels with coils or switchbacks gave the best thermal results, although sometimes at the expense of higher pressure. This study also introduced a mixed-objective 'figure of merit': cooling performance per unit pumping power. However, the pump power tends to be the overwhelming factor, reducing the utility of the figure of merit when cooling performance is the main concern.

Most studies of cooling channels assume square-edged channels, as these are simplest to model; however, precise manufacturing is sometimes infeasible. Fisher and Torrance [44] [45] used the boundary element method to assess heat transfer into channels with rounded corners of varying degree: from rectangular through elliptical to diamond-shaped. They found that moving away from a rectangular cross-section decreases the rate of heat transfer (presumably due to a reduction in surface area), although the optimal spacing between channels becomes smaller, allowing more channels in the same space. Overall, a cooling plate with elliptical channels has heat transfer decreased by approximately 5%.

The optimization studies mentioned above are all based on size or shape optimization algorithms, and as such, can only investigate a limited design space. Several studies of fluid flow and heat transfer have used topology optimization [46] to remove these restrictions. Borrvall and Petersson [47] and Oleson *et al.* [48] have used topology optimization of a fluid-solid domain to determine flow conduits with minimum pressure drop; this results in a typical hydrodynamic geometry. Dede [49] additionally

used heat transfer to define a composite objective of pressure loss and temperature. This method has the power and flexibility to generate innovative designs, however the complexity of the technique limits the accuracy of the simulation, and the resultant designs need significant modification before they would be fit for manufacture.

2.6 Design Optimization

The purpose of design optimization is to determine the design, $\boldsymbol{x}$, that performs best when measured against an objective function, $J(\boldsymbol{x})$, while satisfying a set of constraints, $C(\boldsymbol{x})$. In this formulation, the design must be parameterized, i.e. describable by a set of dimensions, features, properties, etc. that together constitute $\boldsymbol{x}$. The optimization is carried out by modifying the design parameters in such a way that the objective function is improved (increased or decreased, as desired), while staying within the feasible design area described by the constraint function and limits on the design variables. The optimum design is reached when no further improvement is possible within the area bounded by the constraints. Figure 2.3 illustrates some of these points with an example of one design variable. Mathematically, a typical optimization problem can be described as:

$$\left.\begin{array}{rl} \text{minimize :} & J(\boldsymbol{x}) \\ \text{subject to :} & C(\boldsymbol{x}) \leq 0 \\ & C_{eq}(\boldsymbol{x}) = 0 \\ & \boldsymbol{x}_{min} \leq \boldsymbol{x} \leq \boldsymbol{x}_{max} \end{array}\right\} \tag{2.1}$$

Where $C(\boldsymbol{x})$ and $C_{eq}(\boldsymbol{x})$ are the two types of constraint: inequality, and equality respectively. $\boldsymbol{x}_{min}$ and $\boldsymbol{x}_{max}$ define the absolute limits on the design variable parameters.

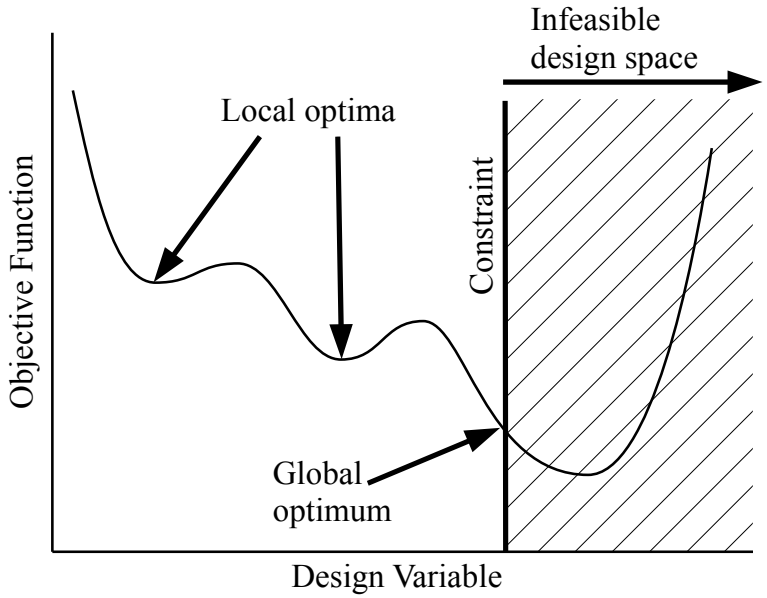

Figure 2.3 – Example of a one dimensional design optimization problem. The function has a global minimum within the infeasible region described by the constraints; therefore, the optimum design is found on the intersection with the constraint.

The critical step in the optimization process is that of sensitivity analysis and design variable update, i.e. to improve the objective function, in which direction and by how far must the design variable be changed? In the one-dimensional example presented here, sensitivity analysis is simply the calculation of the gradient of the objective function with respect to the design variable at the location of the current design. If the objective function is differentiable, then it may be possible to evaluate this analytically, but more often a numerical approach is required. The simplest method is the forward finite difference method, whereby the objective function is evaluated at the current design and at a slightly perturbed design, then a simple calculation will produce the gradient. In multiple dimensions, the same principal can be used, requiring $N + 1$ objective function evaluations, where N is the number of design variables.

Having established the direction of decreasing objective function, the magnitude of the design variable change can be found by a one-dimensional line search. This considers the objective function only in the direction of steepest descent (thus one-dimensional), and attempts to estimate the point at which the function will be at a minimum. The subsequent design iteration then takes on the value of the design variable at this point.

Problems can arise in the optimization process from a variety of sources. In the example shown in Figure 2.3, there are two local minima that the optimization algorithm could end up in. Unable to decrease the objective function further, the conclusion will be that the optimum design has been determined. This reveals an important limitation of design optimization: without knowledge of the entire design-space, there is no way to prove that the optimization result is the global optimum solution. By starting the optimization process from multiple initial designs, the chances of finding it can be increased, but it remains an uncertainty.

Sensitivity analysis and line-search can be misdirected if the objective function is ill-conditioned. When interrogating a totally smooth function, the accuracy of the finite differencing will be increased as the offset between sampling points decreases. However, if the function contains any noise, then the offset will need to be large enough that the noise is filtered out. If not, the sensitivity calculated will be a random function of the noise, and the objective function information will be lost. In multi-variate functions, the steepest descent method is often an inefficient means of reaching the optimum: the local direction of steepest descent rarely points towards the minimum of the function. To improve on the steepest descent method, the conjugate gradient method can be used whereby gradient information from multiple design

iterations is used to go more directly to the minimum. Although this can improve the algorithm efficiency, the optimization of complex functions, or systems where the optimum is determined by a constraint, can still be problematic.

Multi-Objective Optimization

It is frequently the case that a design problem can be formulated with more than one objective function. Sometimes optimization with respect to different objective functions will result in essentially the same solution; for example, component mass and cost are often closely related. In other cases, an improvement in one objective function can only be realized at the expense of another objective function; for example, component mass and structural deflection. For these problems, typical single objective optimization will result in a design that completely fulfills one objective but ignores the other. To better understand the design space, multi-objective optimization can be used to determine intermediate designs and the trade-off between competing objective functions.

The principal difference in the formulation of a multi-objective optimization is that the objective function is composed of two or more of the single objective functions, typically a weighted sum. By controlling the weighting factors, the optimization can be made to favour one or other objective; executing several optimizations with different weighting factors produces several optimum designs that occupy positions between the single objective designs. This set of designs linking the single objective solutions is called a Pareto front and is illustrated in Figure 2.4 with an example of an optimization of two functions, both to be minimized. In this case, the perfect solution (the *utopia* point) would occupy a position towards the origin of the axes, but the

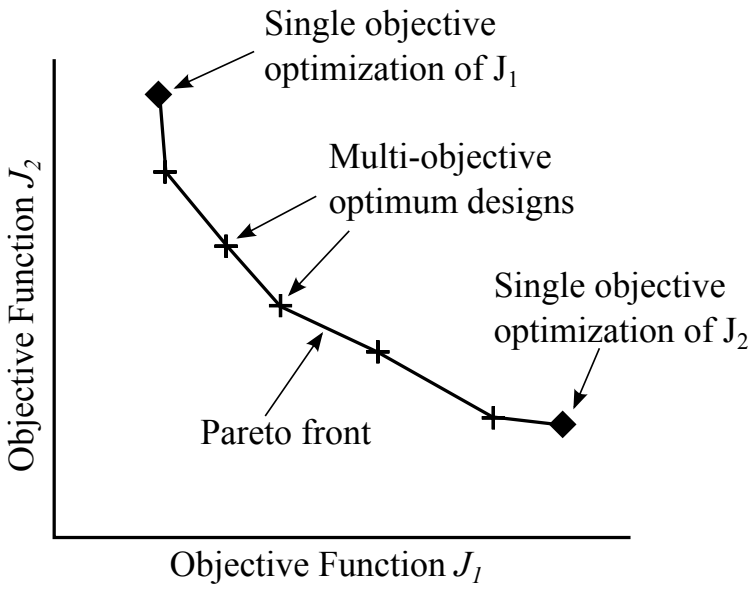

Figure 2.4 – Example of the optimization of two competing objective functions, each of which is to be minimized. The single objective optimum designs form the anchor points of the Pareto front, and the multi-objective optimum designs are found by varying the weighting factors of the multi-objective formulation.

nature of the system is such that designs do not exist here. The closest designs are found along the Pareto front, which must satisfy the condition that no point on the curve outperforms another in both objective functions. Therefore, the designs along the Pareto front describe the sacrifice required in one objective function to achieve a gain in the other objective function.

2.7 Computational Fluid Dynamics

Computational fluid dynamics is a numerical method that can be used to determine the behaviour of systems involving fluid flow, heat transfer, diffusion, combustion, chemical reactions, and other associated phenomena. In its most widely used form, it is a finite volume technique in which the domain of interest is divided into discrete

volumes, called elements or cells, and governing equations are applied to the contents of each volume. In the simplest case, this means applying the principal of conservation to the volume of the cell, i.e. the rate of change of a variable within a cell is equal to the net rate of transport into the cell plus the net rate of creation within the cell.

The basic equations governing conservation of mass and momentum are required for an incompressible fluid system, and additional equations of energy conservation, equations of state, and the Navier-Stokes equations allow the problem to be generalized to the compressible three-dimensional flow of a Newtonian fluid including heat transfer. The conditions in this study required the use of CFD for incompressible flow with heat transfer, and so not all equations were needed.

Up to a certain Reynolds number[1] fluids in motion follow steady, predictable paths and there is smooth transition between fluid layers flowing at different speeds. Fluid systems displaying this behaviour are termed laminar, and in simple cases can be modeled without approximation. However, beyond a critical value of Reynolds number, the flow becomes unstable and turbulence is encountered.

Turbulence is a phenomenon that is observed in all fluid flows above a certain threshold value of Reynolds number. In turbulent flow, random fluctuations in velocity and other flow properties are observed even under constant boundary conditions. Rotating flow structures, or eddies, are present at a range of dimensions that promote much more effective mixing than in laminar flow. There is no economical means to determine the precise state of the fluid, and so a common approach to turbulence is to define a steady average property, and some time dependent fluctuating component described statistically. This introduces many more flow variables, and it is the role

[1] the Reynolds number is a dimensionless property of a fluid flow defined by UL/v where U and L are characteristic velocity and dimension of the flow; v is the kinematic viscosity

of the turbulence model to allow these to be solved.

The application of CFD to solve a fluid flow problem typically involves three steps: pre-processing, solving, and post-processing. The pre-processing stage is the one that requires the greatest input from the user; in this stage a model is generated (often via a graphical user interface), a mesh defined to subdivide the domain into volumes, and boundary conditions selected. It is also down to the user to select the physical models (heat transfer, turbulence, etc.) and settings that are appropriate for the analysis being attempted. The specification of the mesh can have a significant influence on the speed and accuracy of the analysis: a mesh fine enough to resolve all the detail in the flow variables will give a more accurate result than a coarse mesh that can only resolve general trends. However, finer meshes require greater computational resources: memory requirements, time for a single CFD iteration, and the number of iterations to convergence will all increase. Therefore, much of the skill of pre-processing is in constructing a mesh that will be accurate where required, without unnecessary computing cost.

After fully defining the model to be analyzed, the solving stage can begin. The governing equations are applied to each cell of the solution domain and then discretized into a set of algebraic equations. Cells at the edge of the domain with boundary conditions generate additional equations. Given that fluid systems display complex, non-linear behaviour, an iterative numerical method is required to solve the equation system, and the computational resources required are often significant.

The post-processing stage is the reverse of pre-processing: the model and results are presented graphically and interrogated by the user. There is a range of visualization tools available to interpret the behaviour of the system, and any fluid property,

vector, or flux can be plotted and calculated.

Chapter 3

Cooling Plate Optimization: Method

3.1 Problem Statement

3.1.1 Cooling Plate Geometry

A need for effective and efficient thermal regulation of EV batteries has been identified, and although there is reason to believe that optimization methods can be used successfully in liquid cooled devices, none have yet been applied to the particular problem of an EV battery cooling plate. In design problems such as this, the feasible design space is effectively infinite: while the outer dimensions of the plate are constrained by the battery dimensions, the inner geometry describing the coolant channel is continuously variable. Although techniques such as topology optimization have the potential to explore a large region of the design space (albeit at low fidelity), their results are rarely immediately practical. This study intends to investigate whether

numerical optimization can produce practical benefits in the design of cooling plates, and so the design space is restricted to allow a more thorough assessment of designs.

A design optimization process requires that the design be parameterized, i.e. described by a set of parameters: these are most often continuous and numeric, but other data types are possible. The size and complexity of the optimization problem is determined by the number of design variables: those parameters that are controlled to modify the design. With too few design variables, the design space is too limited and optimum designs may be predictable; with too many design variables, the optimization process becomes very time consuming and may have trouble converging.

The selection of the cooling plate design to be optimized should respect these optimization and design space considerations, and ideally offer the opportunity to produce a cooling plate design that outperforms those currently available. For a cooling plate geometry that performs well, and is easily parameterized, the conclusions of the study by Yu *et al.* [34] were considered, which determined that of the six coolant channels shown in Figure 3.1, model D resulted in the lowest average temperature and second lowest temperature variation. However, to parameterize the geometry of this plate, every straight channel section would have a width and position variable, giving a total of 58 design variables. This optimization problem could be tackled given sufficient computing resources, but reducing the complexity of the cooling plate leads to a two-fold benefit: the size of the CFD model required to accurately analyze its behaviour is reduced, and the number of design variables is reduced. These two effects combined lead to substantial computational savings.

Nam *et al.* [50] described an algorithm for constructing cooling channels, which was applied in reverse to this channel to reduce its complexity by a factor of three,

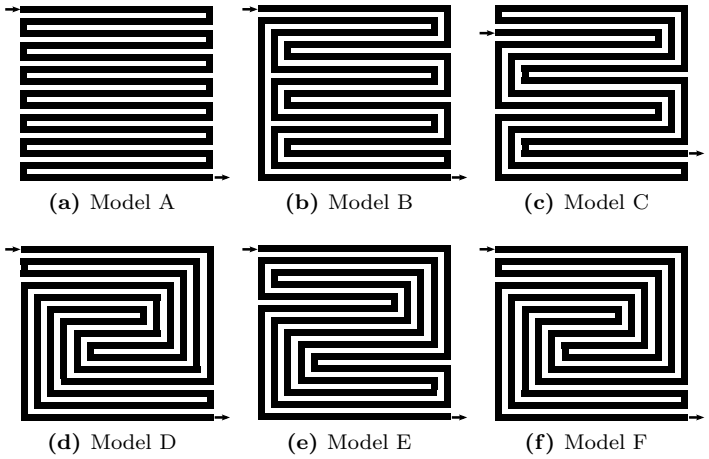

Figure 3.1 – Cooling channel geometry studied by Yu *et al.* [34].

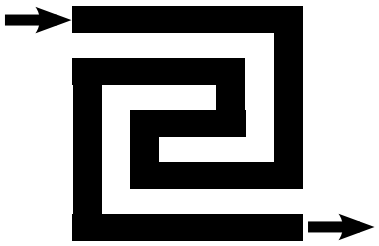

Figure 3.2 – Simplified cooling channel geometry.

leaving the 'double-S' design shown in Figure 3.2. This was described by 18 design variables: two for each straight section of the channel. As presented in Figure 3.2, the cooling channel covers a square area; however, technical information from GM Canada indicated that their battery cells have a size of 160 × 200 mm. To conform to this ratio, the position of the channels was slightly elongated, while maintaining a constant width. This is the design used in the present study as the basic cooling channel geometry.

3.1.2 Boundary Conditions

During operation, the thermal condition of an EV battery are determined by many factors including the instantaneous power demands of the vehicle, the driving history, the level of charge in the battery, and the external environment. The combination of these factors leads to complex transient behaviour that a thermal management system must adapt to. Ideally, any design optimization process should take into account the transient nature of the system, but to simplify the problem, a steady-state analysis was tackled in this study. Technical data received from GM Canada gave an indication of the operating range of a typical battery cell and cooling system: each cell generates heat at a rate of approximately 17 W, and the pumping system circulates coolant at a rate of between 2.5 and 20 l min^{-1} for 100 battery cells.

To define the baseline boundary conditions, a ratio of two battery cells to every cooling plate was assumed. Each cooling plate then absorbs a total of 34 W of heat, equivalent to a flux of 531 W m^{-2} from each side. The coolant flow rate to each cooling plate is in the range of 0.00083 kg s^{-1} to 0.0067 kg s^{-1}. Boundary conditions selected for the baseline analysis were a uniform heat flux of 500 W m^{-2} and a coolant flow rate of 0.001 kg s^{-1} at a uniform temperature of 300 K. However, variation in these conditions may influence the optimum cooling plate design, and so assessment at different conditions is also required.

The battery pack of the GM Chevrolet Volt is insulated to limit the influence of ambient temperature [51]. Although there will inevitably be some heat transfer between the battery and its surroundings, it is convenient for the sake of analysis to assume that it is completely insulated. Therefore, the perimeter faces of the cooling plate were considered to be insulated.

The spatial distribution of heat generation over the battery cell was highlighted in the previous section as a subject requiring investigation. Therefore, an additional objective of this study is to determine how non-uniform heat generation affects the performance and optimum design of the cooling plate.

3.1.3 Objective Function

There are two principal objective functions related to the battery temperature: cooling efficiency, and uniformity of temperature. In this cooling plate analysis, the cooling efficiency was assessed by considering the average temperature of the face through which the heat flux enters the cooling plate—this is the face that is in contact with the battery.

Considering the cooling plate as a steady state heat transfer system (illustrated in Figure 3.3), there is a constant flow of heat from the solid cooling plate to the coolant, driven by a temperature gradient between the plate and coolant. The cooling efficiency can be related to the magnitude of that temperature gradient: a more effective cooling plate will be able to transfer the heat at a lower temperature difference.

Note that in all comparisons of average temperature, it is the increase above the coolant inlet temperature that is considered; i.e. taking the inlet temperature (300 K) to be zero. For example, comparing 308 K to 310 K would be a decrease of 20%.

The temperature uniformity can be assessed by several measures: variation, standard deviation, or range of temperature. In this study, standard deviation was selected as it incorporates temperature data from the entire face instead of just the two points corresponding to the maximum and minimum temperatures.

The third performance criterion for an EV cooling plate is the power consumed

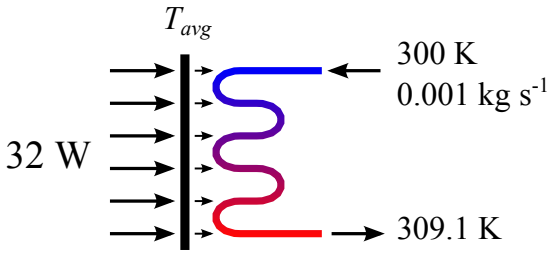

Figure 3.3 – Schematic diagram showing the flow of heat in the cooling plate analysis. A constant heat flux transfers 32 W into the plate, which must then be transferred to the coolant. The coolant enters the plate at a constant flow rate and temperature; by absorbing the heat from the plate, its temperature is raised by 9.1 K. The average temperature of the plate, T_{avg} determines the temperature difference between plate and coolant, and therefore the heat transfer efficacy.

by the circulation of the coolant. The theoretical pump power can be calculated from the product of volumetric flow rate and drop in total pressure (although in reality there are further reductions to account for the pump efficiency) [52]. As the coolant flow rate was specified as a fixed boundary condition throughout the optimization process, the pressure drop from inlet to outlet was used as an objective function in place of the power consumption.

Note that a calculation of power consumption requires the total pressure. This is the sum of static pressure (the pressure that would be experienced by any object moving at the same velocity as the fluid), and dynamic pressure (related to the kinetic energy of a moving fluid). At the conditions investigated in this study, the dynamic pressure is almost negligible—generally less than 1% of the total pressure—therefore total and static pressures are approximately equivalent. Furthermore, as the difference between static and total pressure is determined solely by the fluid velocity, total pressure measurement is only necessary when there is a difference between inlet

and outlet velocity (i.e. when the channel dimensions are different). Therefore, in the current study, using total pressure rather than static pressure will generally have a small effect; however, to allow for future studies that may use different conditions, total pressure was used.

3.2 Cooling Plate Model

For integration into the optimization routine, the cooling plate model must be parameterized and controlled by the optimization algorithm. To meet these requirements, the CFD pre-processor GAMBIT was selected; this has the functionality to be launched in batch mode directly from a command prompt and generate a model based on instructions in a journal file. The journal file used in this study was composed of two parts: a list of parameters written by the optimization algorithm, and a list of commands that construct the model based on these parameters.

3.2.1 Geometry

A GAMBIT journal file was written to create a three dimensional cooling plate constructed from orthogonal volumes. The plate width, height and depth are all defined parametrically, as are the number of volumes in the in-plane directions. In the through-thickness direction there are two layers, one of which is all solid, and in the other layer the volumes are assigned to either the fluid or solid domain to define the geometry of the coolant flow channel. Symmetry boundary conditions act on the top layer to complete the cooling plate. Figure 3.7 illustrates the process of constructing the geometry and Figure 3.4 shows some examples of cooling channel geometries that

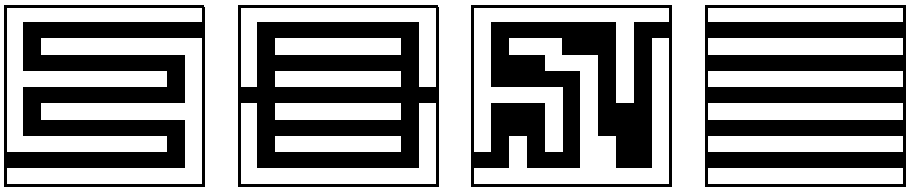

Figure 3.4 – Possible cooling plate geometry constructed with GAMBIT modelling algorithm. The coolant channel is shown in black.

can be constructed with this algorithm.

The cooling channel geometry was selected as described in the Problem Statement, and the GAMBIT journal file was set up accordingly. The outer dimensions of the plate were determined by the size of the battery cells: 160 $\times$ 200 mm. The total cooling plate thickness was 1 mm, and with a symmetry condition a modelled thickness of 0.5 mm. The cooling plate reference design has channels of 20 mm width centred and evenly spaced within the plate. In the through-thickness direction, the channels have a depth of 0.75 mm centred in the plate.

The primary design variables in this optimization are the in-plane position and width of the coolant flow channels. These are controlled by parameterizing the x and y dimensions of the construct volumes while keeping all construction lines orthogonal. The complete set of design variables is illustrated in Figure 3.5 and some possible variations on the selected cooling channel geometry are shown in Figure 3.6.

The faces for the application of boundary conditions were also selected in GAMBIT. All the external faces were assigned a boundary condition: either symmetry, inlet, outlet, heat generation, or insulated wall. In addition, the interface between solid and fluid was selected to make it possible for heat to transfer between domains.

Figure 3.8 illustrates the geometry and boundary conditions of the reference cooling plate geometry.

3.2.2 Mesh

The meshing procedure was also written using parameterized commands for the simple execution of mesh convergence checks and other mesh fine-tuning. Each orthogonal volume was meshed with hexahedral elements with a density defined by several parameters. Solid and fluid domains were treated independently, likewise the through-thickness and in-plane directions. In the fluid domain, boundary layers were used along interfaces with the solid to increase the mesh density close to the wall. In general, the mesh was defined by the number of element intervals per construct volume, and for boundary layers by the number of elements, the ratio, and fraction of construct volume occupied by the boundary layer.

Mesh parameters were selected to provide a combination of accuracy, consistency, and computational efficiency, and are shown in Table 3.1. In total 243 560 elements were used in the mesh; with 231336 of these in the fluid domain. Further discussion of the mesh density is found in the Validation section.

The final command in the journal file is to output the model as a mesh file in a format compatible with FLUENT. Figure 3.7 describes the major steps in the process of mesh creation.

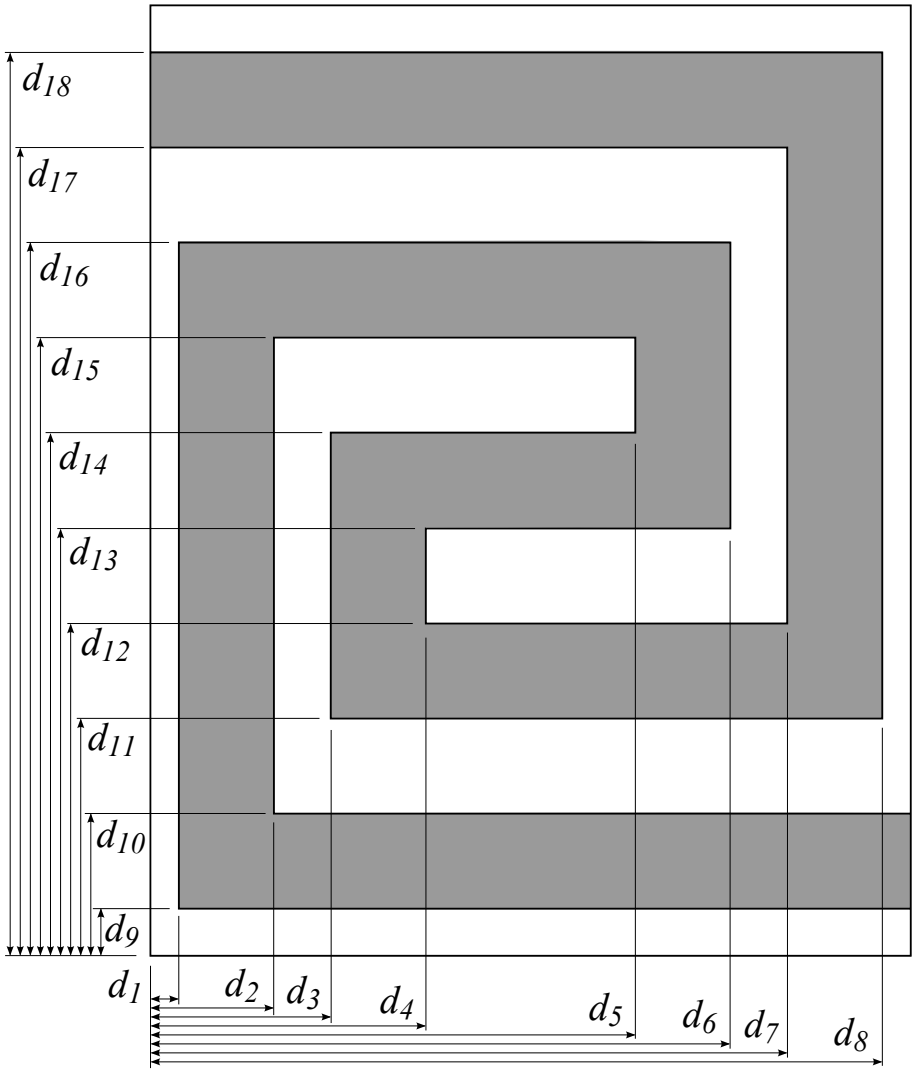

Figure 3.5 – Design variables assigned to the cooling plate geometry.

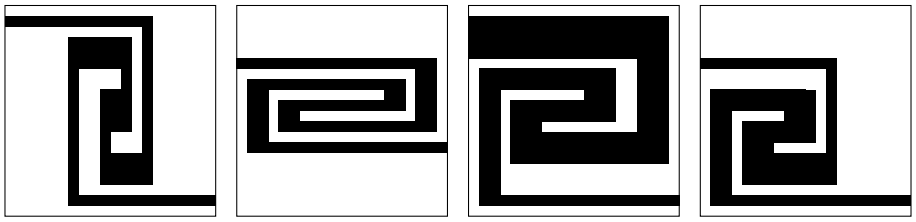

Figure 3.6 – Variations in design variables on cooling plate geometry.

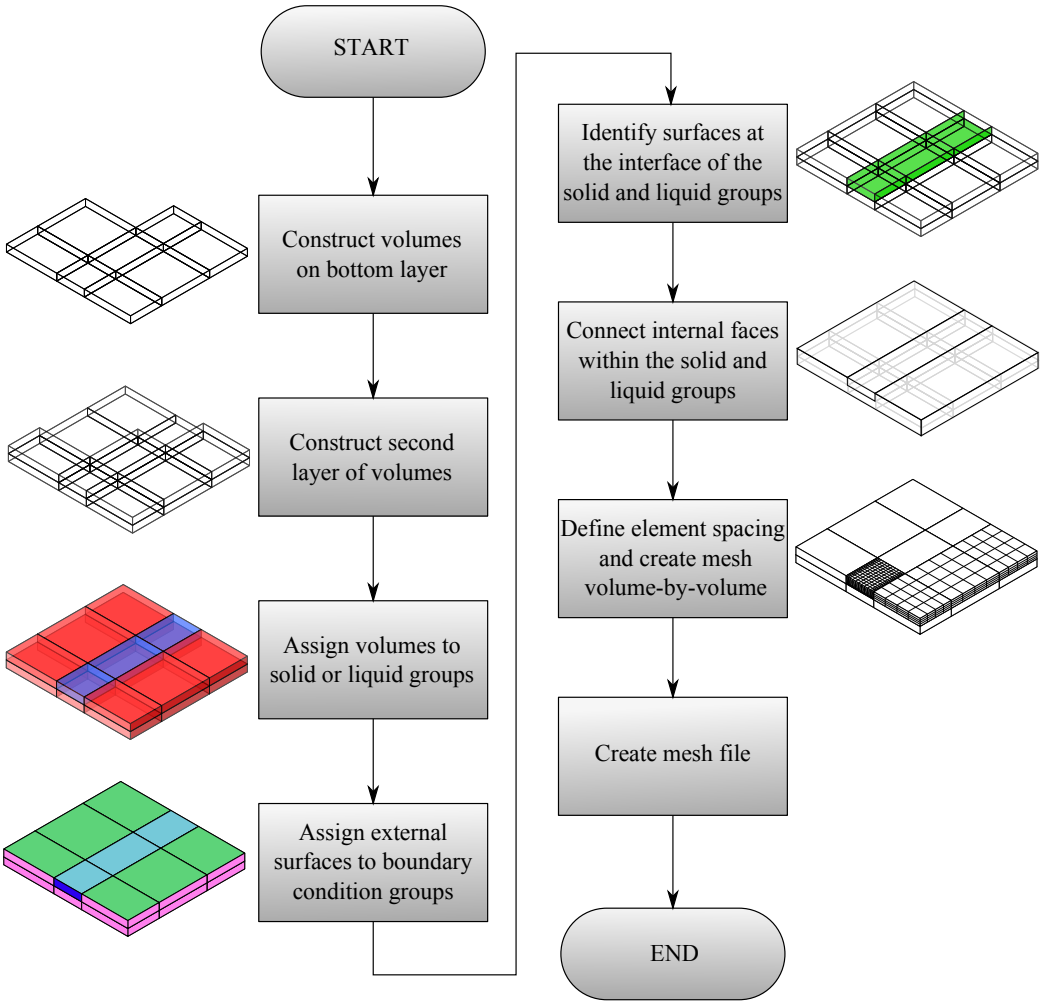

Figure 3.7 – Process of geometry and mesh generation.

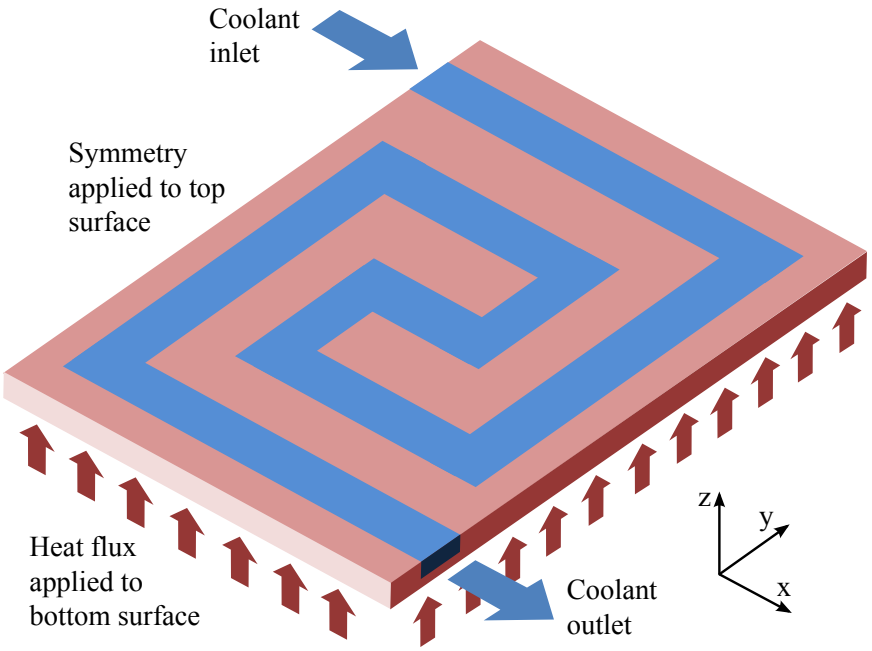

Figure 3.8 – Schematic of the cooling plate analysis model (not to scale).

Table 3.1 – Mesh generation parameters. Element numbers refer to the individual construct volumes.

Mesh Parameter	Value
Solid in-plane elements	4
Solid through-thickness elements	4
Fluid in-plane elements	18
Fluid through-thickness elements	14
Fluid in-plane BL elements	5
Fluid in-plane BL width fraction	0.2
Fluid in-plane BL expansion ratio	1.2
Fluid through-thickness BL elements	4
Fluid through-thickness BL width fraction	0.2
Fluid through-thickness BL expansion ratio	1.2

Table 3.2 – Material properties of 50-50 ethylene glycol – water mix.

Property	Model	Value
Density (kg m^{-3})	constant	1065
Specific heat (J $kg^{-1}K^{-1}$)	polynomial	$2574.7 + 3.0655T$
Thermal conductivity (W $m^{-1}K^{-1}$)	constant	0.42
Viscosity (kg $m^{-1}s^{-1}$)	power law	$0.0069 \times \left(\frac{T}{273}\right)^{-8.3}$

3.3 CFD Analysis

The software selected for the CFD analysis was the commercial package, FLUENT. Like the GAMBIT execution, a FLUENT journal file was written that contains the analysis parameters and commands. It is called in batch mode from MATLAB with the journal file specified as the input; following the analysis, MATLAB reads the output values and uses them in the optimization.

After launching FLUENT, the GAMBIT model is imported and scaled to metre units. The materials used in the analysis are aluminum in the solid domain, and a 50-50 mix of water and ethylene glycol in the fluid. The latter is a user defined material model with properties extracted from the literature of a commercial vendor [53]. Material properties are shown in Table 3.2.

The model boundaries were assigned in GAMBIT, and so the only steps to perform in FLUENT are to specify the flow rate and temperature of the inlet, and the magnitude of the heat flux. Table 3.3 describes all the boundary conditions and states the values used for baseline analyses.

FLUENT offers many settings that determine the accuracy and efficiency of the CFD solver: those selected for this simulation are presented in Table 3.4. In many

Table 3.3 – CFD boundary conditions used for baseline analyses. Note that the mass flow rate is halved to account for the symmetry condition.

Model Boundary	Type	Value
Inlet	Mass flow rate	5×10^{-4} kg s^{-1}
Inlet	Temperature	300 K
Wall	Heat flux	500 W m^{-2}
Outlet	Static pressure	0 Pa

cases, their selection was driven by a need for stable output parameters that allow reliable finite differencing. For example, Figure 3.9 compares second and third order temperature discretization schemes by considering the response of T_σ to small changes in a single design variable. Using the third order scheme, the output is much smoother, and finite differencing will have a greater accuracy. The second order scheme is too unstable for reliable finite differencing: to calculate a representative gradient, the finite differencing sampling points would need to be much further apart, leading to a loss of precision.

FLUENT is executed to solve for laminar fluid flow and energy with convergence criteria on the x- and y-velocity residuals set to 10^{-7} (this value is discussed in the Verification section). Following the convergence of the analysis, FLUENT is directed to write text files containing all the required output data. These are used to calculate the total pressure drop from inlet to outlet surfaces, and area-weighted average and standard deviation of the temperature on the heat flux boundary condition surface.

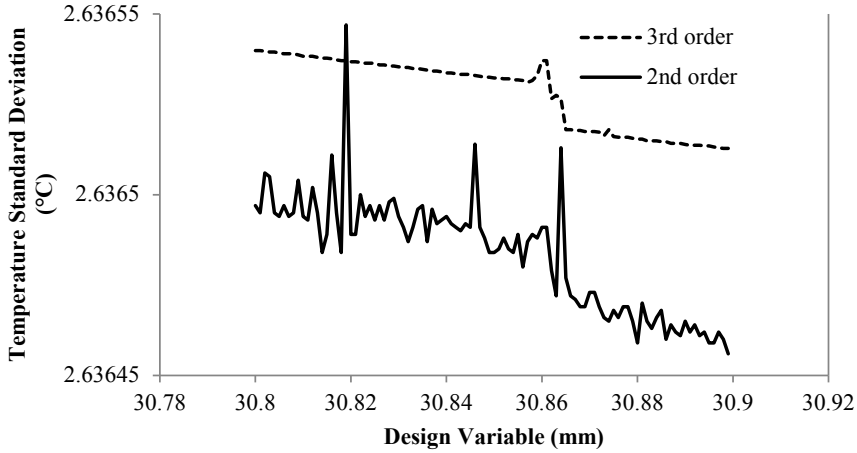

Figure 3.9 – Response of temperature deviation to small changes in a single design variable. Comparison of the 2nd and 3rd order temperature discretization schemes.

Table 3.4 – CFD solver settings. High order solvers were used to improve output consistency. Pressure and momentum under-relaxation factors were decreased slightly from the default to improve solution stability.

Setting	Value
Pressure-velocity scheme	SIMPLE
Gradient discretization	Green-Gauss node based
Pressure discretization	Second order
Momentum discretization	Third order MUSCL
Energy discretization	Third order MUSCL
Pressure under-relaxation	0.2
Density under-relaxation	1.0
Body force under-relaxation	1.0
Momentum under-relaxation	0.6
Temperature under-relaxation	1.0

3.4 Optimization Structure: Single Objective

The previous sections defined a process for generating a model controlled by design variables, and a CFD simulation that analyzes the model and writes output files. The third component is the optimization algorithm, which was written and run in the programming environment, MATLAB.

3.4.1 Formal Optimization Statement

Before describing the process of optimization, the problem to be solved should be clarified. There are three objective functions: the coolant pressure drop, P_{fluid}; the average temperature of the heat generation surface, T_{avg}; and the standard deviation in the temperature of the same surface, T_σ. In the single objective optimization, these were considered one at a time, and a separate optimization executed for each one. The design variables, $\boldsymbol{d}$, are the 18 dimensions needed to describe the shape of the cooling channel in the x-y plane. Constraints on $\boldsymbol{d}$ are used to maintain the topology of the cooling channel and prevent any excessively narrow sections (less than limit Δ). Combining these features, the optimization problem is as follows:

$$
\begin{array}{lll}
\text{minimize} & : & P_{fluid}(\boldsymbol{d}) \\
 & & T_{avg}(\boldsymbol{d}) \\
 & & T_{\sigma}(\boldsymbol{d}) \\
\text{subject to} & : & t_{wall}(\boldsymbol{d}) \geq \Delta \\
\text{where} & : & P_{fluid} = P_{inlet} - P_{outlet} \\
 & & T_{avg} = \dfrac{\int_{A_0} T dA}{\int_{A_0} dA} \\
 & & T_{\sigma} = \sqrt{\dfrac{\int_{A_0} (T - T_{avg})^2 dA}{\int_{A_0} dA}}
\end{array}
$$

where P_{inlet} and P_{outlet} are the area weighted average total pressure on the inlet and outlet faces respectively; A_0 is the bottom surface of the plate, over which T_{avg} and T_{σ} are defined; the constraint function, t_{wall}, is the minimum solid or fluid section thickness.

3.4.2 Optimization Algorithm

The optimization is controlled in the MATLAB programming environment, and makes use of the constrained nonlinear programming function, *fmincon*. From the MATLAB help file, 'fmincon attempts to find a constrained minimum of a scalar function of several variables starting at an initial estimate'. The use of this function automates the tasks of finite differencing, sensitivity calculation, and design variable update. In addition to the function definition and initial estimate, fmincon also accepts as inputs design variable bounds, inequalities or equalities, and a set of options that specify

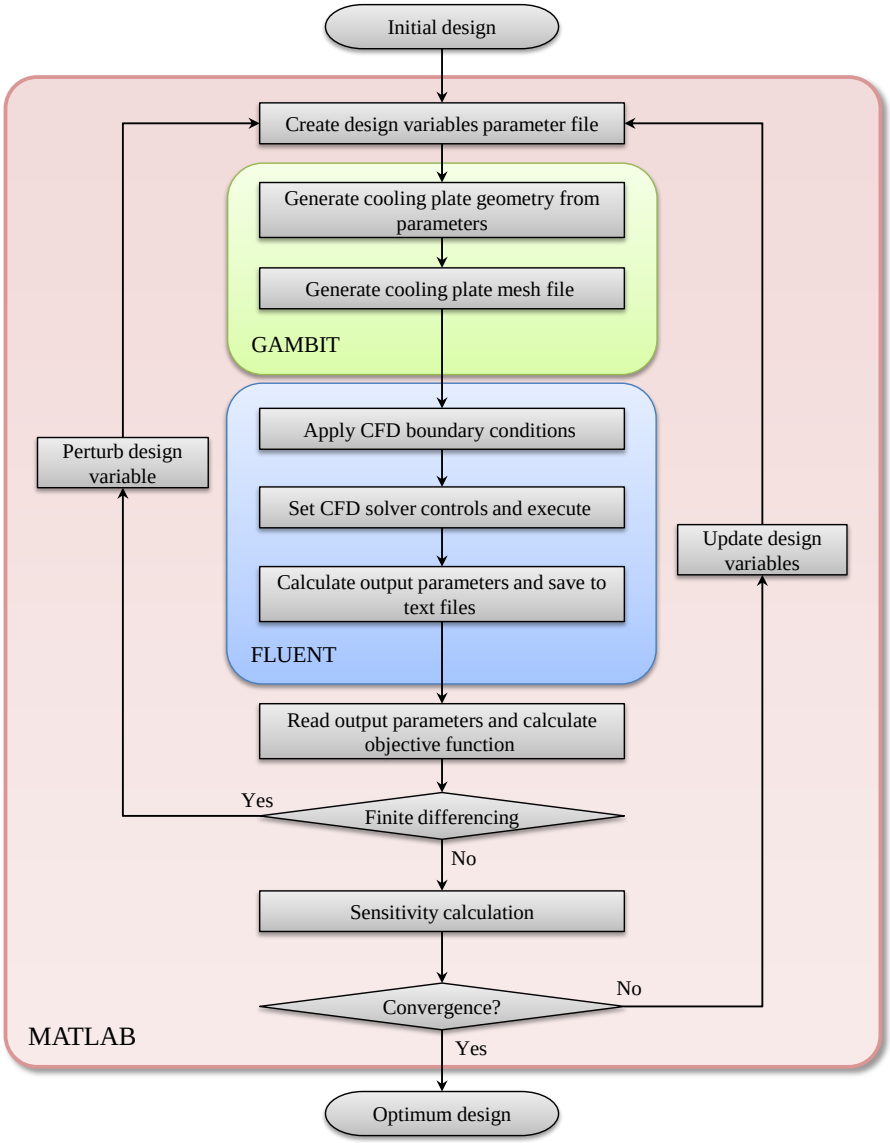

Figure 3.10 – Structure of the optimization process.

the optimization settings. The structure of this optimization problem is quite simple, and the majority of the MATLAB code is concerned with constructing appropriate input files for GAMBIT and FLUENT, and then retrieving the analysis results. This data flow is shown in Figure 3.10.

The MATLAB code is divided into several subroutines described as follows:

`runfmincon` This is the main function, called with inputs specifying which objective function to use, and values of coolant flow rate and heat flux. It returns the set of optimum design variables and associated output parameters. The optimization is set up using an anonymous function of `moo_objective`: this acts as the black-box objective function evaluation. The MATLAB optimizer, `fmincon` is called, with additional inputs of initial design variables, upper and lower bounds, and a `constraints` function. Following the optimization, the optimum design variables are returned as `xsol`, and the `write_outputs` function creates a record of the optimization.

`moo_objective` The purpose of this function is to take the 18 geometry design variables as input, along with all the parameters that describe the cooling plate system, and return a single scalar objective function. This requires calling the `build_design_var` and `fun_eval` functions in sequence before selecting the appropriate objective function determined by the main program input parameter.

`build_design_var` The inputs to this function are the vector of design variables plus the coolant flow rate and heat flux. The `pattern` function recalls the cooling plate configuration parameters, then these are all combined with meshing parameters in a single vector that defines everything required to run the simulation through GAMBIT and FLUENT.

`pattern` This function acts as a catalogue of cooling plate configurations, initial dimensions, bounds, and constraints.

`fun_eval` This function is called with the full set of design variables and controls the execution of GAMBIT and FLUENT, including the file input/output. It first calls the `write_gambit_parameters` function to write a file specifying the geometry and mesh parameters and then executes GAMBIT from the command prompt. It then calls `run_fluent` with an error catcher that detects problems with the license or solver. The final part of the function is to open the FLUENT output files and read in the parameter values.

`run_fluent` To compile the FLUENT input file, the file containing a set of standard instructions is read, then combined with the FLUENT analysis parameters in a separate file. There are some modifications to account for different levels of accuracy and computing set-up, then FLUENT is then executed.

`write_gambit_parameters` This function takes the full set of design variables and writes the file of GAMBIT parameters. Two parameters describe the thickness of the plate and depth of channel; nine describe the mesh density and boundary layers; the solid and fluid sections are defined by a binary series of length determined by the plate configuration.

`write_outputs` This function appends the optimization results to a text file.

3.4.3 Optimization Settings

The `fmincon` function takes several options that determine the behaviour of the optimization algorithm. In many cases, the default settings give satisfactory results, but

there are some settings that were changed to improve the performance.

The default settings assume that the optimization objective is a perfectly smooth function of the design variables; however, Figure 3.9 showed that this is not the case for CFD output. Therefore, the perturbation in design variable used for finite differencing was increased. In the example shown, to reliably capture the downward trend with the *second* order CFD solution, a design variable step of at least 0.1 mm would be required. The *third* order solution is smoother, but still has fluctuations, and so a minimum design variable change of 0.05 mm was specified.

Similarly, the objective function tolerance was increased from the default. It is likely that there will always be a small degree of variability in the objective function, and maintaining a disproportionately small termination tolerance forces the optimizer to keep looking for ever smaller improvements. However, specifying a termination tolerance that is too large may cause the optimization to terminate prematurely. Testing of the optimization routine found that it was common to have design iterations of negligible objective function change followed by further improvement. Therefore, the objective function tolerance was set at a small value to prevent the optimization from stopping in this situation: a normalized limit of 0.00001 was selected (i.e. 0.001% of the initial objective function value).

With a small design convergence criterion, the duration of the optimization is largely determined by specifying the number of design iterations. The optimization behaviour indicated that the converged value sometimes required up to 100 design iterations to be reached, and therefore a three-step process was implemented to decrease the duration of a complete design optimization.

The first two steps use relaxed CFD convergence criteria and a mesh with fewer

Table 3.5 – Variable accuracy levels implemented in design optimization. The time taken for a single objective function can vary depending on the design being analyzed.

Accuracy	Mesh Size (elements)	CFD Convergence Criterion	Time for one objective function evaluation
Low	70976	10^{-5}	120 s
Medium	116672	10^{-6}	360 s
High	243560	10^{-7}	1200 s

elements, then the third step uses the verified settings described below. By this process, the early iterations can quickly find a design in the vicinity of the optimum, and only then is the most accurate and time consuming level of accuracy used. Table 3.5 gives a summary of the three levels of accuracy and indicates the computing time required.

A regime was used running 60 iterations at low accuracy, 40 at medium, and 30 at high, for a total of 130 design iterations taking approximately two weeks for a complete optimization. Several different computer systems were utilized for the optimizations, all with processors of a similar specification: multi-core, and a speed of approximately 3.0 GHz. All the CFD simulations were run with four-core parallelization.

3.4.4 Initial Designs

To mitigate the risk that the optimization algorithm converges on a local minimum, the optimization was performed several times with different initial designs. With 18 design variables and a time-consuming optimization process, it was not possible to use a comprehensive set of initial designs, so a representative random sample was chosen using a technique called Latin hypercube sampling (LHS) [54]. This is a sampling

method designed to select a small, well-distributed sample from a multi-dimensional design-space.

Each design variable is divided into M divisions, and the sample of M designs is chosen ensuring that each division of each variable is represented. In this case, the design variables are not independent; i.e. they must satisfy the topological constraints, so a normalization and scaling step was used. A sample size of eight was selected, and the LHS designs are shown in Figure 3.11. Note that across the eight designs, every channel section displays a range of widths, but the composition of any one design is random.

Due to time and computing constraints, it was not possible for every optimization in this study to be executed eight times with the LHS set of initial designs. However, by applying them to the three single objective optimizations at the standard conditions, it was possible to gauge the consistency and reliability of the optimization process.

3.5 Non-Uniform Heat Flux Profile

The method described above was derived under the assumption that the heat generated by the battery cell is spatially uniform. However, previous studies of battery heat generation [27, 55] indicate that the heat is concentrated in the vicinity of the cell connector tabs. To test whether the heat flux distribution influences the optimum cooling plate geometry, the single objective T_{avg} and T_{σ} optimizations were repeated with a non-uniform heat flux profile. With the exception of a small effect from temperature-dependent fluid properties, the P_{fluid} objective function should be unaffected by the heat flux distribution, and so this objective function was excluded

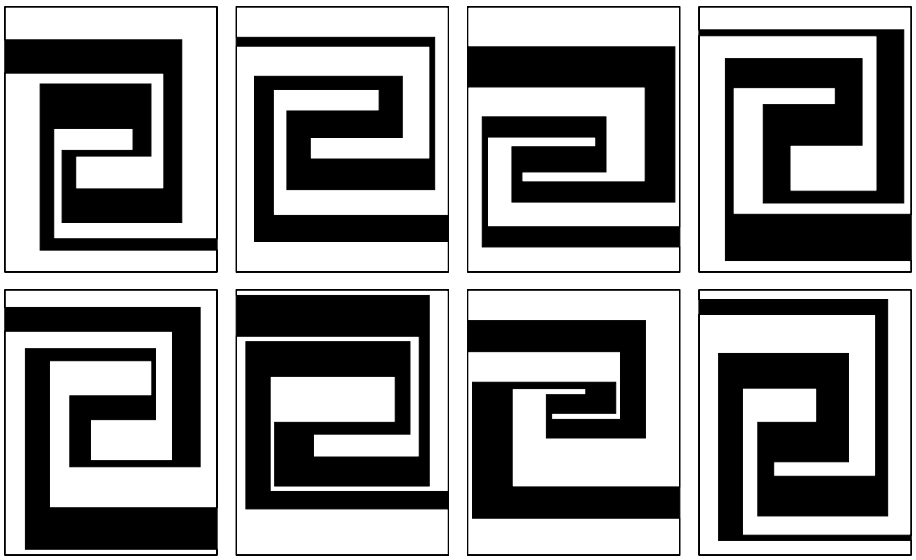

Figure 3.11 – Eight initial designs selected by Latin hypercube sampling. Top row: designs 1–4; bottom row: designs 5–8. The coolant channel is shown in black.

from the study.

3.5.1 Battery Cell Temperature

The work by Kim *et al.* [55] describes the thermal modelling of a lithium-ion polymer battery cell and presents a temperature plot of the cell after a period of discharge. This image was stretched to match the aspect ratio of the cooling plates in this study and digitized on a 10 mm grid. A bi-variate cubic polynomial (Eq. 3.1) was fitted to the digitized profile using linear regression. Figure 3.12 displays the progression from the thermal image captured to the polynomial function.

The steady state relationship between temperature and heat generation in a three-dimensional conductive body is described by the heat equation, Eq. 3.2, in which k is the thermal conductivity (W m^{-1} K^{-1}) and q the volumetric heat generation (W m^{-3}).

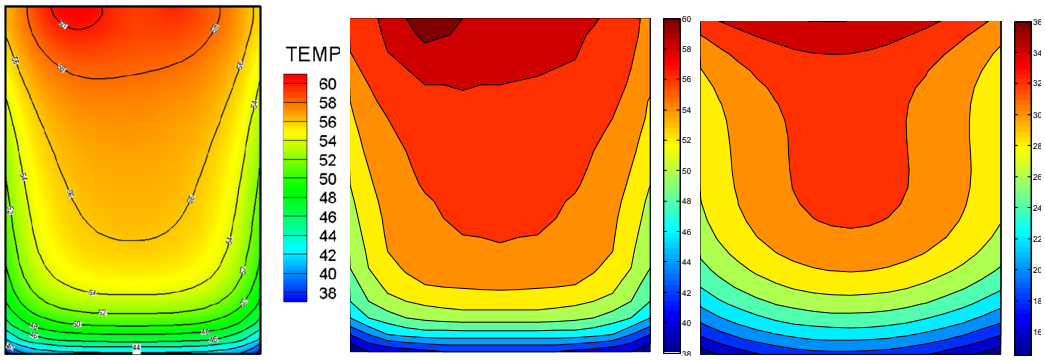

(a) Battery cell thermal modelling; from Kim *et al.* [55]

(b) Sampled temperature distribution

(c) Cubic polynomial temperature distribution

Figure 3.12 – Battery cell temperature distributions.

An approximation to the heat generation distribution was derived by differentiating the temperature distribution twice with respect to x and y (neglecting unknown convection and transient effects), resulting in a linear distribution in x and y (Eq. 3.3). Although there is some dependence on x, it is primarily a linear distribution in the y direction. This approximation predicts a small negative heat generation at the bottom of the cell, but if this is rounded up to zero, then the distribution of heat generation can be represented by a linear gradient increasing from zero at the bottom edge of the plate.

$$T(x,y) = 553x^3 + 7386y^3 + 919yx^2 - 949xy^2 - 892x^2 - 2627y^2 + 126x + 321y + 40 \quad (3.1)$$

$$k\frac{\partial^2 T}{\partial x^2} + k\frac{\partial^2 T}{\partial y^2} + k\frac{\partial^2 T}{\partial z^2} + q = 0$$
$$\text{also written:} \quad k\nabla^2 T + q = 0 \quad (3.2)$$

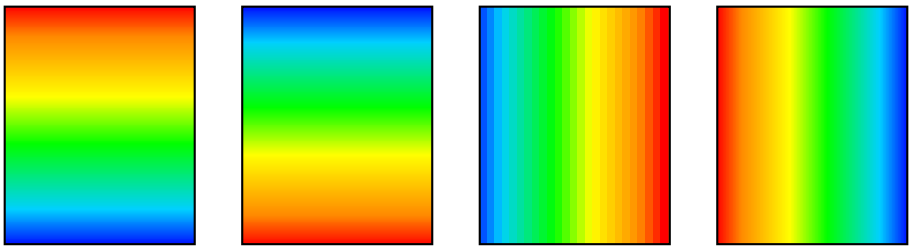

Figure 3.13 – Four linear heat flux gradients from 0 (blue) to 1000 W m^{-2} (red). From left to right: +y, -y, +x, -x (see Figure 3.8 for coordinate system).

$$q(x, y) = k(-1420x - 46148y + 7038) \tag{3.3}$$

3.5.2 Implementation of Non-Uniform Heat Flux

Having established that the distribution of heat generation within the battery cell is approximately a linear gradient, the scope of the investigation was increased by applying all four possible orientations of a linear gradient (positive and negative gradients in both the x and y directions), represented graphically in Figure 3.13. This can be considered equivalent to assessing the four different orientations of the cooling plate. All distributions were scaled such that the total heat flux was equal to a uniform 500 W m^{-2} boundary condition, i.e. the applied gradient was from 0 to 1000 W m^{-2}.

To incorporate the non-uniform heat flux profiles in the FLUENT CFD analysis, they were described by a C++ script in a user defined function (UDF) which was specified as the boundary condition.

3.6 Optimization Structure: Multi-Objective

The implementation of a multi-objective optimization routine requires only small modifications to the optimization structure described above. The fundamental difference is that the objective function, while still being a scalar, is composed of more than one of the performance measures. In this study, the weighted sum method was used to combine several objectives into a single function by summing the objective functions after applying a weighting factor to each one. In the general case, this gives Eq. 3.4.

$$J_{MOO}(\boldsymbol{d}, \boldsymbol{w}) = \sum w_i J_i(\boldsymbol{d}) \tag{3.4}$$

where J_{MOO} is the composite objective function composed of the separate J_i objective functions, each given a weight w_i.

In the case discussed in this study, there are three objectives functions, and for consistency there is an additional requirement that all weighting factors should sum to one. This gives the formulation in Eq. 3.5.

$$J_{MOO}(\boldsymbol{d}, \boldsymbol{w}) = w_1 \widehat{P}_{fluid}(\boldsymbol{d}) + (1 - w_1)\big(w_2 \widehat{T}_{avg}(\boldsymbol{d}) + (1 - w_2)\widehat{T}_{\sigma}(\boldsymbol{d})\big) \tag{3.5}$$

where w_1 and w_2 can assume any value between 0 and 1 to control the relative weighting of the three single objectives. The problem can be simplified by setting w_1 to 0, or w_2 to either 0 or 1, and only varying one of the weighting factors to give a bi-objective system.

In this formulation, it can also be seen that the relative magnitude of the individual

objective functions can influence the behaviour of a multi-objective optimization. i.e. if one objective function is very much larger than the others, it will be preferentially optimized unless the weighting factors are very carefully selected. To avoid this situation, all objective functions were normalized (divided by their initial value) before calculating the composite objective function, indicated by $\widehat{\ }$.

After defining the objective function, the only other modification is to run the optimization several times over a range of weighting factors to determine the Pareto front. This study was limited to bi-objective optimization, and used a total of six weighting factors; i.e. $w = 0.0, 0.2, 0.4, 0.6, 0.8, 1.0$.

3.7 Optimization Verification

There are several aspects of the optimization process for which verification was performed. In some cases, the accuracy of the process has no impact on the optimization results; in other cases accuracy is desirable but not essential; and in other cases, it has a direct outcome on the results:

- The absolute accuracy of the CFD solution is clearly desirable, but local phenomena such as velocity and flow direction are not measured and their influence on the objective function may be negligible.

- The outputs of the CFD simulation that determine the value of the objective function (pressure, mean temperature, temperature uniformity) were the focus of much of the verification effort. By verifying their accuracy, the confidence in the optimization process was greatly increased. However, note that even if there was some error in these outputs, the resultant optimized design may still

be valid.

- The optimization algorithm itself can influence the outcome if inappropriate settings are used; for example, the finite differencing and convergence parameters.

3.7.1 CFD Analysis Convergence

A CFD analysis is an iterative process, and so accurate results will only be obtained when appropriate convergence criteria are used. The most effective way to select the point at which the solution terminates is to specify limits on the calculation residuals: these outputs decrease as the analysis converges, and by using appropriate criteria, the analysis can be stopped when the required accuracy is reached.

To select the appropriate residual limits at which to terminate the solution, several designs were analyzed while monitoring the simulation outputs and residual values. Figure 3.14 shows the results from the reference design: plots of the three objective functions, and a plot of the five residuals. Of the objective functions, the pressure was the last to converge (to within 0.001% of the ultimate value) at an iteration around 250. Cross referencing this iteration with the value of the residuals indicated that this point can be characterized by the following residual values:

Continuity	4.70×10^{-6}
X–velocity	1.06×10^{-6}
Y–velocity	8.22×10^{-7}
Z–velocity	3.32×10^{-9}
Energy	3.73×10^{-10}

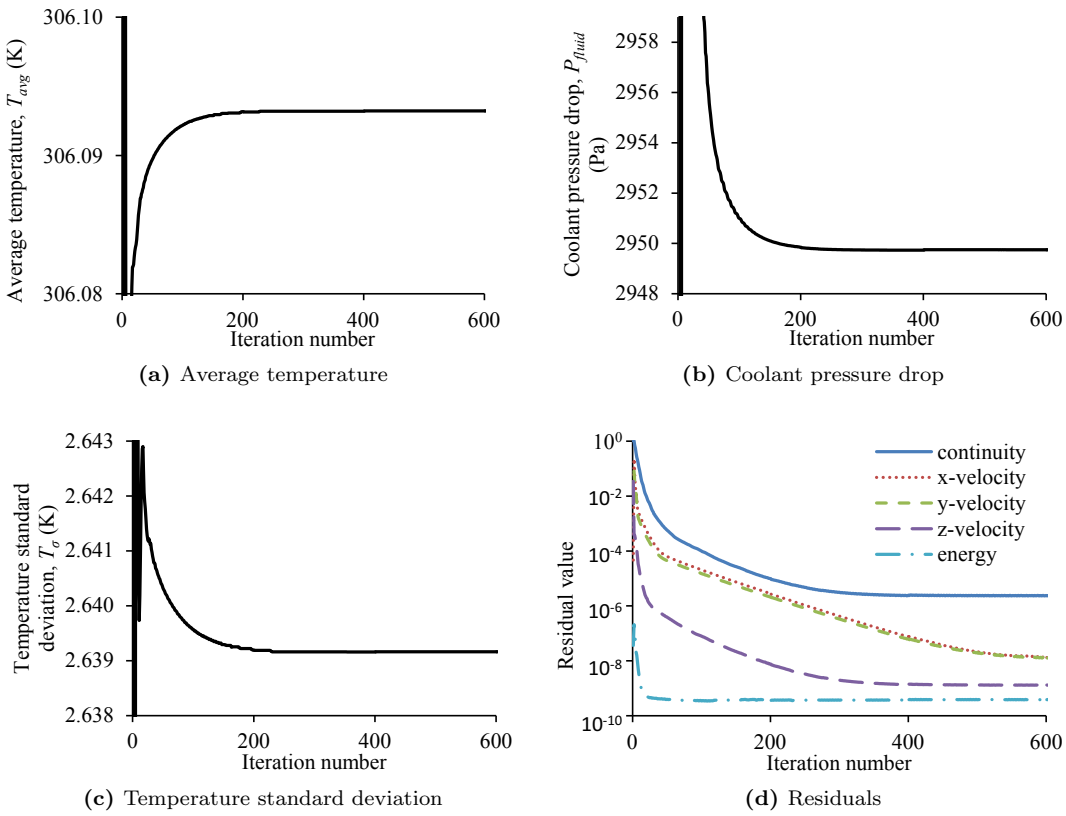

(a) Average temperature **(b)** Coolant pressure drop

(c) Temperature standard deviation **(d)** Residuals

Figure 3.14 – The convergence of objective functions and solution residuals during a CFD analysis.

The residuals plot, Figure 3.14d, shows that the energy residual converged very quickly, and is flat by iteration 250: therefore unhelpful for identifying the termination point. The continuity and z-velocity residuals could be used, but by this point they were also converging. The most accurate residuals to use are the x- and y-velocity: at iteration 250 they were consistently decreasing; specifying a criterion that both x- and y-velocity residuals should be below 1.0×10^{-6} would stop the analysis at iteration 254.

However, different designs display different convergence behaviour, and Table 3.6 shows the corresponding outputs from all three analyzed designs. Of the three designs assessed, the T_σ optimized design was slowest to converge, and so dictates the convergence criteria to be used. Based on this design, the criterion that permits convergence of all three output parameters is a value of 1.2×10^{-7} on the x-velocity residual and 2.1×10^{-7} on the y-velocity residual. For the design optimization process an additional allowance was added to decrease both limits to 1.0×10^{-7}.

An additional convergence check was made to examine the balance of heat flow and mass flow in the converged model; in both cases the net imbalance should be negligible. These were assessed for a sample of CFD analyses and in all cases the heat flow imbalance was found to be less than 0.00001% and the mass flow imbalance was zero.

3.7.2 Mesh Convergence

The accuracy of a CFD analysis can be improved by increasing the number of elements used in the mesh. However, this also increases the computing cost, and for

Table 3.6 – Iterations to convergence of output parameters for three tested designs, and magnitude of residuals at last point of convergence.

		Geometry		
		Reference	T_{avg} Optimized	T_σ Optimized
Iterations to convergence of:	Pressure	250	236	298
	T_{avg}	198	188	368
	T_σ	228	272	437
Residual at last converged objective	Continuity	4.70×10^{-6}	2.72×10^{-6}	3.23×10^{-6}
	X-velocity	1.06×10^{-6}	6.57×10^{-7}	1.23×10^{-7}
	Y-velocity	8.22×10^{-7}	4.64×10^{-7}	2.09×10^{-7}
	Z-velocity	3.32×10^{-9}	3.07×10^{-9}	1.70×10^{-9}
	Energy	3.73×10^{-10}	3.86×10^{-10}	4.47×10^{-10}

optimization (which requires hundreds of function evaluations) this is a critical consideration. Therefore, a study of objective function convergence with respect to the mesh density was performed to determine a mesh size that results in a sufficiently accurate result with acceptable computing cost.

The two domains, solid and fluid, were assessed separately, as their requirements are different. In the solid domain, where there is only heat conduction, very few elements were required for a fully converged solution—four divisions on each parametric volume proved sufficient. To assess the requirements in the fluid domain, an automated routine was used to generate models with between 30 000 and 3 000 000 elements in the fluid domain. Each model was analyzed using the standard boundary conditions and the outputs of P_{fluid}, T_{avg} and T_σ were recorded. These outputs are plotted against the mesh size in Figure 3.15.

The behaviour of P_{fluid} (Figure 3.15a) displays a typical convergence trend to a steady value of around 2975 Pa with a model size of approximately 3 million elements.

The behaviour of the mean temperature and the temperature standard deviation are both less consistent than the pressure output, but also display relatively steady output with meshes of between 2 and 3 million elements. However, with a mesh of this size, every function evaluation would take several hours, making the optimization process too time consuming. For a faster analysis and output still within 1% of the converged value, a mesh of around 250 000 elements was proposed.

Further convergence testing with different plate geometries and different coolant flow rates confirmed that the analysis outputs displayed similar convergence behaviour, and so a mesh of 243 560 elements was used for all optimizations.

3.7.3 Turbulence

The onset of turbulence in fluid flow is predicted by the magnitude of the Reynolds number. This dimensionless constant is calculated by the equation UD/v where U is the flow velocity, D is the hydraulic diameter of the channel and v is the kinematic viscosity of the coolant. In the optimization of the cooling plate, the channel width varies depending on the value of the design variables, thereby varying the Reynolds number. The other parameter that affects the Reynolds number is the coolant flow rate. Table 3.7 shows the Reynolds number calculated at several flow rates and channel widths used in this study.

The critical Reynolds number at which flow transitions from laminar to turbulent is influenced by several factors including the shape of conduit, the wall roughness, and additional disturbances. The general case of a rectangular conduit was studied by Hanks and Ruo [56] who determined that, depending on the aspect ratio, the critical Reynolds number is between 1900 and 2800. On this evidence, all the conditions listed

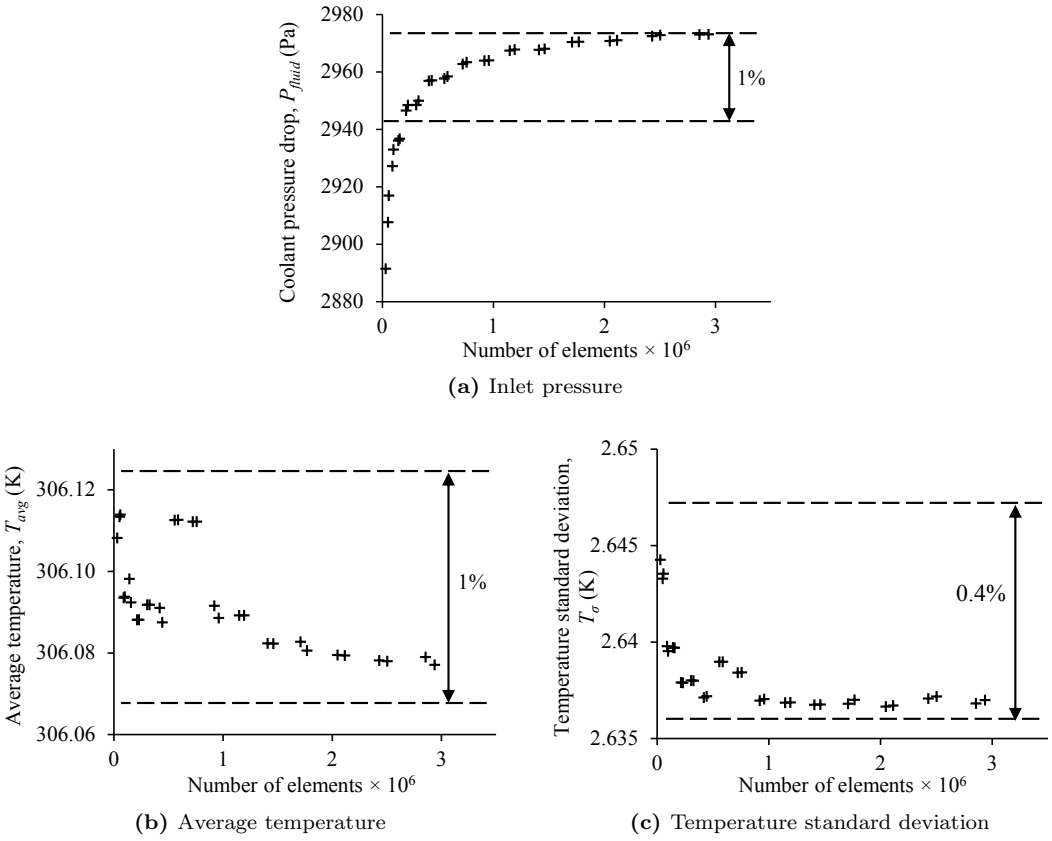

(a) Inlet pressure

(b) Average temperature

(c) Temperature standard deviation

Figure 3.15 – Mesh convergence plots.

Table 3.7 – Reynolds number at selected coolant flow rates and channel widths. The calculation assumes a fluid temperature of 300 K, which leads to a kinematic viscosity of 2.96×10^{-6} m^2 s^{-1}.

		Coolant flow rate (kg s^{-1})		
		0.0004	0.001	0.002
Channel width (mm)	2	92.35	230.88	461.76
	20	12.24	30.60	61.20
	40	6.23	15.58	31.16

in Table 3.7 should be purely laminar. However, the sharp corners of the channel have the potential to reduce the transition point.

Martin *et al.* [57] examined the flow in a square U-bend of a channel with a square cross section at Reynolds numbers between 109 and 872. They found that at Reynolds number greater than 436, flow structures became time-dependent and therefore unable to be accurately simulated with a steady-state analysis. This is the precursor to turbulent flow. Caution must be used when applying these results to the present study: the channel aspect ratios are quite different, and the onset of instability was thought to be influenced by the proximity of the two right-angled bends. Nonetheless, it provides some confidence that an assumption of laminar steady flow is reasonable at the conditions used in this study.

Chapter 4

Cooling Plate Optimization: Results

4.1 Optimization Behaviour

The behaviour of the optimization algorithm was assessed by monitoring the three single objective optimizations from the eight initial designs. Figures 4.1, 4.2, and 4.3 show the progress of the T_{avg}, T_{σ}, and P_{fluid} objective functions respectively over the course of the optimization process from the eight initial designs. All objective functions behaved similarly, with a rapid decrease in the first 20 design iterations followed by gradual improvement—sometimes beyond 100 design iterations. The T_{σ} optimizations were slightly inconsistent in their convergence behaviour, with several examples of the algorithm pausing on an intermediate design before continuing to the converged design. The converged objective functions had a variation of around 2.5% in the case of T_{avg} optimization and 1.0% for the T_{σ} and P_{fluid} optimizations.

Figures 4.4, 4.5, and 4.6 show the geometry of the eight converged designs for the

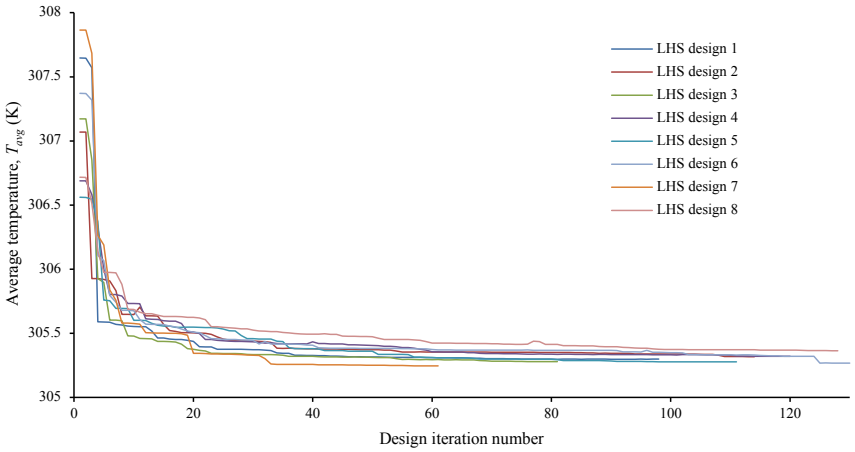

Figure 4.1 – Convergence of T_{avg} objective function from eight initial designs.

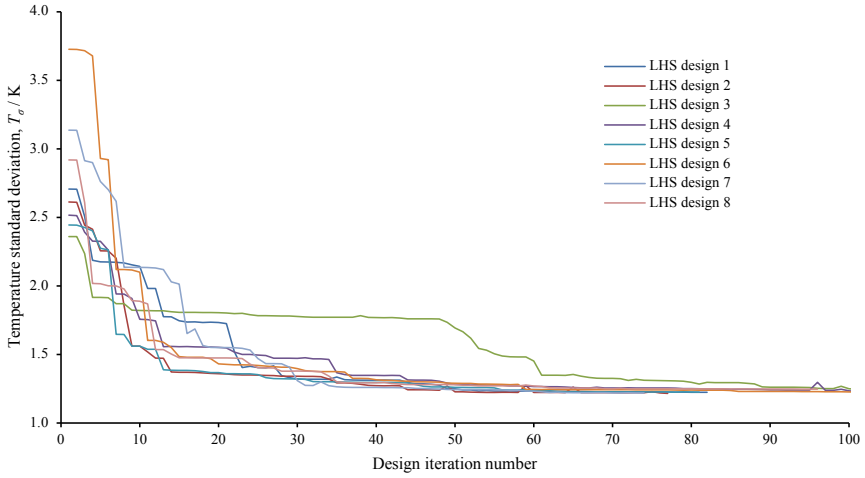

Figure 4.2 – Convergence of T_σ objective function from eight initial designs.

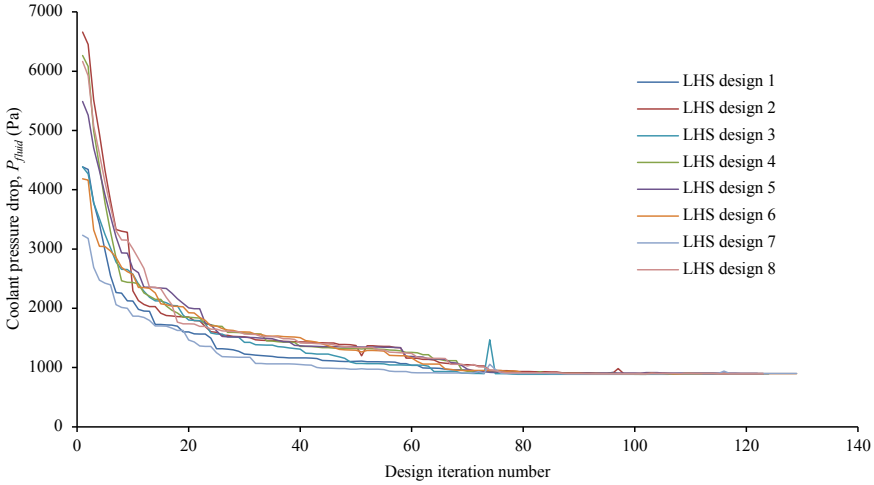

Figure 4.3 – Convergence of P_{fluid} objective function from eight initial designs.

three objective functions, and Figure 4.7 shows the three sets superimposed so that similarities can be more easily seen.

The converged geometry displays a much greater variability than the objective functions would suggest. The T_{avg} converged designs (Figure 4.7a) have almost no coincident geometry, except the outer boundary. However, they all share the feature that the coolant channel occupies close to the maximum possible space. The equivalent T_{σ} figure shows that after the two channel sections closest to the inlet, there is also significant variation between designs. The P_{fluid} optimizations are the most consistent—both in the convergence behaviour, and the converged designs. However, the geometric variability is still significantly greater than the 1% variability in the objective function. This indicates that close to the optimum, the objective functions become quite insensitive, and there are many local optimum designs that have almost identical performance. However, the optimization algorithm is not sufficiently precise to positively identify the global optimum.

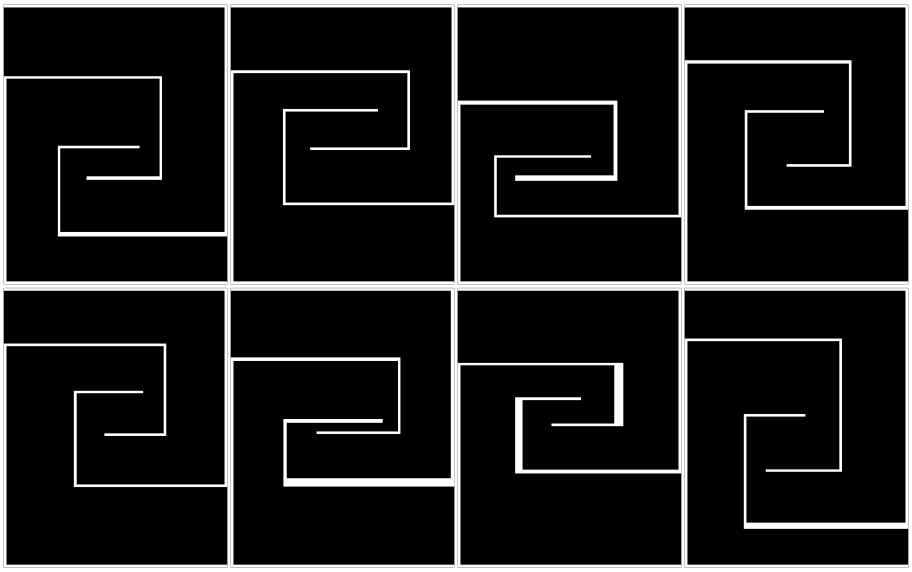

Figure 4.4 – Converged T_{avg} geometry from eight initial designs. The channel is shown in black.

In the following discussions of optimization results, only the design that achieved the lowest objective function is presented.

4.2 Reference Design Analysis

To measure the improvement following optimization, the reference design was assessed with the same procedure as the optimization routine; at a mass flow rate of 0.001 kg s^{-1} and heat flux of 500 W m^{-2}. Figure 4.8a shows the distribution of temperature on the back of the cooling plate, and Figure 4.8b shows the temperature on the symmetry face: this captures both fluid and solid regions. The effect of the coolant can clearly be seen: it enters the plate at the top left corner at a low temperature and exerts a strong cooling force. As it travels through the plate, it absorbs heat and its temperature

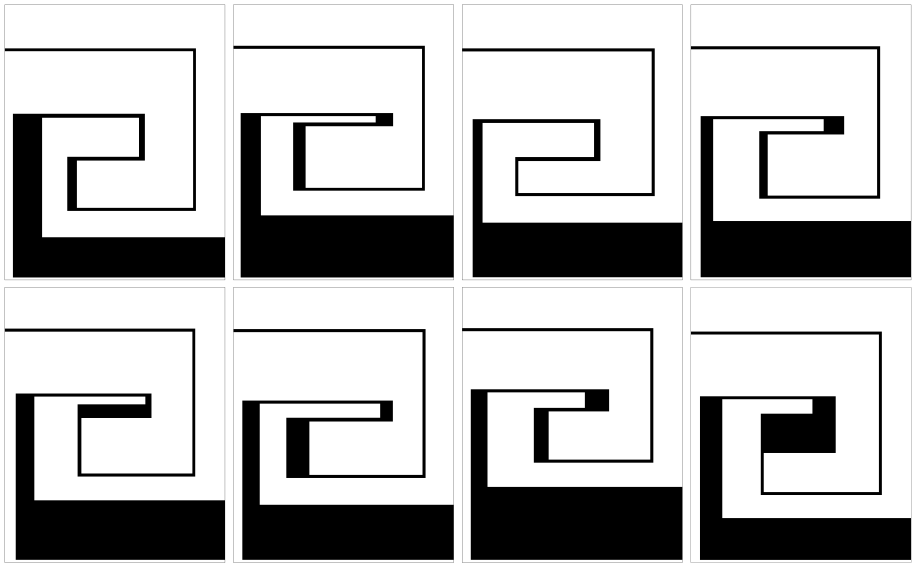

Figure 4.5 – Converged T_σ geometry from eight initial designs. The channel is shown in black.

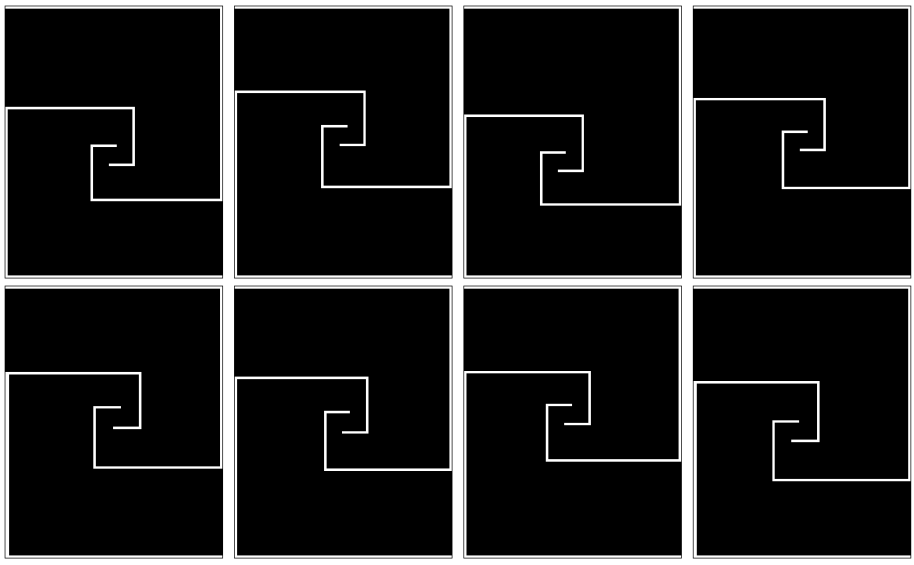

Figure 4.6 – Converged P_{fluid} geometry from eight initial designs. The channel is shown in black.

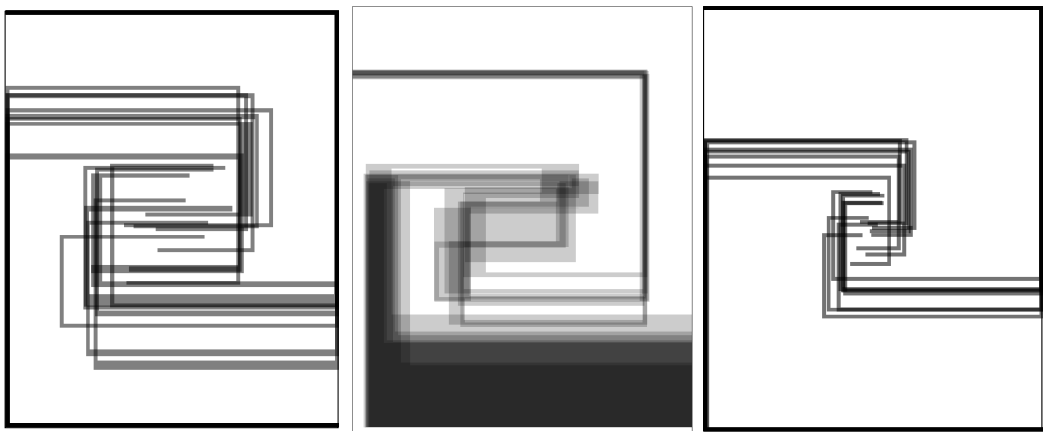

(a) T_{avg} designs **(b)** T_{σ} designs **(c)** P_{fluid} designs

Figure 4.7 – Converged geometry from eight initial designs shown superimposed. For clarity, the T_{avg} and P_{fluid} designs are shown with the coolant channel in white; the T_{σ} designs are shown with the coolant channel in black.

rises, before leaving the plate from the bottom right corner approximately 9 K warmer. This creates a general temperature gradient in the plate from top to bottom.

The temperature of the coolant can be seen in Figure 4.8b: it is only marginally lower than the back face of the cooling plate. This is confirmed by Figure 4.10a which plots the temperature through the thickness of the plate at three locations: the minimum coolant temperature is 0.6-0.8 K lower than the plate temperature at the same location. From the inlet to outlet, the coolant increases in temperature, but the temperature gradient through the thickness is similar, indicating that the local rate of heat transfer into the coolant is approximately equal at different locations in the plate.

The flow behaviour of the coolant is shown in Figure 4.9: Figure 4.9a shows the velocity distribution, and Figure 4.9b the pressure drop. With the exception of the channel corners, the velocity is quite uniform at about 0.1 m s^{-1}. The pressure

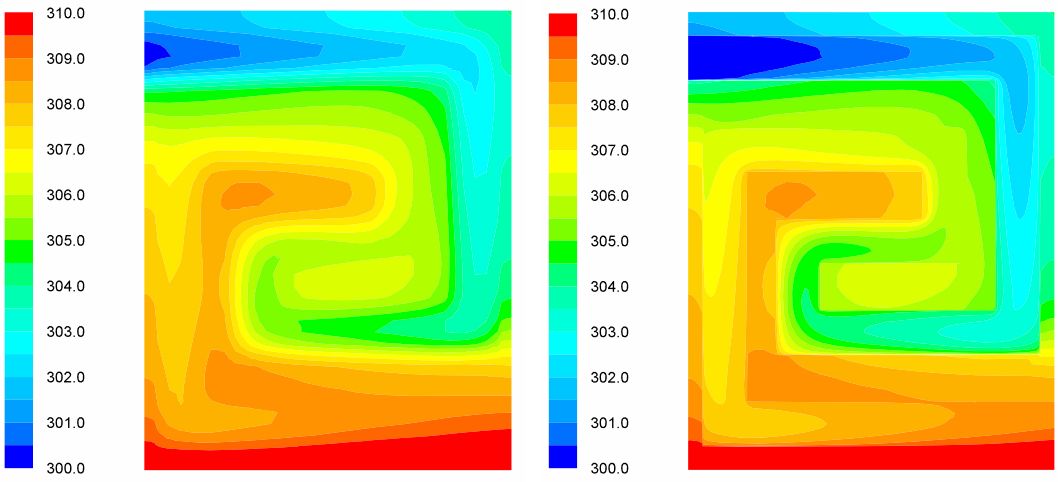

(a) Back of cooling plate **(b)** Centre of cooling plate

Figure 4.8 – Temperature distribution on the reference geometry cooling plate.

decreases constantly from 2951 Pa at the inlet to zero at the outlet; by visual examination of Figure 4.9b, there appears to be no significant additional pressure drop at the corners: this is consistent with the findings of Maharudrayya *et al.* [35] at similar Reynolds numbers.

The output parameters at these conditions are shown in Table 4.1 alongside the results of the single objective optimizations.

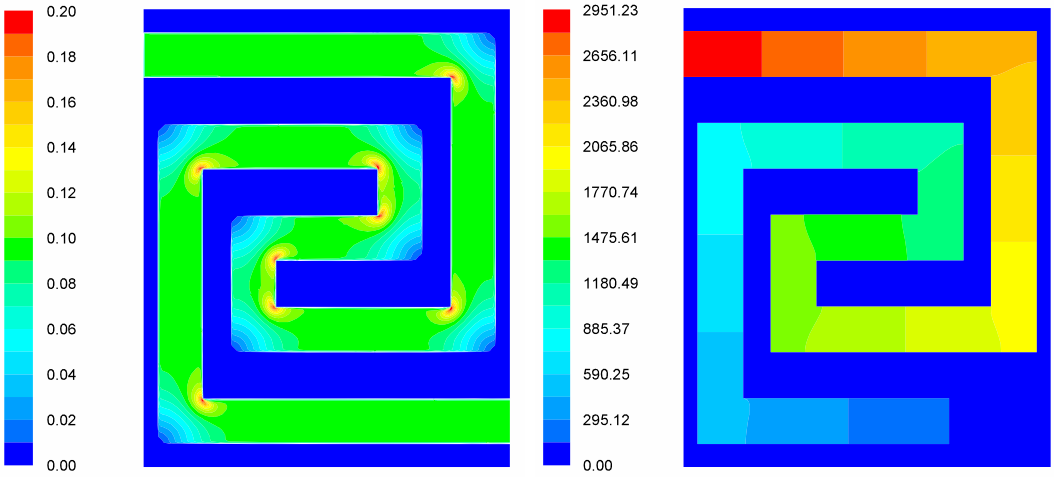

(a) Velocity of coolant **(b)** Total pressure of coolant

Figure 4.9 – Coolant behaviour of the reference geometry cooling plate. Values are shown on the symmetry face.

4.3 Single Objective Optimization

4.3.1 Results

Reference Boundary Conditions

The optimization procedure described in the previous section was executed for each of the three objective functions at boundary conditions of 0.001 kg s^{-1} and 500 W m^{-2}. As discussed above, the starting point for these optimizations was the set of eight LHS initial designs. The nominal termination criterion for the optimizations was an objective function change of less than 0.001% between iterations. However, the optimization progress was often erratic, and user intervention was sometimes required to continue the optimization after premature termination. A more reliable criterion

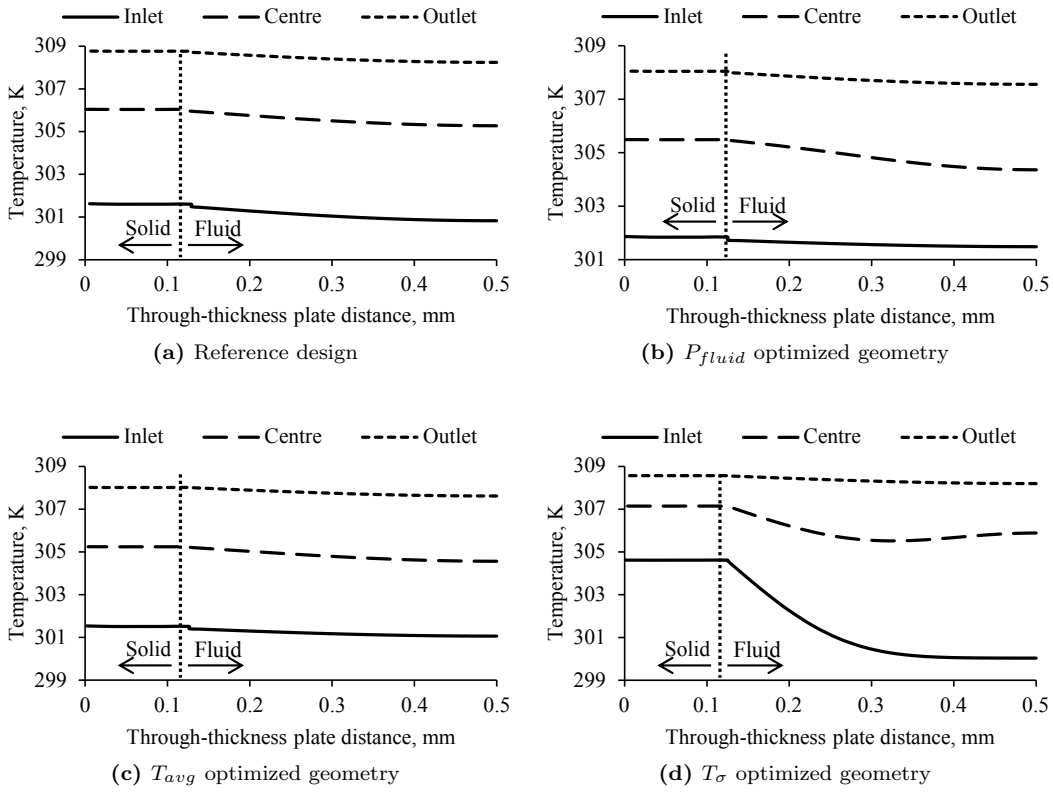

(a) Reference design

(b) P_{fluid} optimized geometry

(c) T_{avg} optimized geometry

(d) T_{σ} optimized geometry

Figure 4.10 – Temperature profiles measured in the through-thickness direction of the reference and three optimized cooling plates. The measurements are through the centre of the channel, and taken at three locations: near the inlet, the centre, and outlet of the plate.

would be to require several consecutive design iterations of negligible objective function improvement.

Table 4.1 shows the objective function improvements achieved in each case, in comparison to the reference design. The P_{fluid} and T_σ optimizations resulted in a greater than 50% improvement over the reference design in their respective objective functions; the T_{avg} optimization showed an improvement of around 14% (this percentage improvement is calculated relative to the coolant inlet temperature). All optimizations resulted in a reduced value of T_σ, but when using T_σ as the objective function, values of P_{fluid} and T_{avg} increased significantly.

The geometry of the optimized designs are shown in Figure 4.11; these clearly show the similarity between the P_{fluid} and T_{avg} designs, both of which have channels expanded to near maximum width, albeit with some difference in the central region. The T_σ optimized design has a channel which is mostly of minimum width, except for the last section of the channel which is quite wide.

The temperature distributions on the back face of the optimum cooling plates are shown in Figure 4.12. Here again ,the similarity between the P_{fluid} and T_{avg} optimized designs is apparent: both show a marked temperature gradient from top to bottom. The T_σ optimized design clearly has a more uniform temperature, although at a higher average: the narrow channel in the top half of the plate restricts the heat flow in this area, thereby equalizing the temperature over the whole plate.

More information on the temperature and heat transfer within the plates is found in Figure 4.10, which plots the through-thickness temperature of each plate at three locations. The temperature gradients in the P_{fluid} and T_{avg} optimized designs are very similar to the reference design: the temperature of the coolant is close to that

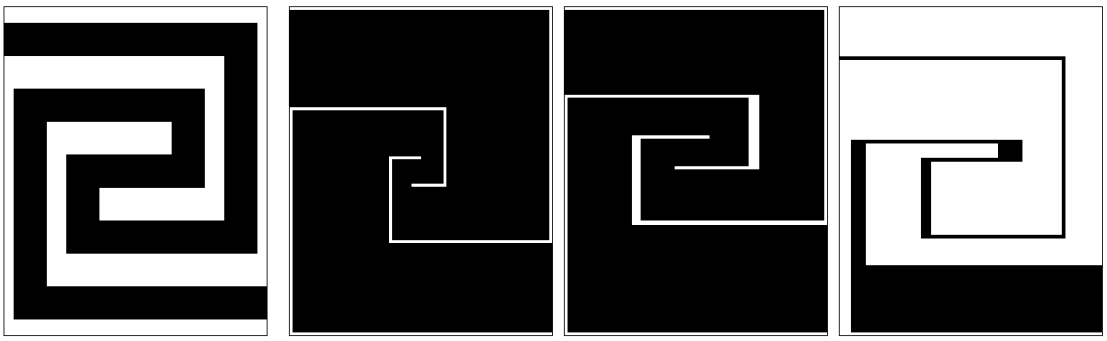

(a) Reference geometry **(b)** P_{fluid} optimized geometry **(c)** T_{avg} optimized geometry **(d)** T_{σ} optimized geometry

Figure 4.11 – Cooling plate geometry of reference design and single objective optimized designs.

Table 4.1 – Objective function values of optimized cooling plates compared with the reference design.

Geometry	P_{fluid} (Pa)	T_{avg} (K)	T_{σ} (K)
Reference	2949	306.09	2.64
P_{fluid} optimized	890	305.33	2.31
T_{avg} optimized	1232	305.25	2.32
T_{σ} optimized	25124	308.31	1.22

of the adjacent solid. In the T_{σ} optimized design, the effect of the narrow channels can be seen: at the first two sampling points, there is a significant difference between coolant and plate temperatures (3–4 K). The large temperature gradient indicates that in these regions, the local rate of heat transfer is very high. In the wide channel section before the outlet, the temperature has almost equalized, and the temperature gradient is negligible.

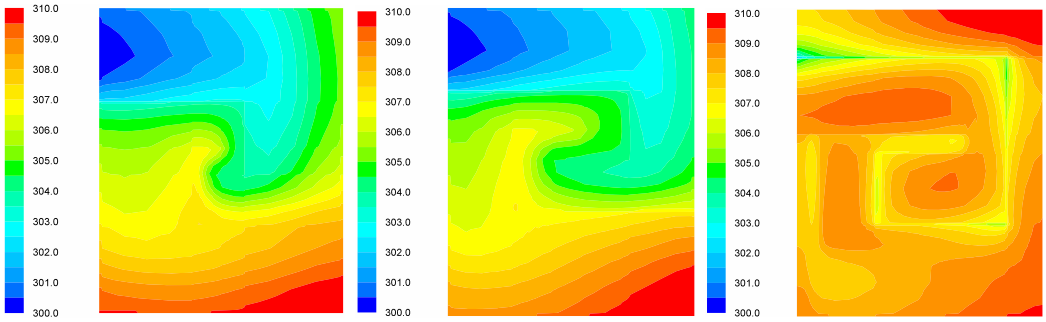

(a) P_{fluid} optimized geometry **(b)** T_{avg} optimized geometry **(c)** T_{σ} optimized geometry

Figure 4.12 – Temperature distribution on the back face of the optimized cooling plates.

Additional Boundary Conditions

The single objective optimization procedure was also applied at different values of coolant flow rate and heat flux. Due to computing constraints only one optimization of each was performed, using the reference design as a starting point. Table 4.2 shows the performance of all additional optimum designs.

To compare the cooling plate geometries optimized at different flow rates, Figures 4.13, 4.14, and 4.15 show the T_{σ}, T_{avg}, and P_{fluid} optimized designs respectively. The T_{σ} designs display a clear trend in the position of the top channel section: as the coolant flow rate increases, the top channel section moves closer to the top edge of the plate. The other channel sections also display some variation, but considering the insensitivity of this region (see Figure 4.7b) this may just be random variation. The T_{avg} optimized designs show some variation in geometry but no noticeable trends; unlike the P_{fluid} optimized designs which are almost identical.

The T_{σ} and T_{avg} optimized designs at different values of heat flux are shown in

Figures 4.16 and 4.17 (P_{fluid} was not used as an objective function as it was not considered to be significantly affected by the heat flux boundary condition). In both cases there seem to be no obvious trend in the designs—some variability is to be expected given the outcome of the study into the initial design dependence.

The clearest trend observed from these additional optimizations was in the T_σ designs at different flow rates. To better understand the mechanism driving the optimum design, the temperature distributions from the three designs, analyzed at their optimum coolant flow rate, are shown in Figure 4.18. As the coolant flow rate increases, it retains more of its cooling effect through the length of the channel: at the lowest flow rate, only the first channel section provides effective cooling, and so this takes a position closer to the centre of the plate. At the highest flow rate, effective cooling is provided for at least half the channel length, and so these sections are spread more evenly over the plate—meaning that the first channel section is pushed towards the edge of the plate.

If this trend can be extrapolated, then it indicates that at a sufficiently high flow rate (in which the outlet temperature is not significantly below the inlet temperature), the optimum design with respect to T_σ will be one in which the channels are of a uniform width and evenly distributed over the plate. In this case, the optimum T_σ design may start to more closely resemble the optimum designs of T_{avg} and P_{fluid}.

Off-design Boundary Conditions

During operation in an EV, a cooling plate will not operate at its design condition all of the time, therefore it is important to assess it at off-design boundary conditions. The other reason for this assessment was to determine if the cooling plate optimum

Table 4.2 – Performance of cooling plates optimized at additional boundary conditions. Results are for the performance at the conditions shown.

Objective	Heat flux (W m^{-2})	Flow rate (kg s^{-1})	P_{fluid} (Pa)	T_{avg} (K)	T_σ (K)
T_σ	500	0.0004	3130	319.33	2.95
	500	0.002	46901	304.11	0.75
	200	0.001	27345	303.28	0.49
	1000	0.001	20535	315.94	2.47
T_{avg}	500	0.0004	481	313.82	6.31
	500	0.002	3438	302.73	1.33
	200	0.001	1618	302.15	1.04
	1000	0.001	1283	310.55	4.95

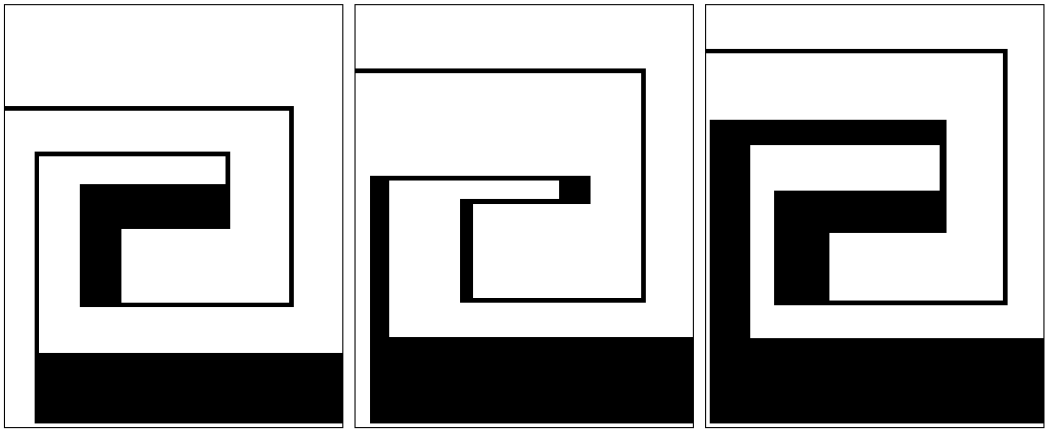

(a) 0.0004 kg s^{-1} **(b)** 0.001 kg s^{-1} **(c)** 0.002 kg s^{-1}

Figure 4.13 – Cooling plate geometry of T_σ optimized designs at three different coolant flow rates.

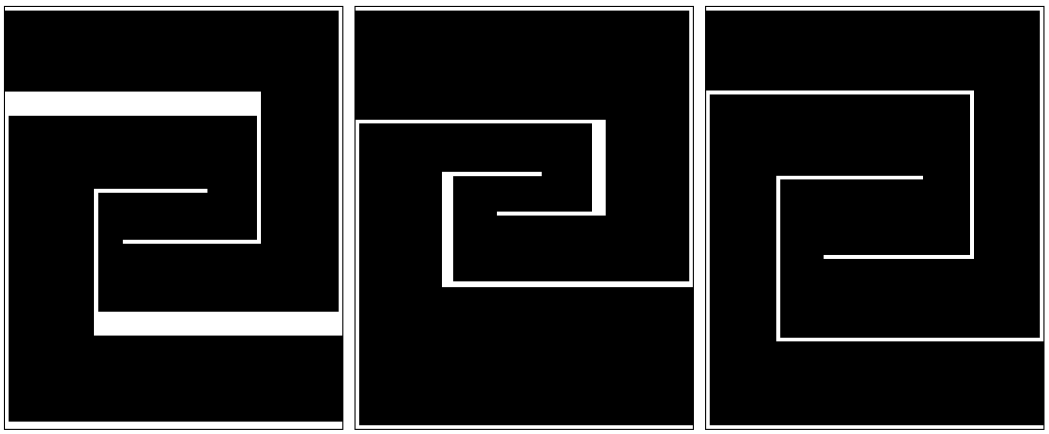

(a) 0.0004 kg s^{-1} **(b)** 0.001 kg s^{-1} **(c)** 0.002 kg s^{-1}

Figure 4.14 – Cooling plate geometry of T_{avg} optimized designs at three different coolant flow rates.

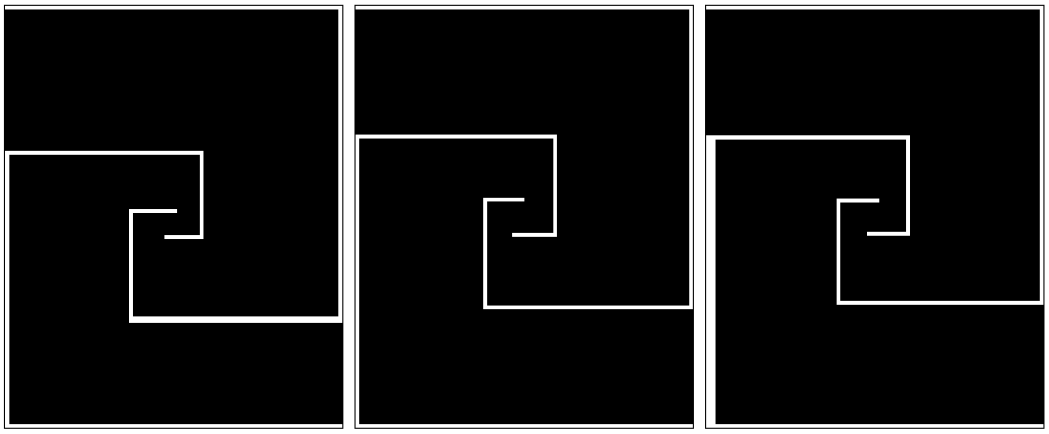

(a) 0.0004 kg s^{-1} **(b)** 0.001 kg s^{-1} **(c)** 0.002 kg s^{-1}

Figure 4.15 – Cooling plate geometry of P_{fluid} optimized designs at three different coolant flow rates.

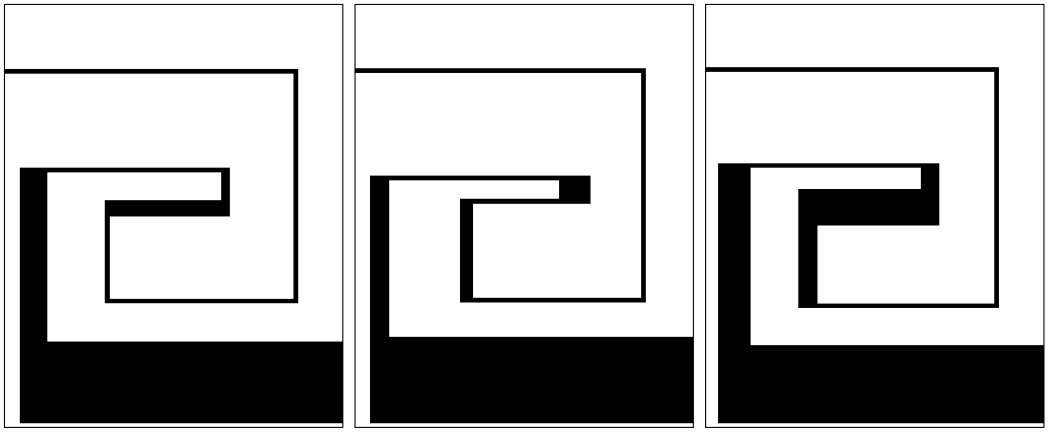

(a) 200 W m^{-2} **(b)** 500 W m^{-2} **(c)** 1000 W m^{-2}

Figure 4.16 – Cooling plate geometry of T_σ optimized designs at three different heat flux values.

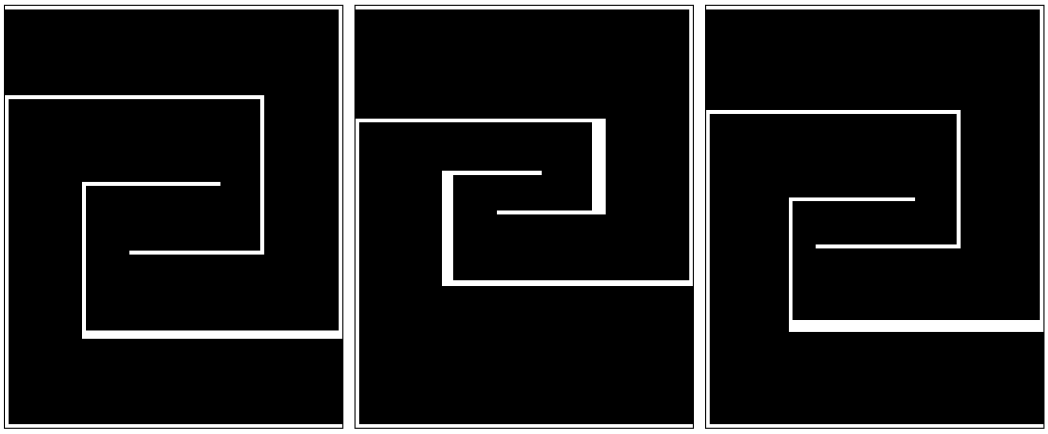

(a) 200 W m^{-2} **(b)** 500 W m^{-2} **(c)** 1000 W m^{-2}

Figure 4.17 – Cooling plate geometry of T_{avg} optimized designs at three different heat flux values.

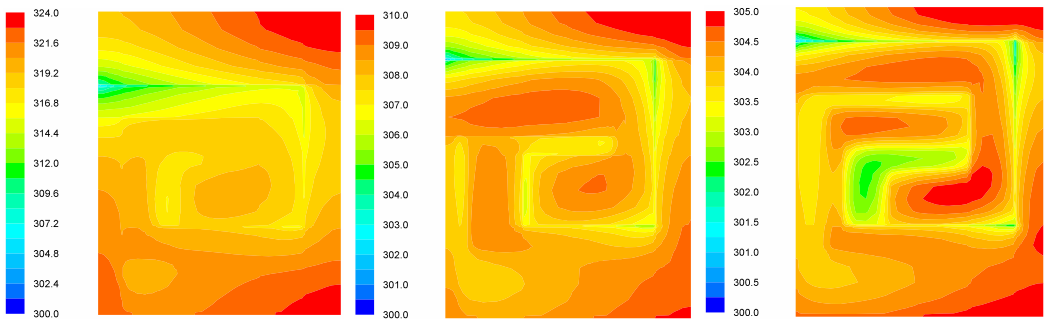

(a) 0.0004 kg s^{-1} **(b)** 0.001 kg s^{-1} **(c)** 0.002 kg s^{-1}

Figure 4.18 – Temperature distribution of T_σ optimized designs at three different coolant flow rates.

design is dependent on the boundary conditions.

The charts in Figure 4.19 show the performance (in terms of T_σ and T_{avg}) of the reference plate over a range of inlet flow rate and heat flux magnitude. Increasing the coolant flow rate reduces both the output functions of T_σ and T_{avg}, whereas increasing the heat flux increases both these parameters. The efficacy of the optimized designs is illustrated by plotting their performance normalized against the performance of the reference design; these charts are shown in Figure 4.20.

For example, Figure 4.20a shows that at low flow rates, the T_{avg} design optimized at 0.002 kg s^{-1} has the highest T_{avg} (about 95% of the reference design); the design optimized at 0.001 kg s^{-1} has the lowest T_{avg} (about 86% of the reference design). As the flow rate increases, all three optimum designs improve relative to the reference design, but the 0.002 kg s^{-1} optimized design improves the most until at a flow rate of 0.003 kg s^{-1}, it has the lowest T_{avg} (about 75% of the reference design). In this case, although there is some indication that a different design should be selected depending on the operating conditions, all three designs perform similarly.

However, the situation is different for the T_σ optimized designs, shown in Figure 4.20b. Here it can be seen that the design optimized at 0.0004 kg s^{-1} performs best at low flow rates, the 0.001 kg s^{-1} design performs best at flow rates around 0.001 kg s^{-1}, and at higher flow rates, the 0.002 kg s^{-1} design is superior. Furthermore, away from their optimization conditions their performance is degraded, although the magnitude of this effect differs among the designs: the performance of the 0.0004 kg s^{-1} design is restricted to low flow rates, but the 0.002 kg s^{-1} design has acceptable (better than reference) performance over the entire tested range.

Variation in the heat flux boundary condition has almost no effect on the normalized performance of the optimum cooling plates. The T_{avg} optimized designs have a performance at about 87% of the reference design, and this is constant at all heat flux boundary conditions. The T_σ optimized designs show a small variation in the normalized value of T_σ, but the design optimized at 500 W m^{-2} is the best of the three at all heat flux boundary conditions.

4.3.2 Discussion

From the geometries of the optimum cooling plate designs, there are several qualitative statements that can be made. The first is that cooling performance is increased by permitting a greater area for cooling channels: the P_{fluid} and T_{avg} designs both have a similarly low average temperature, and both have almost the maximum possible area occupied by the cooling channel. Increasing the width of the channel increases the surface area available for heat transfer, but at a constant mass flow it also reduces the velocity of the coolant, which reduces the local heat transfer rate. This has been confirmed by assessing the temperature gradients through the plate thickness.

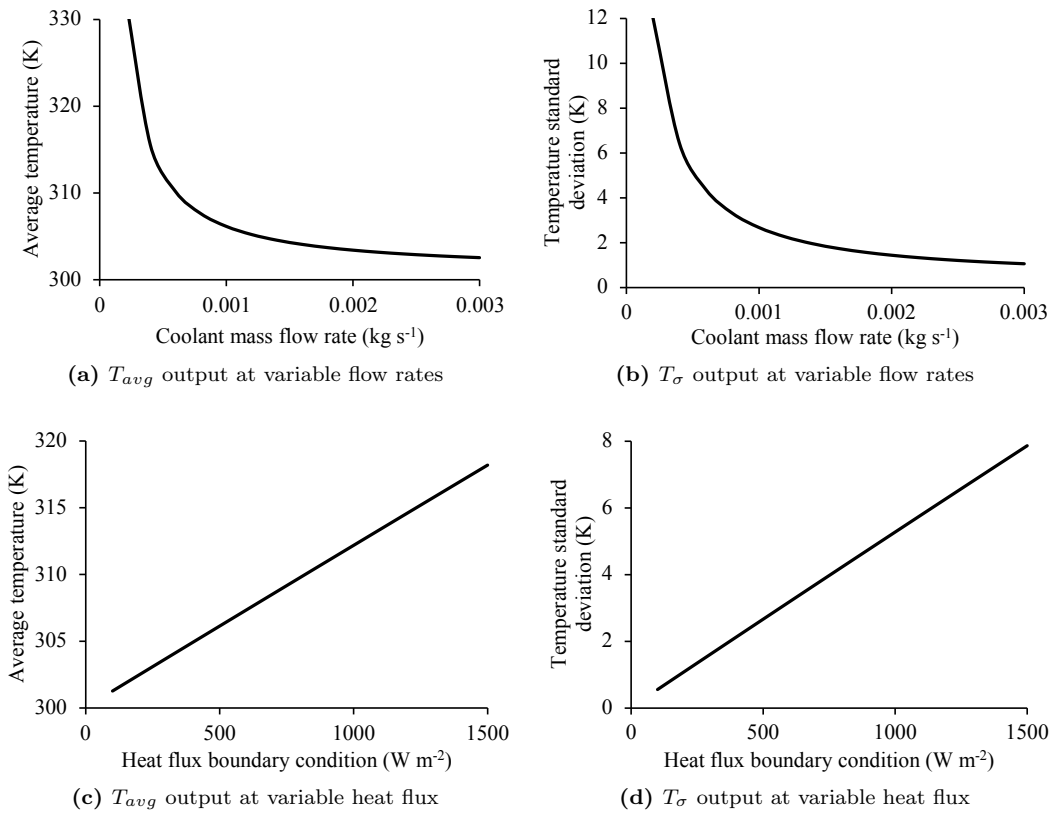

(a) T_{avg} output at variable flow rates

(b) T_{σ} output at variable flow rates

(c) T_{avg} output at variable heat flux

(d) T_{σ} output at variable heat flux

Figure 4.19 – Performance of the reference cooling plate at a range of flow rates and heat flux.

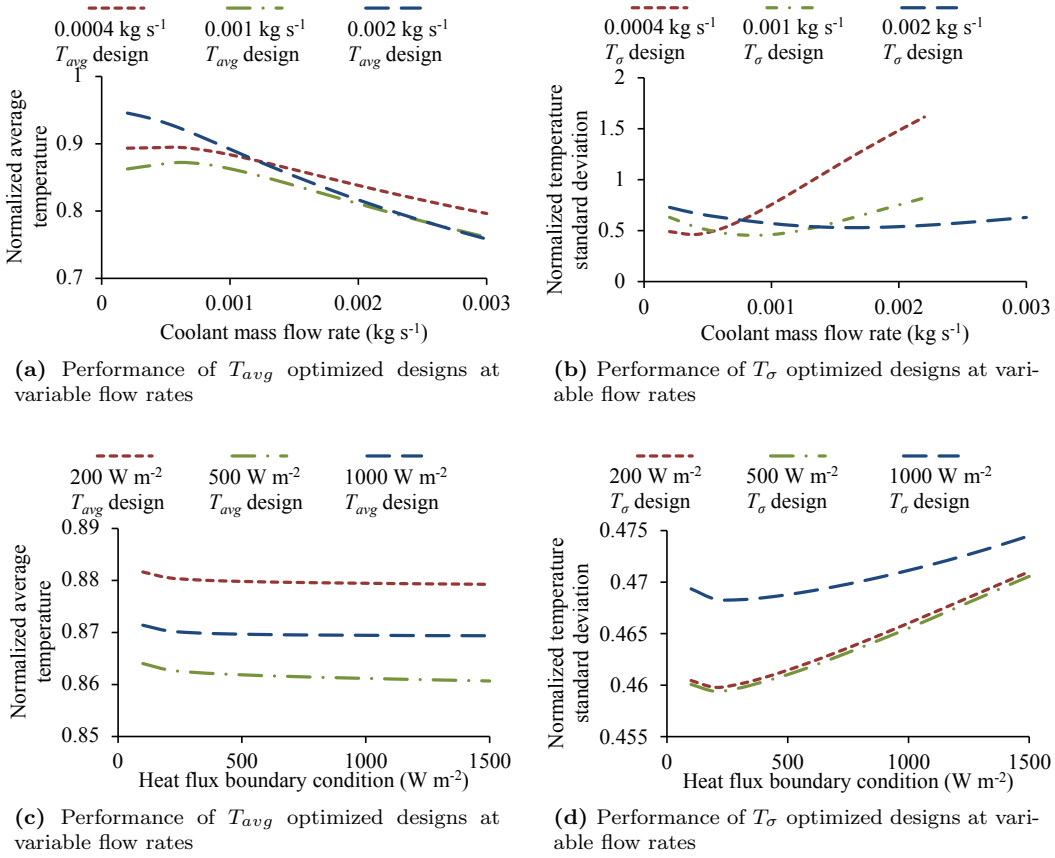

(a) Performance of T_{avg} optimized designs at variable flow rates

(b) Performance of T_{σ} optimized designs at variable flow rates

(c) Performance of T_{avg} optimized designs at variable flow rates

(d) Performance of T_{σ} optimized designs at variable flow rates

Figure 4.20 – Performance of optimized cooling plates at a range of flow rates and heat flux. The cooling plate performance is normalized against the reference cooling plate performance in Figure 4.19

However, in this case, the net effect is to increase the total heat transfer efficiency.

As the channel progresses from inlet to outlet the temperature of the coolant increases, and so the plate temperature shows a corresponding increase to maintain the flow of heat into the coolant. In the reference design, this is apparent by observing the general increase in plate temperature from top to bottom of the plate (Figure 4.8a). The T_σ design must counteract this effect to maintain a uniform temperature, and does so by reducing the width of the channel where the coolant is coldest (close to the inlet). Although the increased coolant velocity raises the local rate of heat transfer, the reduced channel surface area leads to reduced heat transfer in this region, and so equalizes the temperature over the plate.

The pressure optimized design is interesting in that the channel is not uniformly wide: the regions near the edge of the plate are increased at the expense of having a narrow orifice at the centre of the plate. This result is quite specific to the particular plate configuration and constraints, and would be difficult to apply to a plate with a significantly different channel configuration.

The designs optimized with respect to a single objective function clearly show the success of the optimization process. Not only are the performance criteria significantly improved over the reference design, but the converged geometries are in some cases quite specific and non-intuitive. For example, it is apparent from the initial design study (Figure 4.7b) that the T_σ optimized design requires the first two channel sections to be in a specific location to achieve a uniform temperature distribution; an optimization process is needed to locate them. Likewise, although adequate P_{fluid} and T_{avg} performance could be achieved with a design that simply enlarges the coolant area to a maximum, the particular channel widths that lead to the optimum

performance can only be determined by an optimization process.

The study of optimizations at different boundary conditions, and their assessment at off-design conditions reveals another feature of the cooling plate design. The optimizations of P_{fluid} and T_{avg} are generally insensitive to the magnitude of the boundary conditions (coolant flow rate and heat flux); i.e. a design optimized with one set of boundary conditions is also an optimum design at a different set of conditions. (Although note that this cannot necessarily be extrapolated beyond the range used in this study; the onset of unsteady flow and turbulence could alter the optimum designs.)

In contrast, the optimization of T_σ is quite sensitive to the coolant flow rate used in the optimization process; at a different flow rate, the plate operates non-optimally and the penalty increases as the flow rate departs from the optimization conditions. This study has not attempted optimization of an objective function composed of performance at multiple flow rates, but the assessment of these plates indicates that some plates are more 'specialist', and others are 'generalist'; the features of a 'generalist' plate could be selected by an appropriate optimization formulation.

4.4 Non-Uniform Heat Flux Profile

4.4.1 Results

Optimizations of both T_{avg} and T_σ were executed with each of the four non-uniform boundary conditions described in Section 3.5. Figure 4.21 shows the geometry of the four designs optimized for T_{avg}, and Figure 4.22 shows the designs optimized for T_σ.

The geometry of the T_σ optimized designs clearly shows the bias towards the heat

source: compare the orientation of the heat flux distribution in Figure 3.13 with the optimized designs in Figure 4.22. In all cases, the channels are skewed towards the side of the plate where the heat source is concentrated. The geometry of the T_{avg} designs shows no such regular feature: while all of them have almost the maximum area occupied by cooling channels, there does not appear to be a correlation between the layout of the channels and the location of the heat source.

The performance of the optimum designs when analyzed with their corresponding boundary condition is shown in Table 4.3. This shows that particular orientations of the heat flux boundary condition are advantageous for each objective function. To achieve the lowest T_σ value, the heat flux should have a positive y gradient. For the best T_{avg} performance, the opposite is required: a negative y gradient.

However, to demonstrate that it is necessary to perform the optimization with the heat flux distribution applied, a matrix of analyses was produced in which each plate geometry was analyzed with all possible boundary conditions. These are illustrated with charts in Figure 4.23. For every analysis boundary condition, the optimum performance is achieved with the plate optimized using the same boundary condition. In the case of T_σ optimized designs, there is a difference of up to almost six times between the performance of plates analyzed with their corresponding boundary condition, and plates analyzed with an opposing boundary condition. The T_{avg} designs have the same trend, but the variation in objective function is much lower.

4.4.2 Discussion

The results of the study presented here indicate that there are two aspects of cooling plate design that need to be considered when non-uniform heat generation is present.

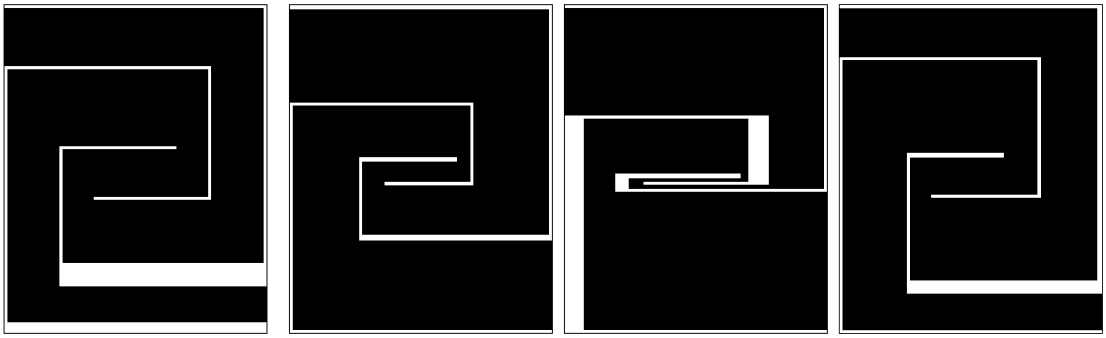

(a) +y **(b)** -y **(c)** +x **(d)** -x

Figure 4.21 – Cooling plate geometry of four designs optimized for T_{avg} with non-uniform heat flux distributions.

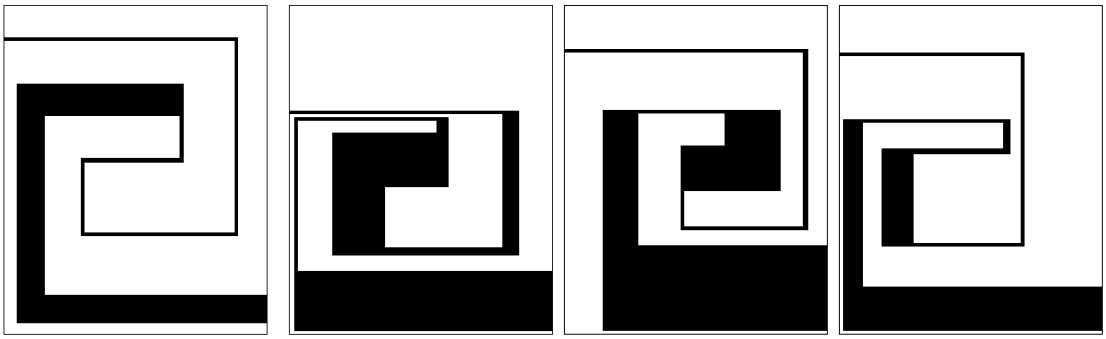

(a) +y **(b)** -y **(c)** +x **(d)** -x

Figure 4.22 – Cooling plate geometry of four designs optimized for T_{σ} with non-uniform heat flux distributions.

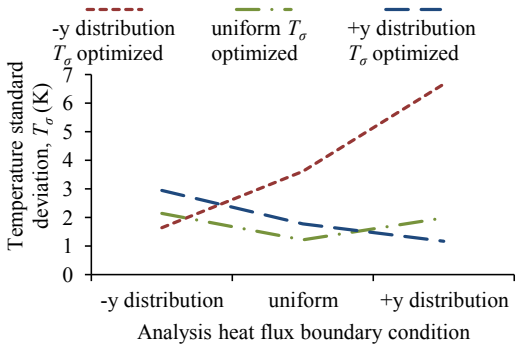

(a) T_σ performance of plates optimized with different heat flux gradients in the y direction

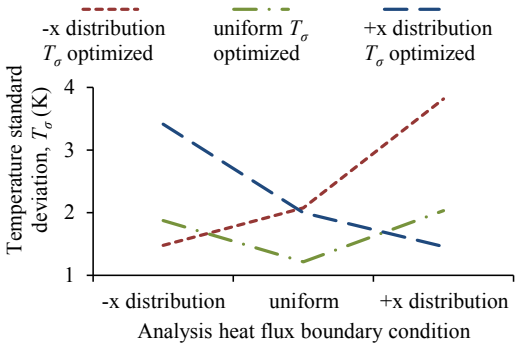

(b) T_σ performance of plates optimized with different heat flux gradients in the x direction

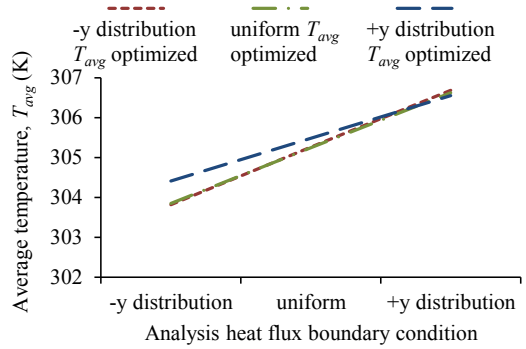

(c) T_{avg} performance of plates optimized with different heat flux gradients in the y direction

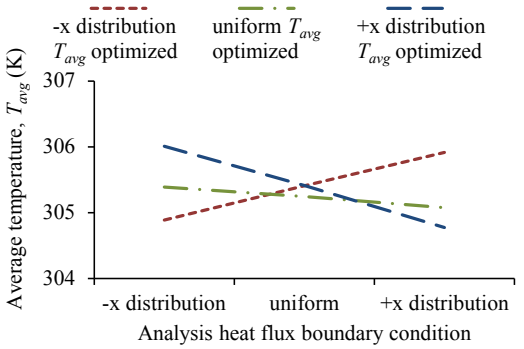

(d) T_{avg} performance of plates optimized with different heat flux gradients in the x direction

Figure 4.23 – Performance of cooling plates optimized with non-uniform heat flux boundary conditions. Each chart shows the performance of plates at three conditions: their optimization condition, and two others.

Table 4.3 – Performance of cooling plates optimized with non-uniform heat flux boundary conditions. Values show performance at the optimization conditions.

Objective	Flux Distribution	P_{fluid} (Pa)	T_{avg} (K)	T_σ (K)
T_σ	+y	25464	308.86	1.17
	-y	19098	307.15	1.64
	+x	20047	307.17	1.46
	-x	23457	308.13	1.48
	uniform	25124	308.31	1.22
T_{avg}	+y	1773	306.55	2.47
	-y	1526	303.82	2.87
	+x	7983	304.77	2.11
	-x	1795	304.89	3.22
	uniform	1232	305.25	2.32

The first is the orientation of the cooling channels with respect to the distribution of heat generation: when temperature uniformity is the objective, cooling plates give better performance when the inlet of the channel coincides with the region of greatest heat generation, as is the case for the positive y gradient. However, when average temperature is the objective, the reverse is true: the inlet of the channel should coincide with regions of lower heat generation.

The second aspect is the decision of whether to optimize the cooling plate with an appropriate heat flux boundary condition, or to use a uniform boundary condition. In the case of the T_σ optimization, there is a significant advantage to be gained in optimizing the plate with the appropriate boundary condition. For example, with a positive y gradient heat flux, there is approximately a 70% penalty in T_σ if the plate is optimized with a uniform heat flux rather than the y gradient flux. The other three heat flux boundary conditions have penalties of 30-40% when the plate used

was optimized with a uniform heat flux. However, it is worth considering that using a uniform heat flux for the optimization gives a consistently low value of T_σ no matter which analysis boundary condition is applied, making it an attractive choice if the distribution of heat flux is not constant.

In the case of the T_{avg} optimized designs, there is less of an advantage to be gained from optimizing the design with the appropriate boundary condition (although the optimization of T_{avg} gave less of an improvement than T_σ to begin with). For the positive and negative y gradients, there is only a slim benefit (approximately 1%) over the design optimized with a uniform heat flux. The heat flux gradients in the x direction exhibit a greater benefit from optimization with the appropriate boundary condition: up to 10% compared with the uniform design.

4.5 Multi-Objective Optimization

4.5.1 Results

The multi-objective optimization algorithm described in section 3.6 was implemented for the objective function pairs of $T_{avg} - T_\sigma$ and $T_\sigma - P_{fluid}$. Boundary conditions used were the same as the single objective optimizations: mass flow rate of 0.001 kg s^{-1} and a uniform heat flux of 500 W m^{-2}. As the single objective optimum designs of T_{avg} and P_{fluid} are so similar, the multi-objective designs would all be approximately the same, and therefore optimization of multi-objective designs between these was not attempted.

Initially four intermediate objective functions were used; at weighting factors of 0.2, 0.4, 0.6, and 0.8 (the single objective optimizations give the designs at weights

of zero and one). However, it became apparent that these points were not enough to give a comprehensive description of the Pareto front, and additional optimizations were used to fill in gaps as necessary. This was implemented in two ways: adding intermediate weighting factors, and introducing constraints on the objective functions. The latter is a technique developed by Kim and de Weck [58] termed the adaptive weighted sum (AWS) method.

Figure 4.24 shows the Pareto front for the $T_{avg} - T_{\sigma}$ optimization. Two charts are shown: the first contains only the designs identified by the initial set of weighting factors, and the second with the additional designs following the application of the AWS method. This chart also indicates the performance of the reference design, for comparison. This illustrates the limitation of the standard weighted sum method, discussed by Kim and de Weck, that the optimum designs tend to cluster towards either end of the Pareto front, and solutions in concave regions are not identified at all. In order to find the additional design points in the centre region of the Pareto front, two constraints were introduced to the objective function to reduce the feasible design space. The geometry of selected designs along the $T_{avg} - T_{\sigma}$ Pareto front is shown in Figure 4.25.

Similarly, Figure 4.26 shows the Pareto front for the $T_{\sigma} - P_{fluid}$ optimization. Chart 4.26a plots the six designs identified by the first pass of the standard weighted sum method: it has five points clustered at a low pressure, and only one point located at the other end of the Pareto front. An additional optimization with a weighting factor of 0.9 produced an intermediate design, and then the AWS method was used to add another two points; all of which are shown on chart 4.26b, along with the reference design performance. The geometry of selected designs along the $T_{\sigma} - P_{fluid}$

Pareto front is shown in Figure 4.27.

4.5.2 Discussion

Both the $T_{avg} - T_\sigma$ and $T_\sigma - P_{fluid}$ multi-objective optimizations revealed intermediate designs that may be more attractive to battery designers than the single objective optimum designs. In both cases there are several designs that outperform the reference design in both of the optimized parameters. From each single objective optimum design at the anchor points of the Pareto front, improvements can be made in the opposing objective function without compromising too much of the primary objective function performance. For example, looking at the $T_\sigma - P_{fluid}$ trade-off, by allowing a 17% increase in T_σ, the value of P_{fluid} can be reduced by approximately two thirds. At the other end of the Pareto front, a reduction of 20% in T_σ can be achieved with only a 525 Pa increase in P_{fluid}.

The geometry of the intermediate designs generally displays a gradual and consistent change from one single objective design to the other. However, there are some examples of adjacent designs that display a bigger difference than their performance would suggest: for example designs 4.27d and 4.27e. This is in the region of the Pareto front where the design takes on some characteristics from both single objective designs, and it is likely that many such intermediate designs exist with similar performance but different combinations of design features.

The location of an ideal design lies in the lower left corner of these Pareto charts; this is termed the utopia point. Although neither multi-objective optimization study has revealed a perfect design, both of them have shown some progress towards this part of the design space. This is particularly true of the $T_\sigma - P_{fluid}$ optimization,

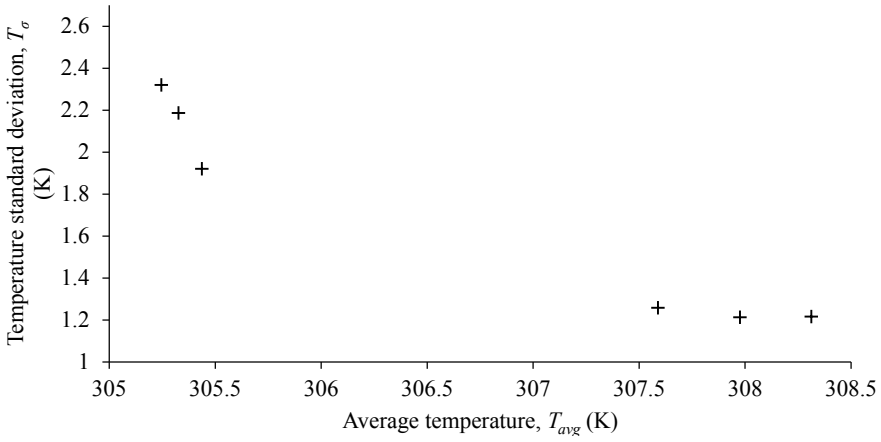

(a) Pareto front showing standard weighted sum designs

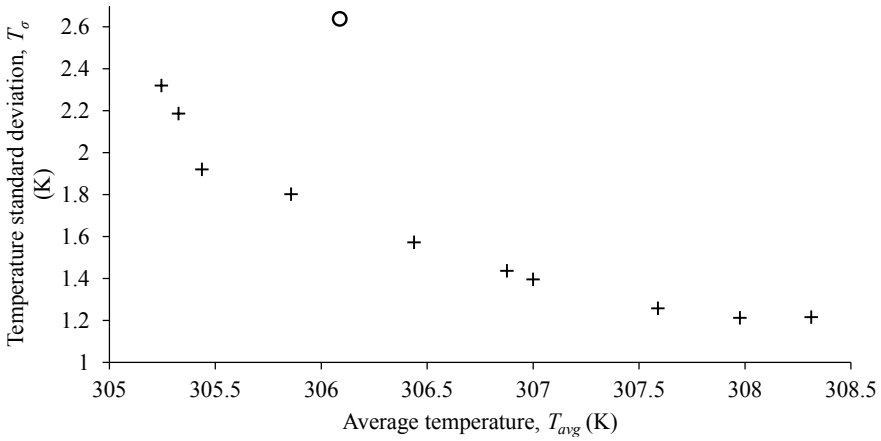

(b) Pareto front including additional designs determined with AWS. Reference design performance is indicated by the circle

Figure 4.24 – Pareto fronts determined as a result of multi-objective optimization of T_{avg} and T_σ.

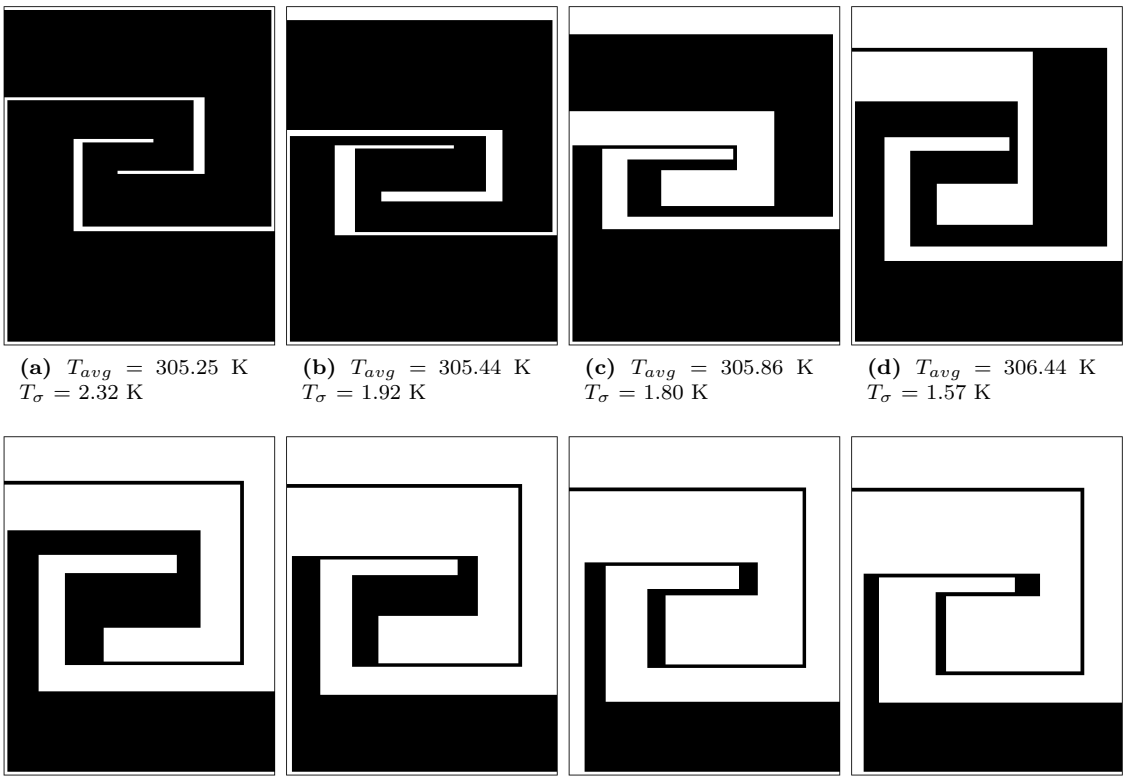

(a) T_{avg} = 305.25 K T_σ = 2.32 K
(b) T_{avg} = 305.44 K T_σ = 1.92 K
(c) T_{avg} = 305.86 K T_σ = 1.80 K
(d) T_{avg} = 306.44 K T_σ = 1.57 K
(e) T_{avg} = 307.00 K T_σ = 1.40 K
(f) T_{avg} = 307.59 K T_σ = 1.26 K
(g) T_{avg} = 307.98 K T_σ = 1.21 K
(h) T_{avg} = 308.31 K T_σ = 1.22 K

Figure 4.25 – Geometry of the multi-objective designs from T_{avg} to T_σ.

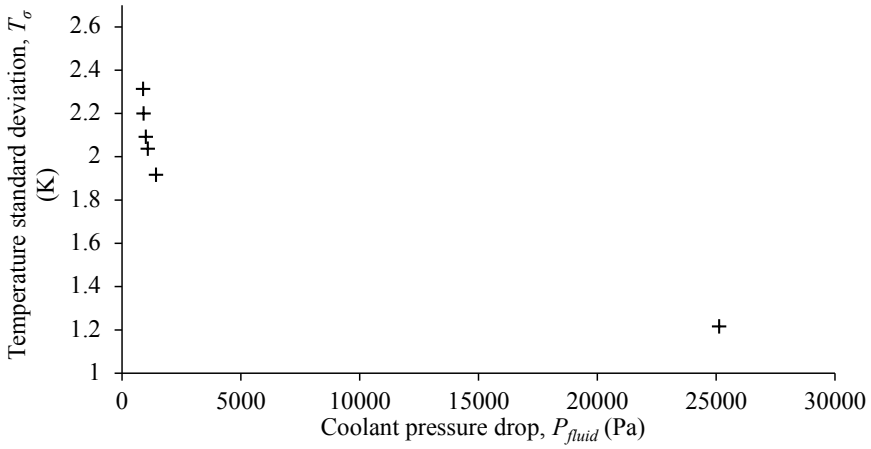

(a) Pareto front showing standard weighted sum designs

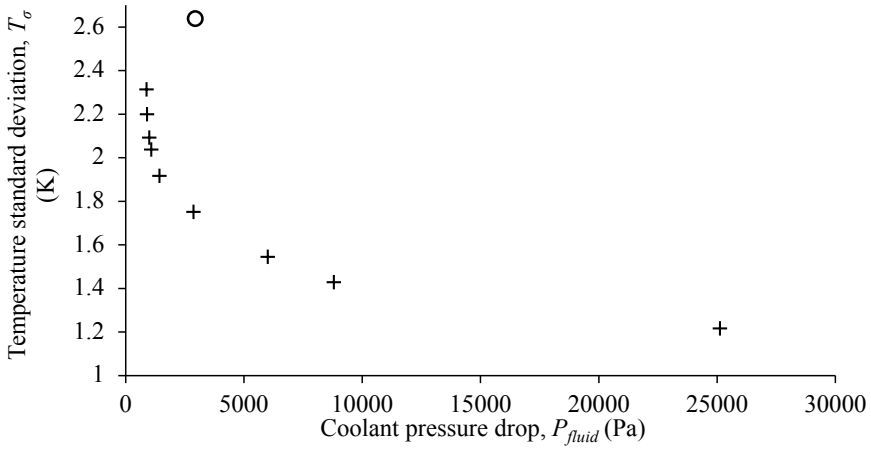

(b) Pareto front including additional designs determined with AWS. Reference design performance is indicated by the circle

Figure 4.26 – Pareto fronts determined as a result of multi-objective optimization of T_σ and P_{fluid}.

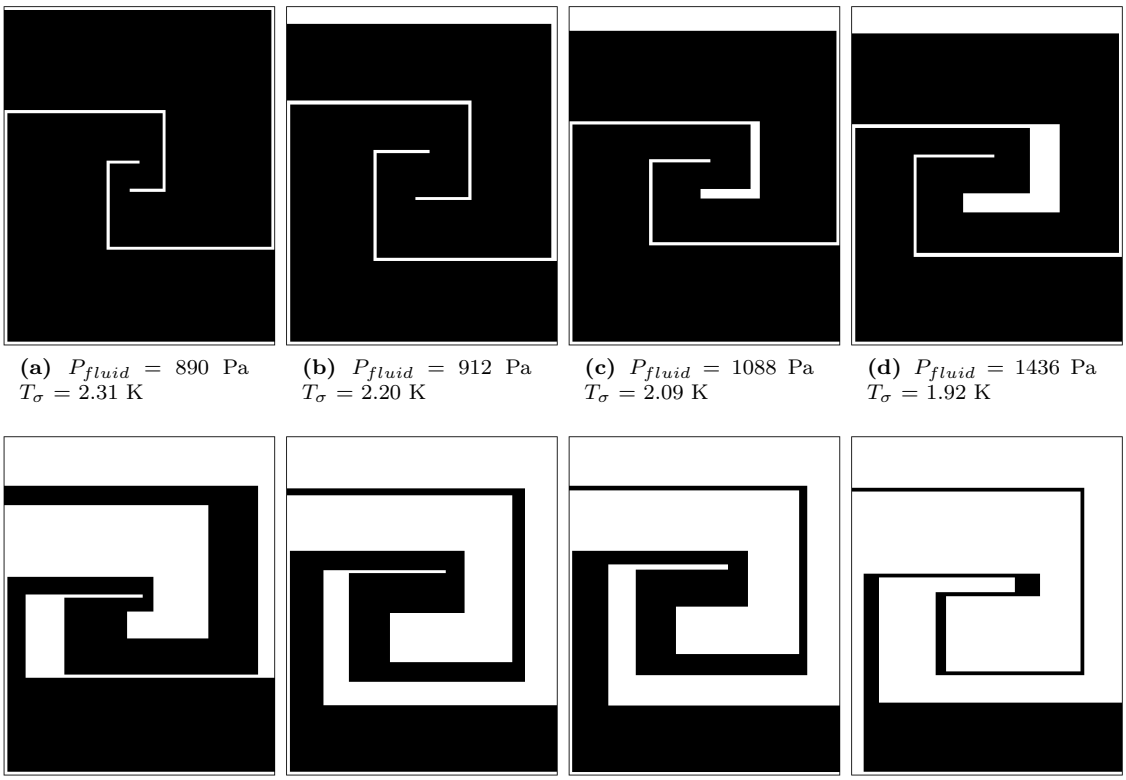

(a) P_{fluid} = 890 Pa T_σ = 2.31 K
(b) P_{fluid} = 912 Pa T_σ = 2.20 K
(c) P_{fluid} = 1088 Pa T_σ = 2.09 K
(d) P_{fluid} = 1436 Pa T_σ = 1.92 K
(e) P_{fluid} = 2877 Pa T_σ = 1.75 K
(f) P_{fluid} = 6020 Pa T_σ = 1.54 K
(g) P_{fluid} = 8814 Pa T_σ = 1.43 K
(h) P_{fluid} = 25124 Pa T_σ = 1.22 K

Figure 4.27 – Geometry of the multi-objective designs from P_{fluid} to T_σ.

where the P_{fluid} parameter can cover a wide range of values. Here the intermediate designs clearly fall on a concave curve with a trade-off much better than linear.

4.5.3 A Consideration of Temperature Uniformity and Power Consumption

Up to this point in the study, the discussion of power consumption has been relatively limited: through the optimization of pressure drop, the power consumption has implicitly been optimized for a particular plate (because the flow rate has been kept constant). However, if a comparison of two designs at different operating conditions is required, then it may be more useful to employ power consumption as a measure of performance, instead of either pressure drop or flow rate.

As an example comparison, consider the design situation in which the reference cooling plate is in use, but a more uniform temperature distribution is required. This is to be attained at the minimum possible increase in power consumption. The study of off-design boundary conditions indicated that one method of increasing the temperature uniformity of the reference cooling plate is to increase the coolant flow rate: this can be seen in Figure 4.19b. Alternatively, the multi-objective optimization showed that it is also possible to reduce either the temperature deviation or pressure drop by modifying the design—this maintains the coolant flow rate, but affects the power consumption via the pressure drop: this can be seen in Figure 4.26b.

Comparing the merit of these two approaches requires that the coolant flow rates and pressure drops be combined to produce a measure of the power consumed by pumping the fluid through the cooling channel. This is achieved simply by multiplying the volumetric flow rate by the drop in total pressure [52].

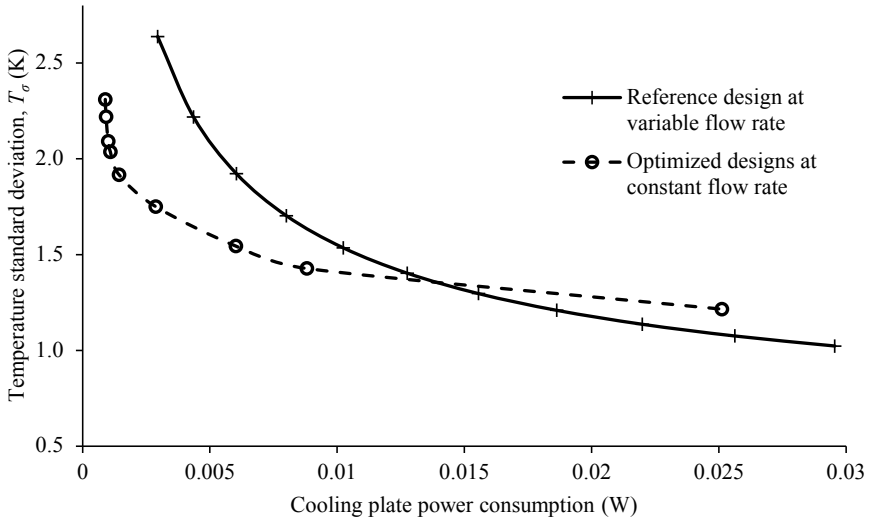

Figure 4.28 – Decrease in temperature deviation plotted against power consumption achieved by increasing the flow rate of the reference design (solid line), and multi-objective optimization of T_σ and P_{fluid} (dashed line).

Figure 4.28 plots the relationship between T_σ and power consumption corresponding to each of the approaches described. The top point on the solid curve represents the reference design at the standard set of operating conditions, and the following points indicate the power increase (by raising the flow rate) that is required to produce an improvement in temperature uniformity.

The alternative is presented by the dashed curve: these points indicate the improvement in temperature uniformity and corresponding power consumption achieved by design optimization—in this case the coolant flow rate is constant, only the pressure drop causes the power consumption to vary. It can be seen that at high values of T_σ, it is very effective to use an optimized design rather than just increasing the flow rate of the reference design: the power consumption can be reduced significantly while still improving the uniformity.

It appears that to further reduce the temperature deviation, it is preferable to use the reference design at a high flow rate; however, there is another option, which is to repeat the multi-objective optimization at a higher flow rate. This will translate the dashed line down and to the right, where there will certainly be an improved design point corresponding to an intermediate multi-objective design at a higher flow rate.

However, this raises another question: starting with the P_{fluid} optimized design, is it more effective (in terms of reducing T_σ at the smallest power increase) to move down the Pareto front towards the T_σ optimized design, or to simply increase the coolant flow rate? To address this question, the chart is redrawn so that instead of the reference design, the performance of the P_{fluid} optimized design is plotted at variable coolant flow rates. This is shown in Figure 4.29.

This chart shows that there are two or three multi-objective designs that offer marginally better performance than the single objective P_{fluid} optimized design, giving approximately a 10% improvement on either T_σ or power consumption. These are the designs shown in Figures 4.27b and 4.27c.

A more efficient method for identifying the optimum designs with respect to power consumption and T_σ would be to change the optimization formulation to add coolant flow rate as a design variable. It would then be possible to determine a true Pareto front in the power consumption–T_σ solution space. However, optimization parameters need to be chosen with care to prevent the flow rate from increasing or decreasing beyond a realistic range (as would happen in the unconstrained optimization of either power consumption or T_σ). This is outside the scope of the current study, but would be a valuable exercise to undertake in the future.

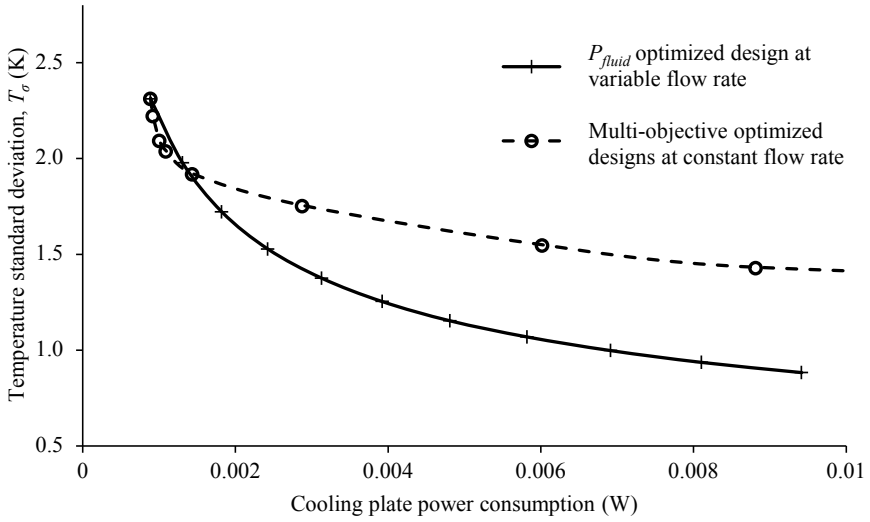

Figure 4.29 – Decrease in temperature deviation plotted against power consumption achieved by increasing the flow rate of the P_{fluid} optimized design (solid line), and multi-objective optimization of T_σ and P_{fluid} (dashed line).

Chapter 5

Summary and Conclusions

5.1 Summary

This study has investigated the performance of an EV battery cooling plate through the optimization of a parameterized plate with 18 design variables describing the location and width of the coolant channel. The research can be broadly divided into the construction, verification, and application of the optimization algorithm; the specific objectives identified in the Introduction are addressed below.

Develop a numerical model that can predict the thermal behaviour of a cooling plate

Utilizing the CFD preprocessor GAMBIT, a script was written that can generate a three dimensional cooling plate model. The script allows any orthogonal cooling plate design to be constructed with parameterized geometry specification; in this study, a single channel design with nine straight sections was used exclusively.

Thermal analysis of the cooling plate was carried out using the CFD solver FLUENT. Boundary conditions of inlet mass flow rate and pressure outlet were used to control the flow of the coolant, and a heat flux boundary was used to represent the heat generated by the battery. Thermal performance of the cooling plate was assessed by measuring the average temperature of the cooling plate, and by the standard deviation of the temperature. Cooling plate power consumption was determined by combination of pressure drop from inlet to outlet and the coolant flow rate.

The analysis model was verified by considering convergence of the output parameters as the mesh density is increased, and by monitoring the calculation residuals to determine the point at which the analysis should be terminated.

Develop a verified procedure for the optimization of an individual cooling plate

Having created the process for the assessment of a single cooling plate, the next step was to incorporate it into an optimization algorithm. Using the GAMBIT-FLUENT analysis as a single objective function evaluation, a process was written in the MATLAB programming environment using a built-in SQP function to handle the optimization, and additional functions to control the file management and user interaction. It allows the optimization of any single or multiple objective function composed of the three output parameters mentioned above; boundary conditions can be selected by the user; the model resolution and optimization precision is also controllable.

The optimization process was verified by conducting multiple optimizations from different initial designs: this was used to assess the dependence on the optimization

starting point and to determine the number of design iterations required to reliably reach a fully converged solution.

Demonstrate that single objective optimization can improve the design of a cooling plate

The single objective optimization solutions all showed a significant improvement over the reference cooling plate: between 14 and 70% depending on the objective function. In this study the objective functions of average temperature and pressure drop are closely linked; both optimizations resulted in a wide cooling channel: maximizing the channel surface area and flow area, but leading to a marked temperature gradient from inlet to outlet. The optimization of temperature uniformity resulted in a very narrow channel through most of the plate, only widening close to the outlet: equalizing the temperature over the plate, albeit at a higher average than the other designs.

Although performance improvements might have been attained without using numerical optimization—by applying engineering judgement, or by trial and error—the optimum designs exhibit specific features that would have been impossible to determine accurately with other design methods. Although the complexity of this design, with 18 variables, allowed for feasible optimization (in terms of computing resources and duration), any significant increase in the model size or number of design variables would directly impact the computing time and therefore the practicality of using numerical optimization of this type.

Perform multi-objective optimization of a cooling plate to determine trade-offs between single objective designs

The single objective optimizations established that the objectives of coolant pressure drop and average plate temperature were closely aligned; i.e. a single design could be derived that meets both objectives (within a small margin). These two objectives compete with the temperature uniformity objective: at the boundary conditions assessed, it was not possible for a design to be optimum for both temperature uniformity and average temperature (or pressure drop).

Using multi-objective optimization, a set of intermediate designs were determined that map out the transition from one single objective to the other. For both objective function pairs assessed, the performance of the intermediate designs was better than a linear trade-off between the single objective designs. In some cases it was possible to make significant improvements in one objective function with only a small penalty in the other objective function.

In practical vehicle design, where many performance constraints and criteria must be balanced, these cooling plates may be a more effective solution than the designs optimized with respect to a single objective function.

Determine the effect of changing boundary conditions: on plate performance, and the optimum design geometry

The investigation into the effect of boundary conditions was divided into two phases: first, simply varying the magnitude of the coolant flow rate and heat flux; second, by skewing the distribution of the heat flux to concentrate it at one side of the plate. In

both cases the performance impact of changing the boundary conditions was assessed, and also the optimization solution at different boundary conditions. This revealed the sensitivity of the optimum design to the boundary conditions.

As expected, varying the magnitude of the inlet coolant flow rate and the heat flux directly affected the cooling plate output parameters: increasing the flow rate reduced both the average temperature and temperature deviation, while increasing the heat flux increased these temperature parameters. When considering designs optimized at different boundary conditions, the optimum geometry was completely insensitive to the heat flux; i.e. the designs were almost identical, and the optimum design at one heat flux magnitude was the optimum design at all other magnitudes. The optimization solution was more sensitive to the coolant flow rate boundary condition: when optimizing with respect to average temperature there was an indication that the optimum design was different at different flow rates (although the effect was small); when optimizing with respect to temperature uniformity, it was clear that a design optimized at one flow rate was sub-optimal at other flow rates.

The study of non-uniform heat generation showed that the orientation of the heat flux had a significant effect on the plate performance. Although the serpentine route of the cooling channel precludes simple statements, it appeared that to promote uniform temperature, the coolant inlet should coincide with the region of greatest heat generation. To reduce average temperature, the reverse is true: the inlet should coincide with the lowest rate of heat generation. These relationships were only demonstrated for the channel geometry used in this study, and for linear heat flux gradients, but comparable effects could be expected other cases.

The optimization with respect to temperature uniformity was clearly affected by

the spatial distribution of heat flux: the cooling channel skewed towards the region of highest heat flux. Furthermore, a design optimized with one heat flux distribution had a significant performance penalty when a different distribution was applied. For the average temperature objective, the conclusion was less clear-cut, although there was still a small advantage to be gained from using the appropriate boundary condition for optimization. The design optimized with a uniform heat flux boundary condition performed moderately well at all applied boundary conditions, so this would make an acceptable assumption where the true boundary condition was unknown.

Determine general cooling plate design principles

The cooling plate design examined in this study was most effective when coolant channels were enlarged to their greatest extent, thereby increasing the area available for heat transfer. Although narrow channels increased the local heat transfer because of faster fluid flow, the overall effect was to reduce the heat transfer efficiency. However, it seems likely that a cooling plate with many narrow channels—thereby maintaining the high local rate of heat transfer, while also having a large total heat transfer area—would perform even better in terms of overall heat transfer efficiency.

Without optimization, there was a clear temperature gradient over the plate surface from inlet to outlet as the coolant absorbed heat and its temperature increased; i.e. the cooling effect due to temperature decreased along the length of the channel. To improve the temperature uniformity, this effect was offset by increasing the channel width along the length of the channel, thereby increasing the cooling effect due to heat transfer area. This is a difficult principle to generalize for any plate design, but balancing the coolant temperature rise against channel surface area should lead

Table 5.1 – Summary of situations in which the optimum design is dependent on the boundary conditions.

		Objective function		
		T_{avg}	T_σ	P_{fluid}
Boundary condition	Flow rate	maybe	Yes	No[1]
	Flux magnitude	No	No	No
	Flux distribution	Yes	Yes	No

[1]Previous studies indicate that at higher Reynolds number this will change

to improvements in the temperature uniformity.

Where the operating boundary conditions of the cooling plate are variable, it may be necessary to take account of this in the optimization process. However, this depends on the applied objective function and boundary condition in question—Table 5.1 summarizes the conditions for which the optimization is sensitive to the boundary condition. For example, for the T_σ objective function, both the coolant flow rate and heat flux distribution affect the optimum design, but the magnitude of the heat flux has no influence.

These principles were derived using a single cooling plate design over a constrained range of boundary conditions. It is reasonable to expect that they can be extrapolated to similar designs at comparable conditions, but beyond this—for example, turbulent fluid flow or multi-channel designs—there could be significant differences in behaviour, and verification would be needed to use them with confidence.

5.2 Limitations and Future Work

5.2.1 Design Space

The primary limitation of this study was the limited design space in which the optimization operated. The geometry generation algorithm with 18 design variables only allowed variation of width and location of the cooling channel sections: the topology always remained the same, and the construction lines were always orthogonal. The reason for these restrictions was to simplify the optimization problem, and to allow the repeatable generation of a consistent structured mesh. Without these bounds, the problem definition and solution could have become infeasible.

Although the topology of the cooling channel was selected following consideration of past studies, there is no reason to expect that it would lead to the absolute global optimum. There are several cooling plate features that have been proposed as the subject of future research including curved channels, multiple and branching channels, non-linear coolant routes, and variation in the through-thickness channel dimension. Some of these could be implemented using the same optimization routine with different parameters, but others would require a completely different modelling approach. In any case, a typical optimization study can only examine a small region of the total design space at a time, and so identification of a true global optimum can never be verified. However, by expanding the investigated design space, more design options are made available and an improved local optimum may found.

5.2.2 Boundary Conditions

Although the ultimate aim of this research is to improve battery performance, the battery cell was not represented in the CFD analysis. The temperature objective functions were measured on the face of the cooling plate in contact with the battery cell, but it cannot be assumed that this accurately reflects the temperature distribution within the battery cell. Thermal properties of the battery cell, and the contact behaviour could influence the temperature in the cell. Particularly in the case of non-uniform heat generation, it is probable that the battery would still maintain a non-uniform temperature distribution, even if the cooling plate interface had been equalized. To address this limitation, the FLUENT model could be increased to model the battery with internal heat generation, and conduction into the cooling plate.

Although this study investigated the optimum designs at a selected set of boundary conditions, they were steady-state single objective designs and not necessarily representative of the full operating conditions. During operation, a cooling plate will be subjected to transient boundary conditions determined by the vehicle driving profile and the external environment. This will dictate the magnitude and distribution of the battery heat generation, and assuming variable pump control, the coolant flow rate will be varied as necessary. A more thorough optimization would apply a representative driving profile and run a transient simulation of the cooling plate performance. However, a transient CFD analysis at each objective function evaluation would be prohibitively expensive, and the results in this study have indicated that of the optimization objectives, only T_σ is significantly sensitive to the boundary conditions. Therefore a suitable compromise for the optimization of T_σ might be to

run several steady state optimizations at different boundary conditions and weight them according to the corresponding driving profile.

5.2.3 Analysis Accuracy

Although there was an extensive verification effort on the CFD analysis, there were still some potential inaccuracies identified. Principal among these was the issue of unsteady or turbulent flow in the coolant. Although the fluid conditions were below the point at which previous studies had predicted the onset of unsteady flow, in some cases the channel geometry may have been severe enough to initiate unsteady flow before this point. This can be seen in the convergence behaviour of the CFD simulation: when sharp 180 degree corners were present, the analysis convergence was slower, and sometimes required a reduction of the under-relaxation factors. Without drastically increasing the CFD complexity, there is little that can be done to directly address this problem. Further investigation may determine settings that produce faster or more reliable convergence, but if the geometry and flow conditions are such that unsteady flow should be present, then the analysis will necessarily be only an approximation.

There are two approaches to the design of the cooling plate that might benefit the accuracy of the analysis: either reduce the Reynolds number and smooth the channels so that the flow is entirely laminar and steady, or purposefully operate in the fully turbulent regime. The first of these options requires slower flow and/or smaller channels: this would occur as a consequence of dividing the flow into multiple parallel channels. It will also be aided by having rounded corners and smooth transitions between conduits: as mentioned previously, this can be a challenge to implement

in a parameterized automated geometry generation process. The second option—to promote predictable turbulent flow—would require a significantly higher Reynolds number, or some geometric features to initiate it. The CFD analysis would be more time consuming, but with a suitable mesh and solver, it could be made reliable enough. By introducing turbulence, the heat transfer coefficient would be greatly increased, although at the cost of greater pressure drop. Some previous studies [59] have suggested that rather than a fully turbulent flow, the ideal heat exchanger design has a laminar main flow with features to introduce secondary vortices to promote mixing and so increase the rate of heat transfer. This is beyond the scope of the current study, but would make an interesting optimization subject for future work.

Another consequence of generating designs with a large variation in fluid and solid section size was that the mesh had sudden transitions in element size. This was a result of the geometry generation routine that can be seen in Figure 3.7: because all the geometry construction lines were orthogonal and continuous, a narrow channel section in one part of the plate could influence the mesh in another perpendicular channel section, as seen in Figure 5.1. This was another factor that led to slower solver convergence, and required lower under-relaxation factors to compensate. A more versatile mesh generation algorithm could overcome some of these difficulties and lead to a faster analysis process.

5.2.4 Optimization Settings

The execution of an optimization algorithm is often a balance between computing resources and accuracy, and this is true in this case. The act of repeating the single objective optimizations with eight initial designs led to a lower objective function

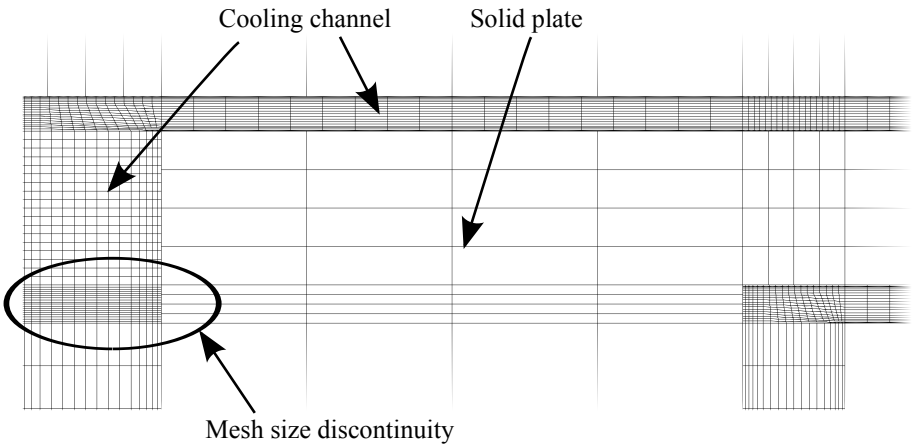

Figure 5.1 – Detail of the T_σ optimized mesh highlighting discontinuities in the element size.

and a design closer to the optimum. However, it was not feasible to employ this procedure for all the optimizations featured in this study: the computing resources available were already fully utilized for several months; to repeat every optimization eight times would be too time consuming. Likewise, the number of design iterations can be increased to arrive at a slightly better objective function, but the returns decrease significantly beyond 100 design iterations.

Due in part to the duration of a single optimization, little effort was directed to modifying settings to improve the performance of the optimization routine. There are several optimization algorithms available in MATLAB, and numerous settings that could be used to adjust the finite differencing, tolerances, and convergence settings of the optimizer. If future work is to involve many cooling plate optimizations then it would certainly be worthwhile to better understand the conditions under which the optimizer performs best.

As with all optimization problems of this nature, it is impossible to say for certain

that the converged design is the true global optimum. The optimization behaviour close to the converged design indicates that the design space is populated by many local minima with approximately equal performance. The repeated optimizations from multiple initial designs increased the chance that one found the true optimum, but without a survey of the whole design space it is impossible to be certain. The magnitude of such a search could be reduced by confining it only to the region in which the optimum is expected, and by pairing design variables where the channel width is expected to be either the maximum or minimum possible. However, a sufficient resolution would still require many thousands or millions of objective function evaluations which would be both inefficient and infeasible.

5.2.5 Validation

A limitation of any purely computational study is in the validation of results. Although every effort has been made to ensure the accuracy of the numerical aspects of the study, true validation can only be achieved by comparison with empirical results. In this study, there have been necessary simplifications of the geometry and boundary conditions, and assumptions regarding the behaviour of the coolant flow; only by performing an empirical assessment of the flow and heat transfer can we determine if these idealizations have been detrimental to the results.

5.3 Conclusions

In order to advance the technology of electric vehicles to the point where they can compete on performance with traditional internal combustion engine vehicles, the

principal area requiring attention is the battery. One aspect of battery design that has potential to bring about performance improvements is the efficiency and efficacy of the cooling system. The present study has addressed this issue through investigation of a specific battery cooling method: cooling plates.

Numerical optimization was used to demonstrate that the performance of a single cooling plate (as measured by temperature and pressure parameters) can be improved and tailored to specific operating conditions and performance objectives. In addition, simple design rules were observed that can be applied in the absence of a rigorous numerical optimization process.

This study was confined to cooling plates with a single channel in a 'double-S' configuration, and although these designs could be implemented with good results, it is likely that a more comprehensive survey of the design-space would reveal designs with superior performance. However, the protocols determined in this study are still of use: they can be applied to any design that can be parameterized and described by a set of design variables.

The ultimate goal of this research is to provide guidelines and design tools to vehicle engineers that enable them to focus the design of the cooling system to address the critical aspects of the battery performance. For example, an optimization study might result in a cooling plate with reduced power consumption. When incorporated into a battery pack, this would leave more energy available for the drive system of the vehicle; thereby increasing the range of the vehicle. Alternatively, optimization could lead to a cooling plate that prevents the formation of hot-spots in the battery cell. This would reduce the activity of corrosive reactions in the cell and so extend the useful lifetime of the battery.

Battery thermal management is one part of a very complex system. To realize its full potential, all components of the electric vehicle will need to operate at optimum performance. This research represents an important step on the path to achieving this goal.

Bibliography

[1] "World motor vehicle production by country and type; 2009-2010 (provisional)," tech. rep., International Organization of Motor Vehicle Manufacturers, 4 rue de Berri, Paris, France, 2010.

[2] "World Oil Outlook 2009," tech. rep., Organization of the Petroleum Exporting Countries, Obere Donaustrasse 93, A-1020 Vienna, Austria, 2009.

[3] International Energy Agency, *World Energy Outlook 2008*. OECD, Paris, 2008.

[4] "Short-term energy and summer fuels outlook; April 2011," tech. rep., U.S. Energy Information Administration, 1000 Independence Avenue, Washington, DC, U.S.A., 2011.

[5] D. Gately, "What oil export levels should we expect from OPEC?," *The Energy Journal*, vol. 28, no. 2, pp. 151–173, 2007.

[6] J. Smith, "World Oil: Market or Mayhem?," *The Journal of Economic Perspectives*, vol. 23, no. 3, pp. 145–164, 2009.

[7] *Deep Water: The gulf oil disaster and the future of offshore drilling. Report to the President*. National Commission on the BP Deepwater Horizon Oil Spill and Offshore Drilling, 2011. ISBN: 978-0-16-087371-3.

[8] R. Hirsch, "The inevitable peaking of world oil production," *Atlantic Council of the United States Bulletin*, vol. 16, no. 3, 2005.

[9] "2007: Technical Summary," in *Climate Change 2007: The Physical Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press, 2007.

[10] J. Kiehl and K. Trenberth, "Earth's annual global mean energy budget," *Bulletin of the American Meteorological Society*, vol. 78, no. 2, pp. 197–208, 1997.

[11] P. Tans, "Trends in atmospheric carbon dioxide." url: `http://www.esrl.noaa.gov/gmd/ccgg/trends/`. NOAA/ESRL. Retrieved June 8, 2010.

[12] J. Larminie and J. Lowry, *Electric vehicle technology explained*. Wiley, 2003.

[13] The Associated Press, "Prius no. 1 in japan sales as green interest grows." url: `http://www.usatoday.com/money/autos/2010-01-08-prius-tops_N.htm`. USA Today. Retrieved July 14, 2010.

[14] Nissan Motor Company, "Nissan leaf: Features and specifications." url: `http://www.nissanusa.com/ev/media/pdf/specs/FeaturesAndSpecs.pdf`. Retrieved April 6th 2011.

[15] General Motors, "2011 chevrolet volt." url: `http://www.chevrolet.com/volt/`. Retrieved April 6th 2011.

[16] "Commuting patterns and places of work of Canadians, 2006 census," tech. rep., Statistics Canada, Ottawa, Ontario, Canada, 2008.

[17] "Low carbon transport: a greener future," tech. rep., Department for Transport, Great Minster House, 76 Marsham Street, London, UK, 2009.

[18] M. Zolot, A. Pesaran, and M. Mihalic, "Thermal Evaluation of Toyota Prius Battery Pack," in *Proceedings of Future Car Congress, Arlington, Virginia*, 2002.

[19] "Voltec battery design and manufacturing," tech. rep., SAE International, Warrendale, PA, USA, 2011.

[20] J. Speight, *Handbook of petroleum analysis*. John Wiley & Sons, 2001.

[21] M. Kutz, *Mechanical Engineer's Handbook - Energy and Power*, p. 910. John Wiley & Sons, third ed., 2006.

[22] Nexergy, "Energy density comparison." url: `http://www.nexergy.com/battery-density.htm`. Retrieved July 16, 2010.

[23] H. Maleki, G. Deng, A. Anani, and J. Howard, "Thermal stability studies of li-ion cells and components," *Journal of the electrochemical society*, vol. 146, p. 3224, 1999.

[24] Y. Chen, L. Song, and J. Evans, "Modeling studies on battery thermal behaviour, thermal runaway, thermal management, and energy efficiency," in *Energy Conversion Engineering Conference, 1996. IECEC 96. Proceedings of the 31st Intersociety*, vol. 2, pp. 1465–1470, 1996.

[25] A. A. Pesaran, A. Vlahinos, and S. D. Burch, "Thermal performance of EV and HEV battery modules and packs," in *Proceedings of the 14th International Electric Vehicle Symposium, Orlando, Florida, December 15–17*, 1997.

[26] A. Pesaran, S. Burch, and M. Keyser, "An approach for designing thermal management systems for electric and hybrid vehicle battery packs," in *Proceedings of the 4th Vehicle Thermal Management Systems Conference and Exhibition*, pp. 24–27, 1999.

[27] A. Pesaran and M. Keyser, "Thermal characteristics of selected EV and HEV batteries," in *Applications and Advances, 2001. The Sixteenth Annual Battery Conference on*, pp. 219–225, IEEE, 2001.

[28] A. Pesaran, G. Kim, and M. Keyser, "Integration issues of cells into battery packs for plug-in and hybrid electric vehicles," 24th International Batterie, Hybrid and Fuel Cell Electric Vehicle Symposium (EVS24), Stavanger, 2009.

[29] J. Shim, R. Kostecki, T. Richardson, X. Song, and K. Striebel, "Electrochemical analysis for cycle performance and capacity fading of a lithium-ion battery cycled at elevated temperature," *Journal of power sources*, vol. 112, no. 1, pp. 222–230, 2002.

[30] A. Pesaran, A. Vlahinos, and T. Stuart, "Cooling and preheating of batteries in hybrid electric vehicles," in *6th ASME-JSME Thermal Engineering Joint Conference Proceedings, Hawaii Island, Hawaii*, 2003.

[31] G. Kim, A. Pesaran, and N. R. E. L. (US), *Battery Thermal Management System Design Modeling*. National Renewable Energy Laboratory, 2006.

[32] F. Chen, Z. Gao, R. Loutfy, and M. Hecht, "Analysis of optimal heat transfer in a PEM fuel cell cooling plate," *Fuel Cells*, vol. 3, no. 4, pp. 181–188, 2003.

[33] J. Choi, Y. Kim, Y. Lee, K. Lee, and Y. Kim, "Numerical analysis on the performance of cooling plates in a PEFC," *Journal of mechanical science and technology*, vol. 22, no. 7, pp. 1417–1425, 2008.

[34] S. Yu, S. Sohn, J. Nam, and C. Kim, "Numerical study to examine the performance of multi-pass serpentine flow-fields for cooling plates in polymer electrolyte membrane fuel cells," *Journal of Power Sources*, vol. 194, no. 2, pp. 697–703, 2009.

[35] S. Maharudrayya, S. Jayanti, and A. Deshpande, "Pressure losses in laminar flow through serpentine channels in fuel cell stacks," *Journal of Power Sources*, vol. 138, no. 1-2, pp. 1–13, 2004.

[36] D. Tuckerman and R. Pease, "High-performance heat sinking for VLSI," *Electron Device Letters, IEEE*, vol. 2, no. 5, pp. 126–129, 1981.

[37] S. Kim, "Methods for thermal optimization of microchannel heat sinks," *Heat transfer engineering*, vol. 25, no. 1, pp. 37–49, 2004.

[38] R. Knight, J. Goodling, and D. Hall, "Optimal thermal design of forced convection heat sinks—analytical," *Journal of Electronic Packaging*, vol. 113, pp. 313–321, 1991.

[39] R. Knight, D. Hall, J. Goodling, and R. Jaeger, "Heat sink optimization with application to microchannels," *Components, Hybrids, and Manufacturing Technology, IEEE Transactions on*, vol. 15, no. 5, pp. 832–842, 2002.

[40] H. Bau, "Optimization of conduits' shape in micro heat exchangers," *International journal of heat and mass transfer*, vol. 41, no. 18, pp. 2717–2723, 1998.

[41] X. Wei and Y. Joshi, "Optimization study of stacked micro-channel heat sinks for micro-electronic cooling," *Components and Packaging Technologies, IEEE Transactions on*, vol. 26, no. 1, pp. 55–61, 2003.

[42] J. Ryu, D. Choi, and S. Kim, "Numerical optimization of the thermal performance of a microchannel heat sink.," *International Journal of Heat and Mass Transfer*, vol. 45, no. 13, pp. 2823–2827, 2002.

[43] J. Kurnia, A. Sasmito, and A. Mujumdar, "Numerical investigation of laminar heat transfer performance of various cooling channel designs," *Applied Thermal Engineering*, pp. 1293–1304, 2011.

[44] T. Fisher and K. Torrance, "Constrained optimal duct shapes for conjugate laminar forced convection," *International Journal of Heat and Mass Transfer*, vol. 43, no. 1, pp. 113–126, 2000.

[45] T. Fisher and K. Torrance, "Optimal shapes of fully embedded channels for conjugate cooling," *Advanced Packaging, IEEE Transactions on*, vol. 24, no. 4, pp. 555–562, 2002.

[46] M. P. Bendsøe and O. Sigmund, *Topology optimization: theory, methods, and applications*. Springer Verlag, 2003.

[47] T. Borrvall and J. Petersson, "Topology optimization of fluids in Stokes flow," *International Journal for Numerical Methods in Fluids*, vol. 41, no. 1, pp. 77–107, 2003.

[48] L. Olesen, F. Okkels, and H. Bruus, "A high-level programming-language implementation of topology optimization applied to steady-state Navier-Stokes flow,"

International Journal for Numerical Methods in Engineering, vol. 65, no. 7, pp. 975–1001, 2006.

[49] E. Dede, "Multiphysics topology optimization of heat transfer and fluid flow systems," in *Proceedings of the COMSOL User Conference, Boston, MA*, 2009.

[50] J. Nam, K. Lee, S. Sohn, and C. Kim, "Multi-pass serpentine flow-fields to enhance under-rib convection in polymer electrolyte membrane fuel cells: Design and geometrical characterization," *Journal of Power Sources*, vol. 188, no. 1, pp. 14–23, 2009.

[51] "Modeling of battery pack thermal system for a pug-in hybrid electric vehicle," tech. rep., SAE International, Warrendale, PA, USA, 2011.

[52] A. Parr, *Hydraulics and pneumatics: a technician's and engineer's guide*. Elsevier, third ed., 2011.

[53] MEGlobal, "Ethylene glycol product guide." url: `http://www.meglobal.biz/literature/product_guides/MEGlobal_MEG.pdf`. Retrieved November 23, 2010.

[54] M. McKay, R. Beckman, and W. Conover, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, vol. 42, no. 1, pp. 55–61, 2000.

[55] U. Kim, C. Shin, and C. Kim, "Modeling for the scale-up of a lithium-ion polymer battery," *Journal of Power Sources*, vol. 189, no. 1, pp. 841–846, 2009.

[56] R. Hanks and H. Ruo, "Laminar-turbulent transition in ducts of rectangular cross section," *Industrial & Engineering Chemistry Fundamentals*, vol. 5, no. 4, pp. 558–561, 1966.

[57] J. Martin, P. Oshkai, and N. Djilali, "Flow structures in a U-shaped fuel cell flow channel: quantitative visualization using particle image velocimetry," *Journal of Fuel Cell Science and Technology*, vol. 2, p. 70, 2005.

[58] I. Kim and O. De Weck, "Adaptive weighted-sum method for bi-objective optimization: Pareto front generation," *Structural and Multidisciplinary Optimization*, vol. 29, no. 2, pp. 149–158, 2005.

[59] A. Jacobi and R. Shah, "Heat transfer surface enhancement through the use of longitudinal vortices: a review of recent progress," *Experimental Thermal and Fluid Science*, vol. 11, no. 3, pp. 295–309, 1995.