LUND UNIVERSITY

Direct air cooled hollow windings

Performance enhancement of electrically excited machines

Estenlund, Samuel

2024

Document Version:
Förlagets slutgiltiga version

[Link to publication](#)

Citation for published version (APA):
Estenlund, S. (2024). *Direct air cooled hollow windings: Performance enhancement of electrically excited machines* (1 uppl.). Division of Industrial Electrical Engineering and Automation, Faculty of Engineering, Lund University.

Total number of authors:
1

Direct air cooled hollow windings

Performance enhancement of electrically excited machines

SAMUEL ESTENLUND
FACULTY OF ENGINEERING | LUND UNIVERSITY

Direct air cooled hollow windings
Performance enhancement of electrically excited machines

Direct air cooled hollow windings

Performance enhancement of electrically excited machines

by Samuel Estenlund

LUND
UNIVERSITY

Organization	Document name
LUND UNIVERSITY	**DOCTORAL DISSERTATION**
Division of Industrial Electrical Engineering and Automation	Date of disputation
Box 118	2024-06-14
SE–221 00 LUND, Sweden	Sponsoring organization
Author(s)	
Samuel Estenlund	

Title and subtitle

Direct air cooled hollow windings: Performance enhancement of electrically excited machines

Abstract

This doctoral thesis presents the work the practical verification of a cooling concept involving direct air cooling of hollow windings for electric machines. The focus is on rotor windings for electrically excited synchronous machines for traction of electric vehicles. The work describes the cooling concept and compares it with the most common cooling concepts used today. The design of a experimental prototype includes the design of a modular dovetail joint rotor core to enable the installation of prewound coils.

The prototype is tested in order to verify the cooling concept. Experimentally a continuous rotor current of 175 A per turn is achieved with a winding hotspot temperature of 150 °C at an airflow of 1800 l/min. This corresponds to a copper current density of 21 A/mm^2. The thermal behavior is assessed with the help of a theoretical model which is validated with the experimental measurements. The model predicts that the prototype could handle a continuous current of 213 A per turn with a winding hotspot temperature of 155 °C at an airflow of 3200 l/min, which corresponds to a copper current density of 25 A/mm^2.

With this maximum rotor current the electromagnetic performance of the electrically excited machine is compared to a permanent magnet machine through finite element simulations. The results show that the electrically excited machine can produce a higher peak torque, but with a lower overall efficiency. However, the efficiency of some operating points above base speed, the electrically excited machine is higher, which means that it could be more efficient over drive-cycles operating much in those operating points, such as those representative of heavy-duty trucks. Among the conclusions it is stated that with sufficient cooling (such as that provided by the cooling concept in this work) the electrically excited machine can help relieve the electric vehicle market from its dependency of rare earth elements.

Key words		
direct cooling, electrical vehicle, electrically excited machine		

Classification system and/or index terms (if any)		

Supplementary bibliographical information		Language
		English

ISSN and key title		ISBN
		978-91-985109-9-7 (print)
		978-91-985109-8-0 (pdf)

Recipient's notes	Number of pages	Price
	196	
	Security classification	

I, the undersigned, being the copyright owner of the abstract of the above-mentioned dissertation, hereby grant to all reference sources the permission to publish and disseminate the abstract of the above-mentioned dissertation.

Signature _____   Date ___2024-04-30___

Direct air cooled hollow windings

Performance enhancement of electrically excited machines

by Samuel Estenlund

LUND UNIVERSITY

Cover illustration front: A close-up photograph of the prototype machine of this work, indexed to a Lund University color palette.

Who is wise and understanding among you?
Let them show it by their good life
by deeds done in the humility
that comes from wisdom.

James 3:13

Contents

List of publications

I **Performance estimation of hybrid excited machine with alternating pole configuration**
A. Reinap, M. Alaküla, S Hall, **S Estenlund**
2014 International Conference on Electrical Machines (ICEM), Berlin, Germany, 2014, pp. 2638-2643

II **PM-less machine topologies for EV traction: A literature review**
S. Estenlund, M. Alaküla, A. Reinap
2016 International Conference on Electrical Systems for Aircraft, Railway, Ship Propulsion and Road Vehicles & International Transportation Electrification Conference (ESARS-ITEC), Toulouse, France, 2016, pp. 1-6

III **Model of air cooled windings for traction machine**
S. Estenlund
2018 XIII International Conference on Electrical Machines (ICEM), Alexandroupoli, Greece, 2018, pp. 1278-1284

IV **Air cooling of an EMSM field winding**
S. Estenlund
TEIE-1087, Licentiate thesis, Lund University, 2018

V **Hollow direct air cooled windings: Rotor design process**
S. Estenlund
2020 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM), Sorrento, Italy, 2020, pp. 489-496

VI **Dovetail Design for Direct Cooled Rotor: Design and Manufacturing**
S. Estenlund, A. Tokat, J. Engqvist and M. Alaküla
International Conference on Electrical Machines (ICEM), Valencia, Spain, 2022, pp. 2121-2127

VII **Hollow direct air cooled rotor windings: Experimental verification**
S. Estenlund, A. Reinap
International Conference on Electrical Machines (ICEM), Torino, Italy, 2024 - *Submitted*

VIII **Hollow direct air cooled rotor windings: Cooling assessment**
S. Estenlund, A. Reinap
IEEE Transactions on Transportation Electrification - *Draft*

Preface

This dissertation is written as a monograph, which means it is written as a book covering all the work, as opposed to a compilation thesis which mainly consists of a collection of published papers. Parts of the work has been published, mainly as conference papers, but my aim has always been for this work to result in this monograph covering the entire design, manufacturing, testing and assessing of a cooling concept prototype, rather than producing individual papers.

The work has taken a long time to complete, long enough for me to see our four children being born, to buy and renovate two houses and to see my political party double its support in two consequential local election, all since the work started. These milestones along the time axis are of course also a contributing factor to the time is has taken to finish the thesis. (To the list of contributing factors one could also add the Covid-19 pandemic which interrupted some of the prototyping work)

All in all, during this time I have been given the chance to grow a lot - As a husband, father, house owner, public debater and politician, but also as a teacher, researcher and public outreacher. For this chance I am truly grateful, and also eager to continue the work to help others to grow through teaching and supervising students, to overcome technical obstacles from the energy transition we are in the middle of, and to make the fruits of this work reach the public.

Contribution statement

The entire work of the monograph presents the work and conclusions from the author alone whenever nothing else is clearly stated. From the publications listed on a previous page, the roles of the other authors have been as described as follows.

This Doctoral thesis is preceded by a Licentiate thesis (publication **iv**) which summarizes the firth half of the entire PhD (**i-iii**), while this dissertation focuses on the work after the Licentiate thesis. Publication **v**: S. Estenlund is the only author. Publication **vi**: A. Tokat, (co-supervisor at the time), review and editing. J. Engqvist, Solid Mechanics LTH, performed the tensile tests and shared the test data with S. Estenlund, review and editing. M. Alaküla, (main supervisor), review and editing. Publication **vii**: A. Reinap, (co-supervisor), experimental setup, review and editing. Publication **viii**: A. Reinap, (co-supervisor), fluid dynamics simulations and analysis (not part of this thesis), review and editing.

Acknowledgements

I personally wish to thank my supervisor, Mats Alaküla, for believing in me as a teacher, researcher and public outreacher. You have become a good friend and I treasure our conversations together about life, politics, society and our faith in the One with the Plans. I hope there will be many more of those.

Thank you to my co-supervisors; Avo Reinap for the everyday questions, chit-chats and laughs that you have provided, along with answers to my own questions; and Fran Marquez for the role-model you have been while you have grown from a doctoral student into a full-fletched senior and researcher during our time together. Also a very special thanks to Gustaf Olsson for your part in making sure this thesis could be finalized and defended.

There are many more who have played important roles in this work. Mikael Hörndahl at Industrial Production, Jonas Andersson at Arema-Mekano and Jonas Engqvist did so for the design, testing and manufacturing of the dovetail rotor prototype. Getachew Darge is another person who provided a crucial support in this work.

There are also many more colleagues at the division who all deserve much more than just to be mentioned like this: Fadi Abdallah, Philip Abrahamsson, Rasmus Andersson, Johan Björnstedt, Hannes Bydén, Hans Bängtsson, Max Collins, Leonardo Colombo, Gabriel Domingues, Bobbie Frank, Sebastian Hall, Morten Hemmingsson, Zhe Huang, Conny Högmark, Mattias Ingelström, Alice Jansson, Ulf Jeppsson, Anton Karlsson, Finn Landegren, Huan Li, Mats Lilja, Lars Lindgren, Gunnar Lindstedt, Carina Lindström, Yuri Loayza, Meng Lu, Martin Lundberg, Gabriel Malmer, Imran Maqbool, Hamoun Pourroshanfekr Arabani, Ramesh Saagi, Olof Samuelsson, Bengt Simonsson, Jörgen Svensson, Alexandra Tokat, Amir Torkiharchegani, Akanksha Upadhyay, Maira Viana Einarsson, Edvin Wallander, Henriette Weibull, David Wenander, Ulrika Westerdahl, and Per Widek. You have all in different ways been there with advice, jokes or just as someone to share a cup of coffee or an interesting lunch conversation with during workdays and it has all been valuable to me. Thank you!

Outside work I want to mention Dr. Per Arvidsson och Dr. Samuel Bengtsson specifically for the inspiration you have been as you have walked this path and further before me. To all of my friends at Kristna Föreningen Filippi and Laurentiistiftelsen, Kristdemokraterna, Väggarp and other places where I spent a lot of time during the last 10 years - If I were to start mentioning individual names from these places there would be no end to this preface. My parents, parents-in-law, siblings with families and siblings-in-law with families also have a big part in this work. Finally, the people who matter most to me and who have fought this battle night and day with me: Sara, I still do not understand what I have done to deserve to be your husband. We made it! And to my best teachers and biggest blessings: Noomi, Amos, Mika and Salome. Thank you. Without you none of this would matter.

Popular summary in English

The global market of electric vehicles is growing rapidly. This is an important part of the transition to a fossil free society. However, for this transition to be successful and have limited environmental effects it is important that sustainable materials are used in the electric vehicles. The permanent magnets that are used in over 80 % of today's electric vehicles contain rare earth elements which is one such critical raw material that may become short in supplies and cause a significant environmental impact. A solution to this is to use electric machines with copper-wound electromagnets instead, but these machines can seldom match the performance of electric machines with permanent magnets, since the electromagnets will more easily become overheated.

For this reason, the work of this thesis aims to verify a novel cooling concepts of the electromagnets that can improve their performance. This cooling concept is called direct air cooled hollow windings and means that the electromagnets are wound with hollow windings so that a coolant, in this case air, can be fed inside the windings allowing higher electric currents without overheating. With this cooling concept permanent magnet-less electric machines can match the performance of permanent magnet machines and for some types of vehicles even be more efficient.

In order to verify this cooling concept in practice a new type of rotor for electric machines has been developed, which allows the installation of the direct air cooled hollow windings. This new type of rotor includes dovetail joints, which means that the rotor can be taken apart and put together like a jigsaw puzzle. This joint must be able to handle the centrifugal forces of the spinning machine while also maximizing the performance of the electromagnets.

Once the prototype of the cooling concept is manufactured and assembled it is tested experimentally. The tests show that the concept does works as expected in practice. The simulated performance of the electromagnets of the prototype are compared to that of permanent magnet machines. The results show that the electrically excited synchronous machine can be more efficient for vehicles like heavy duty trucks and output a higher maximum torque. The hollow windings have been shown experimentally to continuously handle a copper current density of 21 A/mm^2 in the prototype, which according to theory can be increased to 25 A/mm^2 with a higher airflow.

This research can thus play an important role in the transport electrification transition by contributing to the use of materials with lower environmental impact and less risk of supply shortage. The direct air cooling concept can also potentially be used to enable the use of aluminium windings instead of copper windings, which would further improve these aspects. The use of a dovetail joint rotor may also allow for easier recycling and changing of materials in the electric machines during, or after, their life-cycle.

Chapter 1

Introduction

1.1 Background

The global road transport sector is taking quick steps from fossil fuels to electrification. In 2023 almost 14 million electric vehicles (EVs), including battery EVs (BEVs) and plug-in hybrid EVs (PHEVs), were sold globally compared to 2 million in 2019, now making up 18 % of every car sold globally [1]. In Norway, Iceland, Sweden and Finland, a majority of all personal vehicles sold in 2023 were electric (93 %, 71 %, 60 % and 54 % respectively). In Denmark, Belgium, China and the Netherlands EVs made up around a third or more of the vehicles sold, which may turn into a majority within this or the next few years.

This transition accelerates the demand for high power density and efficiency in the drive-trains. At the same time, the drive-train industry needs to find sustainable ways to supply the market, so the transition really becomes the positive evolution for the environment and climate that it is meant to be. The strive for power dense drive-trains has led to permanent magnet (PM) machines being represented in 82 % of all EVs in 2022 [2]. This is a problem for two reasons.

The first reason is the impacts that PM manufacturing and the extraction of the rare earth elements (REEs) needed for the PMs have on the environment and climate [3], especially since the global REE mining has increased by 150 % from 2018 to 2023, and by 280 % since 2003 [4] and will have to continue to grow as the EV market expands.

The second reason is that the global EV market is heavily dependent on one single country, which can lead to volatile prices and supplies. In 2023 69 % of the global REE mining took place in China [4]. This share has decreased from above 95 % in 2005-2011, due to a rapid increase of REE mining in the US, Myanmar and Australia. However, around 90 % of

the global refining of REEs into PMs still takes place in China [5], which means the REEs mined elsewhere in the world are still exported to China for refining.

This has led to the European Commission proposing an EU Critical Raw Material Act in March 2023 [6], in order to ensure the European supply of raw materials critical for energy transition and digital technologies. Of all the raw materials listed as critical, REEs have the highest supply risk [7, 8]. One way to decrease this supply risk for the EV industry is to make the EV market overall less dependent on these critical raw materials by turning to technologies that use other materials, such as PM-less electric machines.

This constitutes the motivation to look into electric traction machines for EVs without PM: The environmental impact of the manufacturing of PMs, the EV market's dependency on one single country, and the importance of ensuring sufficient raw materials for the energy transition.

1.2   Licentiate thesis summary

This section presents a summary of the licentiate thesis [9] preceding the work presented in this doctorate thesis.

1.2.1   PM-less machines for EV traction

The scope of the licentiate thesis was to study the options to PM machines for EV traction. In 2017, 17 of the 20 most sold EVs (BEV and PHEV) in the world used PM machines, despite the permanent magnets being expensive and the PM market proving to be uncertain due to 97 % of all neodymium for permanent magnets coming from China at the time. On top of that, extraction of neodymium is very unfriendly to the local environment.

A literature study shows there are several options to PM-less electric machines for EV traction, but not all are as promising. Switched reluctance machines and synchronous reluctance machines are low-cost and robust machines, but there are difficulties to obtain the power and torque density required by the EV market. Induction machines are a more promising option, already present in the 2012 EV market in Tesla Model S. The induction machine is robust too and it can reach comparable torque densities but with lower efficiency than those of the PM machine. Control is more challenging compared to synchronous machines, particularly as performance and efficiency at higher speeds should not be sacrificed. The electrically excited machine too has the ability to reach torque density comparable to that of PM machines, but with lower efficiency. However, if the rotor losses can be managed, it has several advantages over PM machines. An example of the most promising machine topologies for EV traction are shown in Figure 1.1.

(a) IPMSM   **(b)** EMSM   **(c)** IM
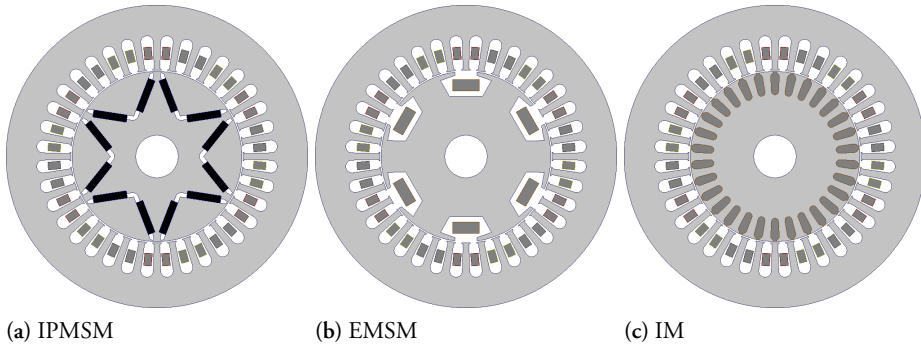
Figure 1.1: The machine topologies found most suitable for EV traction in the literature study of the licentiate thesis: The interior permanent magnet synchronous machine (IPMSM), the electrically magnetized (or in this thesis, excited) synchronous machine (EMSM or EESM) and the induction machine (IM).

The licentiate thesis thus concludes that the PM-less topology to study is the electrically excited machine, with a special focus on thermal management of the rotor losses.

1.2.2 Direct cooling concept

Direct cooling is here interpreted as cooling applied directly where the losses appear in an electric machine, and thus also directly where the heat is generated. At the division of Industrial electrical engineering and automation (IEA), work with direct cooling started already with the concept of laminated windings in 2012 [10, 11, 12]. The work concluded that laminated windings are really promising in theory, but difficult to implement due to the laminated windings getting dented much easier in practice than in CAD software (Figure 1.2). This results in thermal hotspots where the laminates are too close together to allow a sufficiently high air flow through.

Figure 1.2: Example of rolled laminated winding with uneven spaces for air flow.

With this experience at hand, another design iteration was needed for direct cooled machine windings. If the conductors can be shaped as dog-bones, or hollow, the coolant can flow

along the windings in robust and consistent channels with consistent cooling. This concept was combined with the need for thermal management of rotor losses in electrically excited machines for the work of the rest of the licentiate thesis.

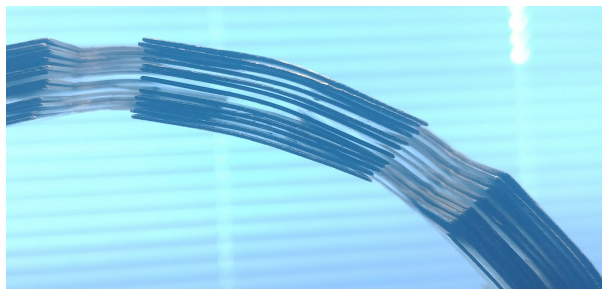**Figure 1.3:** Left: Cross section of two rectangular winding with concave indents along the sides. Right: Cross section of two rectangular hollow windings with inner circular coolant channels.

1.2.3 Direct cooling model

In order to verify the feasibility of this new design iteration of direct cooled electrically excited synchronous machine (EESM) rotor windings, a model was developed. The model is a combined empirical-lumped parameter model, iterating back and forth between two models (as shown in Figure 1.6). The empirical model calculates the heat transfer coefficient from the copper to the coolant air in the cooling channels, along with the bulk air temperature along the channels (as shown in Figure 1.4). The lumped parameter thermal network calculates the copper temperatures as well as the temperature dependent resistive copper losses (as shown in Figure 1.5). The model assumes that the windings are thermally isolated from the surrounding, so that all the heat from the copper losses are carried away only by the coolant air.
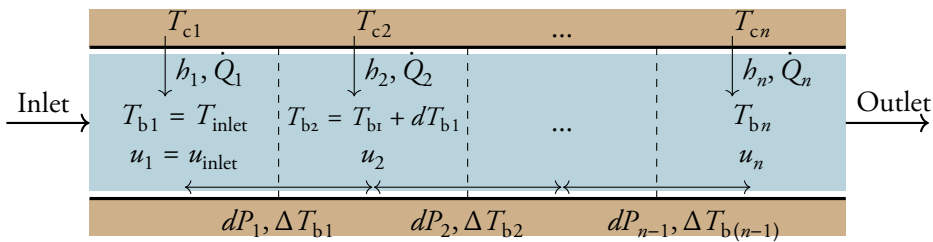
Figure 1.4: Sketch of the sectioned pipe and how some of the quantities develop along the direct cooled winding in the empirical forced convection model.

The cooling model is experimentally verified, with a hollow conductor with outer dimensions of 6x6 mm, an inner radius of 2 mm and a length of 870 mm, and then used to compare different winding dimensions. The model shows that it is theoretically better with a fill factor of about 70 % in order to reach the maximum current density, using as
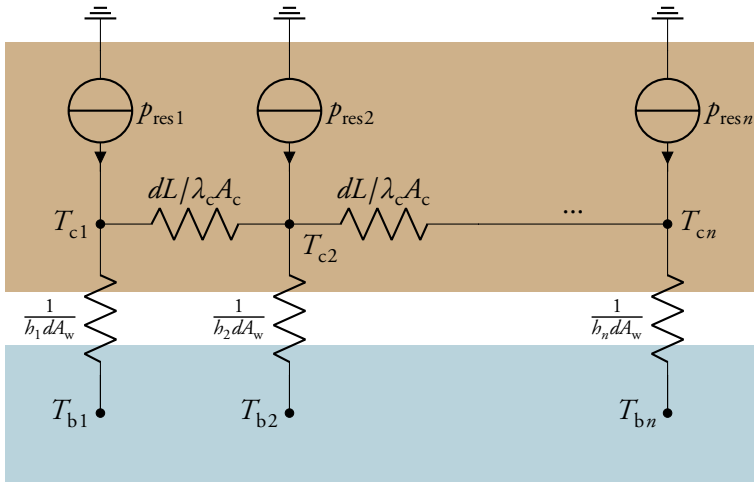
Figure 1.5: Lumped parameter thermal network of the direct cooled winding.

thin conductors as possible. In other words, really thin conductors with consistent cooling channels making out around 30 % of the cross section area, will allow better cooling and thus higher current densities without overheating the windings. However, in order to avoid the mistakes of the preceding work with the laminated windings, the smallest dimensions are avoided. What works good in theory will meet unforeseen challenges in practice. Thus, aiming for too thin conductors with too narrow cooling channels may prove not to be a robust design in practice since real conductors are not necessarily perfectly straight and debris or dust may block narrow cooling channels.

Results from the cooling model are presented in Chapter 2.2.1 and 2.4, where it is used to predict the performance of the cooling concept as well as motivating the coolant choice.

1.2.4 Direct cooled rotor design and results

In order to study the implications on the performance in using direct cooled windings in an electrically excited rotor, a reference machine is chosen (see Table 1.1). The machine is an interior permanent magnet (IPM) machine designed by Dr Rasmus Andersson and used by Volvo [13, 14]. The purpose has not been to produce a full prototype, but to design an EESM rotor for the stator of the IPM machine and compare the simulated performance with that of the IPM rotor.

Three rotor design geometries are studied in order to find an optimal electrically excited rotor. Two of the rotor design geometries are shown in Figure 1.7: One with a constant rotor tooth width, one with a constant winding slot width, and a final option where the slot is at an angle in between the other two geometries. These three options are studied

$$u_{\text{inlet}}, dP_{\text{linear}}$$
$$T_{\text{b,inlet}}, T_{\text{c,constant}}$$

$T_{\text{c,constant}}$

$T_{\text{b}1..n}$

| 1st empirical iteration |

$dP_{1..(n-1)}$

$h_{1..n}, T_{\text{b}1..n}$

| 1st lumped parameter it. |

$R_{1..n}$

$T_{\text{c}1..n}$

$T_{\text{c}1..n}$

$T_{\text{b}1..n}$

| 2nd empirical iteration |

$dP_{1..(n-1)}$

$h_{1..n}, T_{\text{b}1..n}$

| 2nd lumped parameter it. |

$T_{\text{c}1..n}$

$T_{\text{c}1..n}$

$T_{\text{b}1..n}$

| 3rd empirical iteration |

...
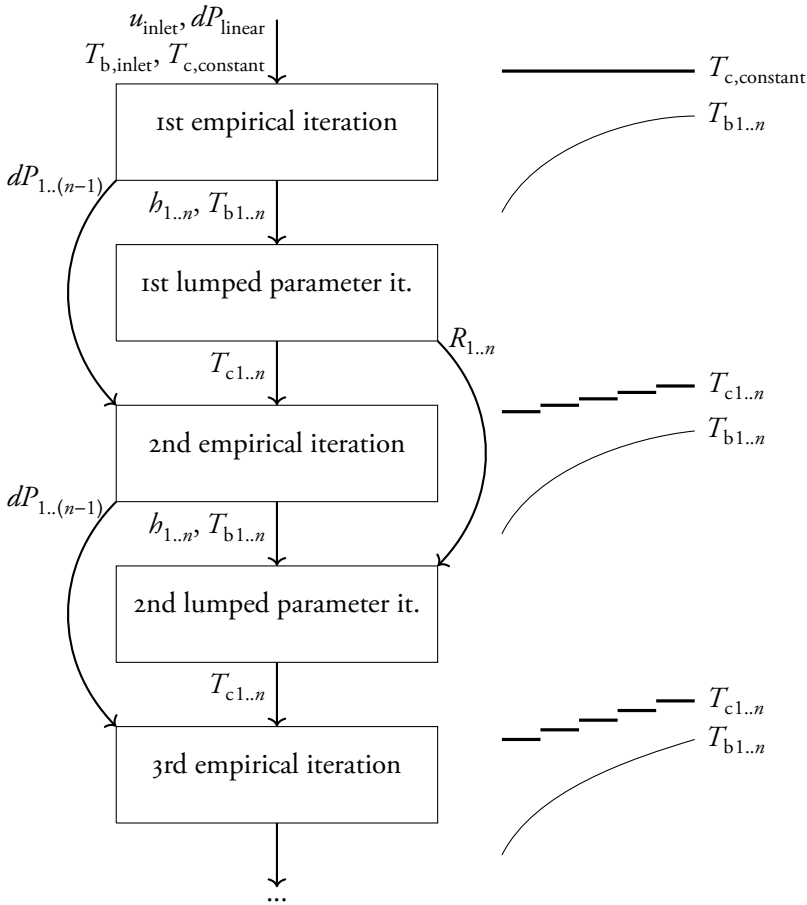
Figure 1.6: Flow chart of the combined empirical-lumped parameter model process.

with a sweep of various slot widths and heights, in an optimization process where finite element model (FEM) simulation results are combined with data from the cooling model to calculate the full losses in each operating point (copper losses, core losses and air pump losses for the coolant air). For each slot dimension of each rotor geometry a full $T\omega$-map with efficiency is produced. The optimal design was found to be one of option 2, with constant slot width, since it maximizes the space for copper and iron in the rotor without any dead space in the angle between the coils of two poles.

The results of the final design are compared with the original IPM machine. The comparison shows that the EESM rotor can make the machine provide a higher torque, higher power and higher speeds, but with lower efficiency than the IPM rotor in almost all operating points.

Table 1.1: Reference machine specifications

Overall	Peak continuous power	110	kW
	Peak maximum power	180	kW
	Supply voltage	650	V
	Maximum speed	15 000	rpm
	Active length	224	mm
	Airgap length	1	mm
Stator	Outer radius	100	mm
	Inner radius	60	mm
	Slots	36	mm
	Cooling method	End turn oil spray	
	Coolant	Gearbox oil	
Rotor	Outer radius	59	mm
	Inner radius	31	mm
	Poles	6	mm
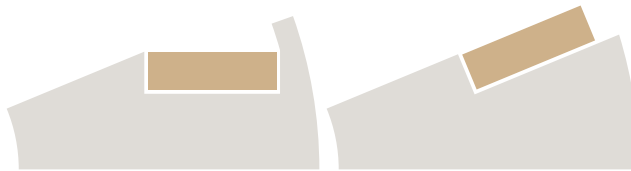
Figure 1.7: Left: Option 1 with constant tooth width. Right: Option 2 with constant slot width

1.2.5 Licentiate thesis conclusions

In the licentiate thesis it is concluded that it may be possible to make a direct cooled EESM that outperforms PM machines when it comes to peak torque, output power, and in some operating points efficiency. The direct cooling concept in the final rotor design allows a slot current density of 14 A/mm^2. The licentiate thesis concludes that the work shows good signs for the forthcoming work with the direct cooling concept, but that there are also some important knowledge gaps to be filled.

The rotor design did not take manufacturability into account, nor did it consider the thermal management of the whole machine, but only the direct cooled rotor windings. Some practical hurdles remained unconsidered such as how the current supply to the rotor windings should be provided, how the coolant air is produced and fed to the windings, how is a rotor designed so that the hollow windings can be installed, how these practical solutions affect the electromagnetic and mechanic aspects. The conclusions from the Licentiate thesis work [9] clearly verify that the objectives of the rest of the work to the doctorate thesis must

be finding out whether direct air cooled windings for an EESM rotor works in practice.

1.3   Objectives

The objective of this doctoral thesis is to verify a concept of hollow direct air cooled windings for an electrically excited rotor for EV traction in practice. This is done with practical and industrial perspective throughout the work so that a suitable manufacturing process is developed rather than striving to complete a prototype with maximized performance at any cost. The experimental verification aims to verify the thermal aspects of the cooling concepts, which requires controllable and measurable air flows and excitation currents. This means that the design models can be verified by calibrating the models with the experimental results.

1.4   Thesis outlines and limitations

The thesis follows the path from the idea of the concept in theory to the realization and validation of the concept in practice. *The first chapter* explains that the purpose of this thesis is to verify a concept of hollow direct air cooled windings for an electrically excited rotor for EV traction in practice.

The second chapter presents the target prototype of the thesis, the machine for which a new electrically excited direct cooled rotor is to be designed, built and tested as described in later chapters, along with the expected performance of the finalized prototype compared to an IPM rotor in the same stator. The prediction is done with a thermal model which is verified and calibrated with the experimental data at the end of the work. The chapter goes on to explain the reasons to study direct cooling, hollow conductors, and air as coolant. This constitutes the wider background of the work.

The third chapter explains the design of the prototype from a practical angle. It includes an analysis of how the cooling requirement will put constraints on the rotor design, including practical feasibility as well as electromagnetic performance. The conclusion is that a modular rotor topology is needed, introducing a mechanical challenge which is treated in *the fourth chapter*. *The fifth chapter* involves some unique manufacturing challenges and techniques for manufacturing the modular rotor and for the prototype as a whole, and finally assembles the whole prototype for testing.

The sixth chapter presents the experimental tests of the prototype, the equipment and methods used and the resulting data. The theoretical thermal model of the prototype is verified with the experimental results and then used to better assess and understand the function of

the cooling concept. From that the conclusions are drawn along with reflections on future work in *the seventh chapter*.

1.5   Contributions

Chapter 2 presents one comparison between different cooling methods, one between different coolants and one between the performance and efficiency of PM and electrically excited (EE) machines.

Chapter 3 presents a winding method suitable for hollow direct cooled windings.

Chapter 4 presents a novel topology for electrically excited rotors, namely the dovetail joint rotor, which enables prewound coils to be installed on the rotor.

Chapter 5 presents the implications of using wire discharge machining on laminated steel.

Chapter 6 presents an accurate theoretical thermal model of a direct air cooled machine and suggests several concepts for even more effective direct cooling of electric machine windings.

Chapter 2

Background of the direct cooling concept

Figure 2.1: Previous page: Photo of hollow windings with open end turns for parallel airflow.

2.1   Chapter overview

In Chapter 1 it is stated that the goal of this thesis is to verify whether a certain concept of hollow direct air cooled windings for the rotor windings of an electrically excited synchronous machine works in practice. In this chapter the motivation of the work is presented more in depth:

- The target machine, used as prototype for the experimental verification, is presented. The performance level expected from the prototype with the hollow direct air cooled rotor windings and how that compares to an IPM rotor.

- Common cooling concepts used in electric traction machines today are compared in a literature study and a numerical study to point out the strengths and weaknesses in order to stress the need to study and verify direct air cooled hollow windings.

- The coolants that can be used for direct cooling in hollow windings are compared in order to show why air is used as coolant in this work.

2.2   Prototype machine presentation
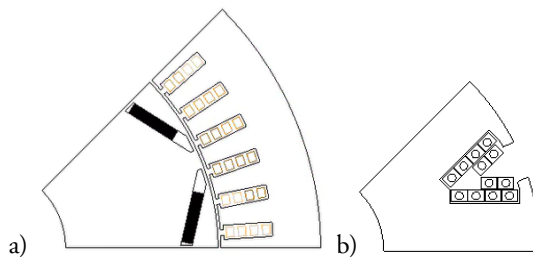
Figure 2.2: a) Equivalent design of the original machine with an IPM rotor. b) Initial design of the prototype direct cooled EE rotor.

The aim of the work is to produce a prototype and validate the concept of hollow air cooled windings for an EE rotor. This is achieved by designing and manufacturing a rotor for an existing stator. The stator in question is provided by Chalmers University of Technology and is originally designed for an IPM rotor. The design of the stator and IPM rotor are

protected by immaterial rights, but an equivalent stator and IPM rotor are shown in Figure 2.2. The equivalent IPM rotor is designed to match the original rotor with a no-load airgap flux density of 0.6 T (calculated as the peak value of the fundamental). The machine specifications, both the general and those specific to the equivalent IPM rotor and the prototype EE rotor respectively, are listed in Table 2.1.

Table 2.1: Target machine specs

Overall	Output power	56.7	kW
	Supply voltage	360	V
	Maximum speed	10 000	rpm
	Max torque	120	Nm
	Active length	120	mm
	Airgap length	0.75	mm
Stator	Outer radius	87.5	mm
	Inner radius	61.5	mm
	Slots	48	-
	Winding type	Hairpin	
	Turns	2	-
	Max current	225	A
	Cooling method	Housing water jacket	
Rotor	Outer radius	60.75	mm
	Inner radius	21	mm
	Poles	8	-
Equivalent IPM rotor	Airgap flux density	0.6	T
	PM dimensions	3x13.4 mm	
Prototype EE rotor	Winding type	4x3 mm hollow	
	Turns	6	pp
	Max cont. current	195	A
	Cooling method	Direct air cooling	
	Cooling channel	1 mm radius	

2.2.1 Initial EE rotor design

An initial design of the prototype EE rotor is made in order to enable an early performance comparison with the equivalent IPM rotor. This design is not determined from any analyses such as electromagnetic performance, manufacturability and structural integrity, which means the design is expected to be changed in the upcoming design process.

The initial design is based on a few constraints. The stator yoke is 12 mm thick, which

determines that the rotor teeth should be 24 mm wide and the rotor yoke at least 12 mm thick, so that the magnetic flux saturation begins simultaneously in both the rotor and the stator. In addition the tooth tips thickness is set to be 3 mm to ensure some for mechanical integrity of the rotor. The rotor windings are chosen to be 4x3 mm hollow windings from Luvata [15] with a circular hollow channel having a 1 mm radius. The initial design resulting from these constraints is shown in Figure 2.3, with the constraints summarized to the left and the design with the best possible winding configuration shown to the right.
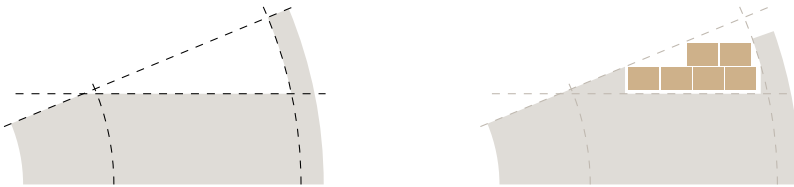

Figure 2.3: Left: Boundaries for slot area marked with dashed lines, and the resulting available slot area as white. Right: Possible winding configuration with 6 turns of 4x3 mm windings within the boundaries.

In order to determine an initial thermal limitation of the rotor windings, the direct cooling model of the licentiate thesis is used [9, 16]. The constraints set for the cooling include allowing a maximum conductor temperature of 160 °C and an inlet coolant air temperature of 60 °C. The resulting maximum continuous rotor current is 195 A per turn, corresponding to a copper current density of 23.5 A/mm^2. The temperatures of the windings and the coolant along the active length of the machine at 195 A rotor current is shown in Figure 2.4. The results are based on a number of assumptions: all 6 turns are cooled equally, no thermal exchange happens between the windings and the rest of the rotor and any rotor core losses present can be handled by the stator housing water jacket cooling. This constitutes a conservative case, since in a real rotor some heat will escape from the rotor windings through the shaft and via the airgap and not only through the coolant.

The winding temperature limit of 160 °C is set in order to avoid overheating the varnish of the conductors. The temperature index of windings refers to the maximum temperature the varnish take for 20 000 hours of operation [17]. The temperature index depends on the classification acquired and is today often in the range of 180-240 °C [18]. This leaves a margin to the limit of 160 °C used in this work, for three reasons. First, a margin to the temperature index prolongs the lifetime of the windings [19, 20]. Second, a safety margin for unexpected temperature hotspots in the windings. Third, it means that the results of this work is not only applicable for the highest temperature classifications of windings, but also for lower classifications. In the end, the application of a certain machine will set the prerequisites for the maximum temperature used, which includes a trade-off between lifetime and performance. The limit used in this work makes is not a universal limit, but it can be applicable for a large variety of applications for the mentioned reasons.
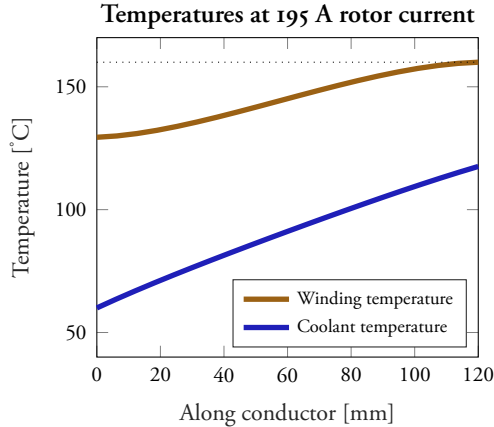
Figure 2.4: Temperature distributions along a winding of length a 120 mm winding with a current of 195 A according to licentiate thesis model. Dashed line marks the maximum allowed winding temperature of 160 °C.

2.2.2 IPM and EE rotor comparison

A simulation study was performed in order to compare the reference machine with an IPM rotor with 0.6 T airgap flux density and an EE rotor with the basic design described in the previous section. The simulations were run in ANSYS Electronics 2021 R1. The core mesh near the airgap and other important areas had a length of 1.5 mm, while most of the cores had a mesh length of 5 mm (which is tested to be as accurate as any finer mesh but faster to simulate). Eddy current and core loss models are activated and all conductors are simulated as solid copper with skin effect present. Only positive torques are compared.

None of the designs are thoroughly optimized. The EE rotor is only set up through the aforementioned basic restrictions and some minor optimization of the teeth-tip gap-width over the slots. The IPM rotor studied here is not the original IPM rotor of the reference machine (due to immaterial properties), but is designed to have the same airgap density and then optimized with the width of the V-shape and the width of the iron bridge between the magnet barriers. The resulting designs are shown in Figure 2.2.

The results are shown in efficiency maps in Figure 2.5. The IPM rotor can achieve the target output power of 56.7 kW over a wider speed range and has a higher peak efficiency as well as efficiency overall, compared to the EE rotor. However, the EE rotor can produce a higher maximum torque. A closer look shows that the efficiencies are comparable for many operating points and even better for the EE rotor in some.

Figure 2.6 shows a comparison between the efficiencies of the reference machine with the two studied rotors. A large area is covered in bright bronze, which represents operating points where the IPM rotor is more than 1 % more efficient than the EE rotor. The mint
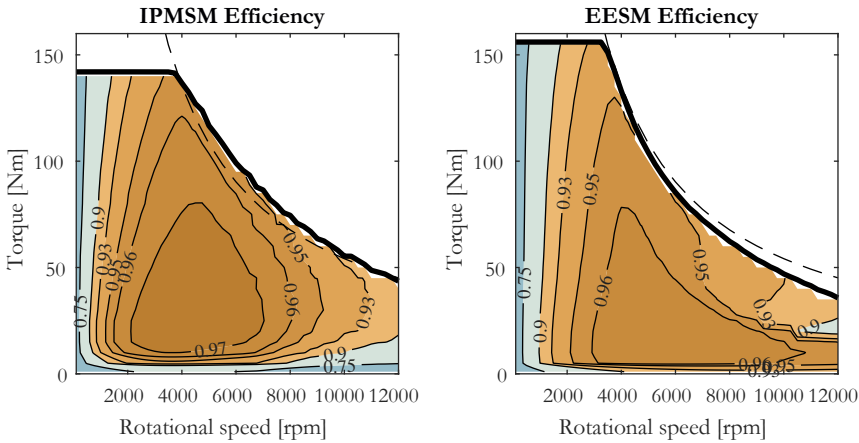
Figure 2.5: Simulated efficiency for the interior permanent magnet synchronous machine (IPMSM) and EESM respectively. Dashed line marks target continuous power of 56.7 kW.

colored (green-isch grey) areas show operating points where the efficiency difference is less than 1 % between the machines. Bright blue areas mark operating points where the EE rotor is more than 1 % more efficient than the IPM rotor. Dark bronze colored areas are operating points that only the IPM rotor can reach, while dark blue operating points can only be reached by the EE rotor. Considering the machine is originally designed for an IPM rotor, it may be surprising that a EE rotor can match the IPM efficiency in so many operating points and even be more efficient in some high speed, low power, points. Both machines have the same stator windings after all, but the EE rotor has rotor windings too. The rotor windings are even allowing some quite extreme current densities due to the direct cooling, which should intuitively generate significantly more losses than those of the IPM rotor.

To understand this better, the currents of the machines are studied (Figure 2.7). It is clear from the IPM stator current that the current increases gradually for higher torques at lower speeds, while at higher speeds the current must be increased to maintain the same torque in field weakening. The IPM rotor can produce low torques (< 20 Nm) with less than 50 A of current up to around 7000 rpm, but then it needs to increase the current even to be able to produce this very low torque level.

The EE machine works differently. The stator current increases less with the torque compared to the IPM at low speeds. The IPM needs 150-175 A to produce 100 Nm while the EE needs only 125-150 A. When the IPM reaches its maximum torque of 142 Nm, the current is 225 A, but at the same torque the EE needs less than 200. This is partly because the EE rotor with 195 A field current can produce more magnetic flux linkage than the IPM magnets, but also because the rotor current can be increased gradually, and not just producing a constant flux like the permanent magnets. The winding losses are proportional to the
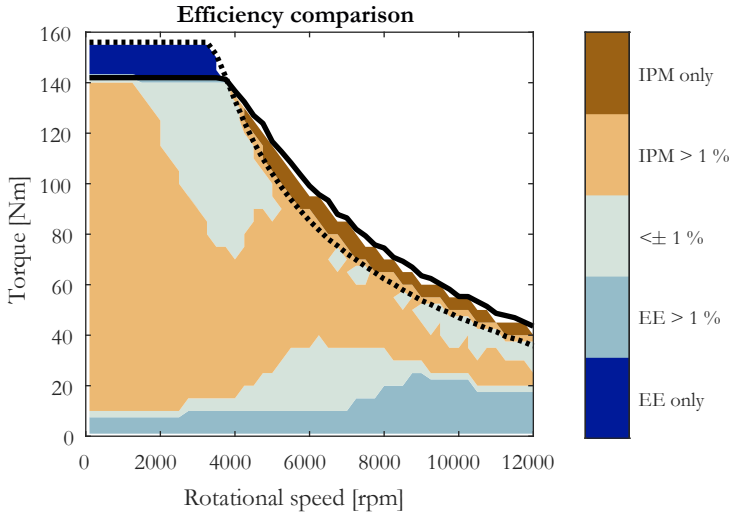
Figure 2.6: Comparison of IPM and EE efficiencies. Mint color means both machines are within 1 % in losses. Bright bronze means IPM is more than 1 % more efficient than EE. Dark bronze means only IPM can reach the operating point. Same for bright and dark blue but for EE. Dotted and black lines show maximum torques for EE and IPM respectively.

current amplitude squared and the EE machine will thus generate less losses by gradually increasing both the stator and rotor currents instead of keeping one at its maximum while keeping the other lower. This is the reason for the mint area at high torque levels in Fig 2.6 - Even though the EE machine has losses in both the stator and rotor windings, neither of them are utilized as heavy as the stator current of the IPM machine, so the total losses are still comparable (with less than 1 % difference in efficiency).

At higher speeds, the difference between the EE and IPM operations are even larger. One reason is the difference in reluctance between the two topologies. The IPM rotor has a larger q axis inductance than d axis inductance, while the EE rotor has the opposite. This means that stator current of the EE normally works in the first quadrant ($i_d > 0$, $i_q > 0$), as opposed to the IPM stator current that works in the second ($i_d < 0$, $i_q > 0$). But as the speed increases, the voltage cuts off the first quadrant and forces the EE stator current to work increasingly towards negative i_d currents (This is the main reason for the lower power output at high speeds for the EE machine). However, this can be prevented by decreasing the rotor current. A lower rotor current will decrease the torque output, but also the induced voltage, allowing the stator current to keep the i_d current higher than if the rotor current was kept at its maximum. The optimal field weakening operation for EE machine is to combine stator current and rotor current field weakening [21]. At higher torque levels, the IPM is more efficient because of the reluctance advantage, when the EE stator current needs to use much stator field weakening to reach the desired torque (Figure 2.7 and 2.8). But at lower torque levels, the IPM machine needs to maintain a high negative d-axis stator current, while the EE machine can instead reduce both the stator and the rotor currents, thus turning out to

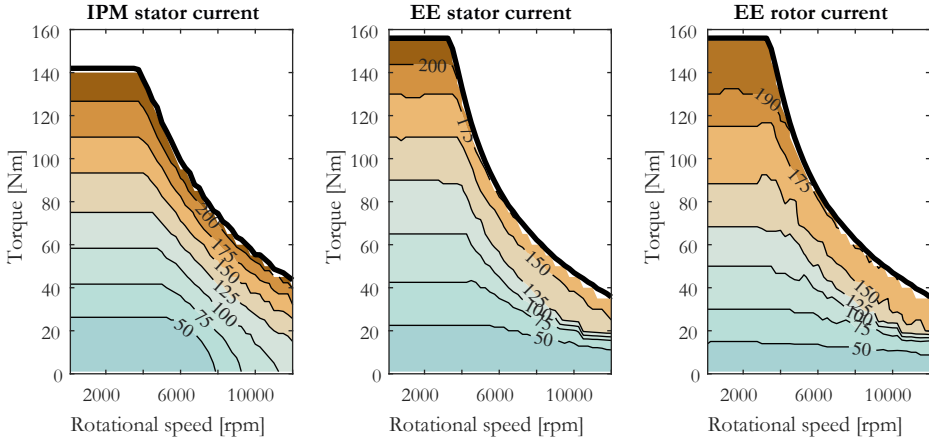be more than 1 % more efficient in these operation areas.

Figure 2.7: Stator and rotor current magnitudes for IPM and EE.
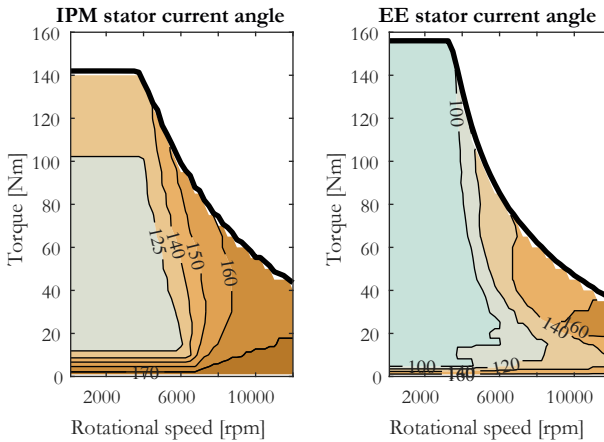
Figure 2.8: Stator and rotor current angles for IPM and EE in degrees.

To summarize the comparison, IPM and EE machines are two different machine topologies with different properties, advantages and disadvantages. The IPM is often praised for its high efficiency and performance compared to induction machines and EE machines, the most relevant competitor topologies for EV traction. But with direct cooled rotor windings, an EE machine can reach higher peak torque, reach comparable efficiency levels in many operating points and even be significantly more efficient in some.

2.3   Cooling concepts for electric traction machines

Thermal management is one of the most crucial aspects involved in the design of electric machines. This is especially true for electric machines in EV traction, because of the demand for high power density. Depending on the topology and application, there can be various heat sources of an electric machine, but the windings are always an important part to the thermal management since they are both a predominant heat source, and a thermal delimiter in the varnish of the windings which have a temperature limit that will cause a break-down of the machine if exceeded.

Historically, electric machines are cooled at the outer surface of the housing, either with passive or active convection. In EV traction, the demand for a more weight- and volume-efficient machine has lead to some more advanced ways of cooling the machines and some of them will be described here.

The thermal management of machine windings includes three parts: the heat source, the heat conduction and the heat convection. In many energineering fields there is an Ohm's law analogy to make, and thermal management of electric machines is not an exception, so these parts are shown as an electrical circuit (or a simple lumped parameter model) in Figure 2.9. Good thermal management is achieved by optimizing all these parts.
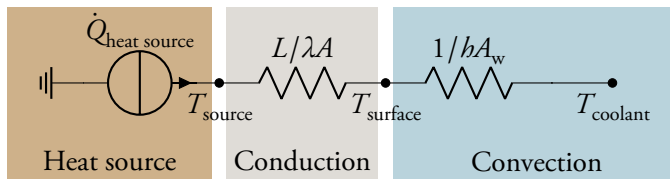

Figure 2.9: Simple lumped parameter thermal network of thermal path from heat source to cooling in an electric machine.

The heat source is the resistance of the windings and the heat that is produced there as losses when a current flows through them. This resistance can be decreased by increasing the winding cross section area. However, the goal of the electric current in the machine is to produce a magnetic flux which will generate the torque of the machine. However, increasing the cross section area of the windings means there is less space for the iron core which is needed to carry the magnetic flux. The trade-off between the cross section area of the windings and the iron core is thus a fundamental challenge in electric machine design. The resistance can also be decreased by increasing the amount of copper in the slots. This can be achieved by using hairpin windings instead of conventional wires, but hairpins may also cause increased losses due to the skin effect, which creates another intricate dilemma in electric machine design. However, this work does not focus on either of those balancing acts.

The convection is where the heat is transferred from the machine to a coolant which in turn

transports the heat away from the machine. Historically this has been done with passive convection to the air surrounding the machine, but in order to reach a higher power density of the machine, forced convection is needed. Convection is described in detail in section 2.4. As shown in Figure 2.9 it relies on two factors, the heat transfer coefficient (h) and the surface area between the solid material and the coolant (A_w).

The heat conduction is the thermal path between the heat source and the convection, and it is the main subject of this section. As shown in Fig 2.9 it depends on the length (L) and the cross section area of the path (A), as well as the thermal conductivity of the path (λ) (much like how the electric resistance depends on the length, area and electric conductivity). In order to produce a sufficient level of torque in an electric machine, a high current is desired. But the higher the current, the higher the resistive losses are (seen as $\dot{Q}$ in Figure 2.9 and equation 2.1).

$$T_{source} - T_{coolant} = \dot{Q}(\frac{L}{\lambda A} + \frac{1}{hA_w}) \tag{2.1}$$

For a given value of $L/\lambda A + 1/hA_w$ as well as a given coolant temperature ($T_{coolant}$), a certain amount of losses will results in a certain temperature at the heat source (T_{source}). This means that in order to allow a higher current (in order to increase the performance of the machine), $L/\lambda A + 1/hA_w$ must be decreased. In literature different ways to achieve this effectively in electric machines for EV traction are described.

2.3.1 Cooling concepts in literature

One way to achieve high heat dissipation is to design the housing as a cooling jacket with channels through which a coolant can be pumped. This coolant is normally water, like in the original stator housing of the machine used for the prototype in this work, as well as in the works of many others [22, 23, 24, 25, 26, 27], but there are also designs with oil [28] and air [29]. The purpose is to increase the heat transfer coefficient (h) significantly from the levels of the passive convection that otherwise takes place around the machine, while still keeping the cooling out of the way from the active parts of the machine, where the amounts of copper and iron can be maximized and balanced. A disadvantage with this concept is that it makes the thermal path from the windings to the cooling jacket quite long, as the heat has to be conducted from the windings to the stator, which is often not optimized for thermal conductivity but rather for electrical insulation (between windings and core), and further on through the stator yoke to the housing.

Another way to improve the cooling without meddling with space inside the active part of the machine is to cool the end turns of the windings. A common way to do this is by spraying oil on the end turns [30, 31, 13, 32], but can also be implemented by leading coolant

pipes through or along the end turns [33, 34, 35] or by deflecting an airflow to pass effectively over the end turns [36]. This potentially shortens the heat path, especially since the thermal conductivity along the copper windings is very good. However, if the active length of the machine is too big, the temperature gradient can still be large. End turn cooling is more effective if the surface area of the end turns is as large as possible. This can be achieved with additive manufacturing, which enables shaping the end turns for maximized surface area [37, 38]. Additive manufacturing can also be used to prolong the end turns more and enclosing them in a heat ex-changer which maximizes the heat transfer from the end turns [39].

The heat path can be improved without being made physically shorter by improving the thermal conductivity with potting. Potting means filling out empty spaces of the machine, such as in slots and around end turns, with materials like epoxy, silicone, ceramics or composites providing electrical insulation, mechanical support and better thermal conductivity [40, 41, 42]. The impact of potting is significant on the thermal management, but it also makes it more difficult to recycle the materials from the machines.

The thermal path can also be shortened by compromising with the space inside the active part of the machine in order to add cooling channels closer to the heat sources (which could be other than just the windings, although winding cooling is the scope if this work). This can be done in the stator by leading coolants through channels or gaps inside the stator core either axially [43, 44, 45], radially [46] or tangentially [47]. These solutions can prove to be effective for cooling other specific heat sources too, like core losses, permanent magnets or the windings of other machine topologies like axial flux machines. For rotor cooling, the coolant can be lead through channels in the shaft, not least present in a patent from Tesla [48].

The clearest example of minimizing the heat path, for winding cooling, is to lead the coolant in direct contact with the windings along their whole length. This can be done by having a single cooling channel somewhere in each winding slot [49, 50, 51, 52, 53, 54]. However, this still means that the insulation varnish of the windings, as well as the air in the space between the individual windings, will contribute to the thermal resistance between the heat source and the coolant. This can in part be overcome with potting [55], or by integrating more coolant pipes in each winding slot [56, 57]. These ideas get close to the concept of direct cooling, which aims to apply the cooling directly at the heat source with little or no thermal resistance in the heat path.

The works that go all the way with the direct cooling concept include laminated windings [10, 11, 12], which is a concept that has proven to be more difficult to realize in practice. With the rise of additive manufacturing, this might change. Additive manufacturing has the potential to reach extreme current densities [58, 59], but it leaves questions about production costs and material properties of the 3D-printed coils. Copper wires can be extruded

at a lower cost into a U-shaped conductor, providing a coolant channel in the concave part of the conductor, but the insulation varnish is still present in the concave part, reducing the effectiveness of the cooling [60, 61]. The varnish may even gather up and be thicker in this concave channel than around the rest of the conductor.

This demonstrates that hollow conductors is a viable solution, where varnish can be applied on the outer surface of the conductor, insulating the individual turns from each other electrically, but still allowing the coolant to come in direct contact with the copper inside the hollow channel, leaving no thermal resistance between the heat source and the coolant whatsoever. Hollow conductors for electric traction machines have been the subject of some research [62, 16, 63, 64, 65, 66, 55, 67, 68] but these projects are still in a theoretical modelling stage or on an experimental stage with smaller machines or individual conductors. However, Luvata [15] extrudes hollow copper conductors and endorse using them for direct cooling, and the company Dynamic E Flow [69] builds electric machines with direct cooled hollow conductors, indicating hollow windings for direct cooling in electric machines are at a more mature level in industry than in academia.

2.3.2 Cooling concepts comparison

In this section three different cooling concepts are studied: end turn oil spray, housing water jacket and direct air cooling with hollow conductors, shown in Figure 2.10. The aim is to point out the differences between different cooling concepts and identify some of their main challenges. The aim is not to determine which cooling concept is better, since they can be effective in different applications and also be combined in different ways to form an even better thermal management than any concept could on their own.

For this reason, one example of each of the three cooling concepts has been designed with the same resulting cooling outcome, namely the same winding hotspot temperature with the same coolant temperature.

All examples have the the same basic dimensions, consisting of a 108 mm long generic rectangular core (which could represent both a stator or rotor core). The core is 10 mm wide and 27 mm high and the slot contains three turns of a 4x3 mm copper winding (either solid hairpins or hollow with a 1 mm radius inner channel). A direct current of 335 A per turn is applied. The calculations of the examples are carried out in lumped parameter models as shown in Figure 2.11.

The *oil spray* example is shown in a radial cross-section in order to show the end turns in the housing where the oil spray is applied. The oil spray is assumed to reach half of the surface area of the end turns of each turn. According to [70] heat transfer coefficients up to 3500 W/m^{2}K are achievable with this cooling concept. The length of the end turns is
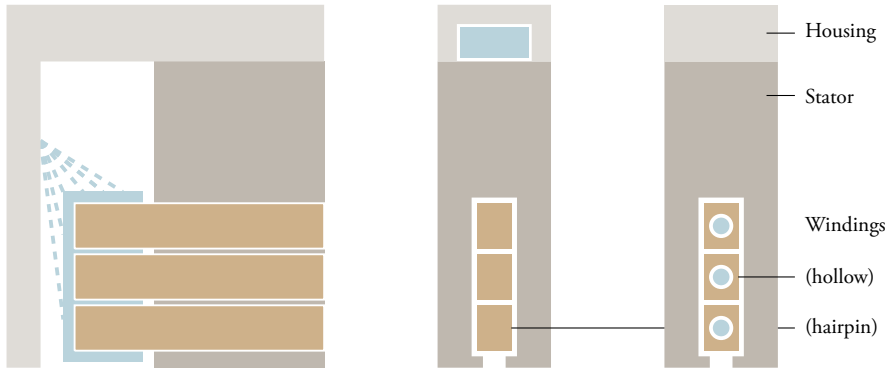
Figure 2.10: Cross section sketches of the three studied cooling methods, including descriptions. Left: end turn spray cooling (radial cross-section). Middle: Housing jacket cooling (transverse cross-section). Right: Direct winding cooling (transverse cross-section).

assumed to be 35 mm. The length of the whole machine (108 mm compared to the 120 mm of the prototype machine) is a compromise to achieve a total temperature gradient of 100 degrees from the coolant to the winding hotspot.

The *water jacket* example has a water coolant channel of 6x2.5 mm in the housing. The winding varnish is assumed to be 0.35 mm thick, in order to achieve a total temperature gradient of 100 degrees. According to [24] heat transfer coefficients up to 9000 W/m^{2}K are achievable.

The *direct cooling* has the same 4x3 mm hollow conductors with an inner channel radius of 1 mm. According to the convection model used in this work, heat transfer coefficients up to 500 W/m^{2}K are achievable with these winding dimensions.

Figure 2.11 shows the lumped parameter thermal networks used for the calculations. For the oil spray example the resistive losses are separated into to nodes with a quarter of the total losses generated in each node.. The water jacket example has two thermal resistances, one for the insulation varnish and/or slot liner, and the other for the stator core. The direct cooling example has no thermal resistance, but a higher electrical resistance due to the hollow conductors.

The modeling data are listed in Table 2.2, along with the resulting thermal resistances, heat flows and total temperature gradient ΔT from the winding hotspot to the coolant.

For water jacket cooling the main issue is the thermal resistances from the heat source to the coolant. The cooling channels in the jacket can be optimized for a very high convection, but minimizing the thermal resistance between the windings and the cooling channels are a bigger challenge. In this example the thermal resistances in the stator core and than that in the winding insulation and slot liner are of roughly equal size. If conventional smaller round wires are used in the windings instead of three large hairpin conductors, the thermal
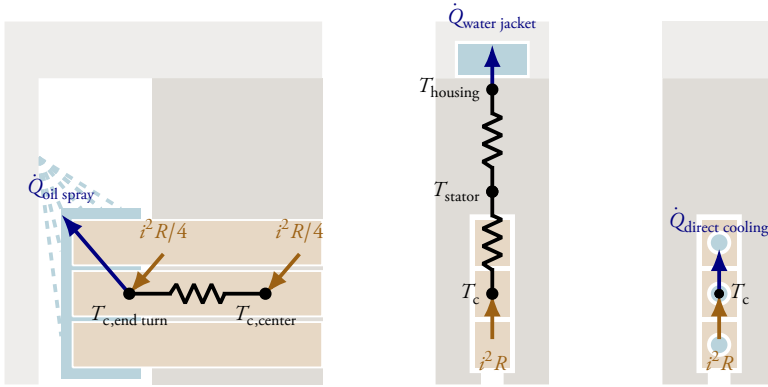
Figure 2.11: Cross section sketches of the three studied cooling methods, including thermal networks. Left: end turn spray cooling (radial cross-section). Middle: Housing jacket cooling (transverse cross-section). Right: Direct winding cooling (transverse cross-section).

Table 2.2: Cooling method comparison

		End turn oil spray	Water jacket (Insul. + Core)		Direct cooling	
Conduction	L	*avg* 30	0.35	15	0	mm
	λ	400	0.20	22	-	W/mK
	A	36	2920	1080	-	mm^2
	$\frac{L}{\lambda A}$	*avg* 2.08	0.45	0.68	0	K/W
Convection	max h	3500		9000	500	W/m^{2}K
	A_{w}	735		1190	2040	mm^2
	$\frac{1}{hA_{\mathrm{w}}}$	0.39		0.093	0.98	K/W
Total heat flow	$\frac{L}{\lambda A} + \frac{1}{hA_{\mathrm{w}}}$	1.33		1.33	0.98	K/W
	R_{winding}	74.7		74.7	101.2	µΩ
	$\dot{Q}$	75.3		75.3	102	W
	ΔT	100		100	100	K

resistance between the individual wires will add up, resulting in a temperature hotspot in the center of the slot.

For oil spray cooling, the challenge is split between the convection and the thermal resistance. The length of the machine is the driving factor behind the thermal resistance (which is why the length of the examples in this section had to be shortened slightly from the 120 mm of the prototype). The coolant's flow rate and access to as much surface area as possible of the end-turns are the driving factors behind the convection.

For direct air cooling, the issue is the extra resistance, since some copper is removed to
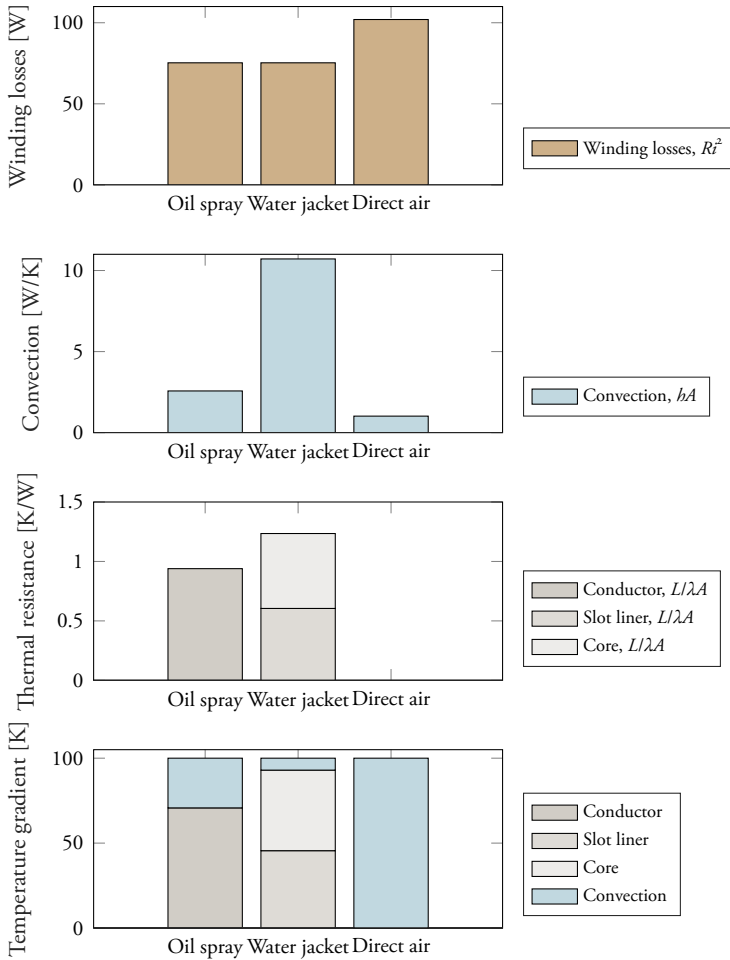
Figure 2.12: Visual representation of the data in Table 2.2, with the winding losses, total convection, thermal resistance and the temperature gradient for the cooling concepts water jacket, oil spray and direct air cooling respectively.

give room for the inner cooling channels. This, and the lower heat transfer coefficient and thus larger temperature gradient to the coolant, are compensated by the lack of thermal resistance between the heat source and the coolant.

Figure 2.13 shows the temperatures in the different examples, for various values of the heat transfer coefficient. For each case two temperatures are shown: the winding hotspot temperature and the temperature at the surface in contact with the coolant. For the oil spray example this means the temperature difference between the end turns and the windings in the middle of the length of the machine is shown. For the water jacket example this means that the difference between the winding temperature and the housing is shown. For the direct cooling example there is no such temperature gradient, since the coolant is in direct

contact with the heat source, so the direct cooling case only shows one set of temperatures. Each cooling concept is shown up to the maximum heat transfer coefficients stated above (500 W/m^{2}K for direct air cooling, 3500 W/m^{2}K for oil spray and 9000 W/m^{2}K for water jacket).
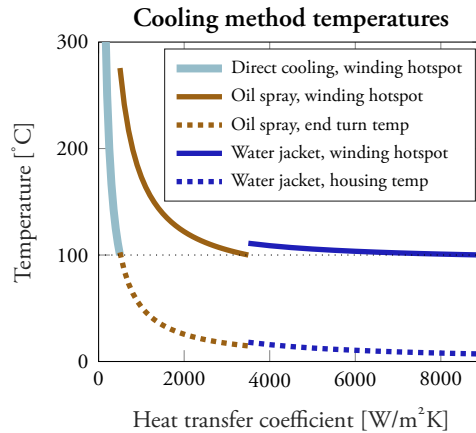

Figure 2.13: Temperatures of the three cooling concept examples (oil end turn spray, water housing jacket and direct air cooling in hollow conductors. The hotspot temperature of the windings and the temperature of the surface in contact with the coolant, are shown where they are different.

2.3.3 Cooling concepts conclusion

As described already in the beginning of the cooling concepts comparison the aim of this study is not to show that any of the cooling concepts are better than the other. Their effectiveness depend on the application and other requirements. The higher resistance of the direct cooled hollow windings will in general result in higher overall losses and lower efficiency of the machine they are used in. Oil spray requires a shorter machine length to be effective. Water jacket cooling is common in industry and literature for a reason, but its advantage when it comes to easily maximized convection in the stator does not help when it comes to rotor cooling.

Coming to a conclusion it should be kept in mind that the target of this work is the rotor windings of an electrically excited machine. The rotor winding losses can not be cooled effectively with a housing water jacket, because of the airgap between the rotor and the stator making out a major thermal resistance. The coolant can be pumped through the shaft as an alternative with a smaller thermal resistance, but the small radius of the shaft makes it impossible to reach the same surface area in contact with the coolant as with a housing water jacket. Oil spray cooling may work in rotor winding end turns, but oil spray in general is a well-explored concept already. Direct air cooling, however, is still much of an unexplored territory, which is what this thesis seeks to remedy.

2.4 Fluid coolants comparison

One of the most important objectives in the design of a cooling system for an electric traction machine is to maximize the heat flow from the heat source and removed from the machine. In section 2.3.1 the importance of optimizing the *conduction* of the heat is shown. This is accomplished by reducing the thermal resistance in order to minimize the temperature gradient between the heat source and the convection where the heat is dissipated. The concept of direct cooling is about completely removing the conductive path.

This section investigates the convection and focuses on the heat transfer coefficient and the properties of different coolants.

2.4.1 Convection and heat transfer coefficient

The next part, the *convection*, depends on the heat transfer coefficient, h, described in Equation 2.2. Convection happens when heat is transferred from a solid material to a fluid. The fluid can be either a liquid or a gas. The heat transfer coefficient depends on how the fluid flows, as well as the properties of the fluid used.

$$h = \frac{\text{Nu}\,\lambda}{D_{\text{h}}} \tag{2.2}$$

Equation 2.2 shows that the heat transfer coefficient depends on the Nusselt number Nu, the thermal conductivity of the fluid λ and the size of the channel the fluid flows in, represented as the hydraulic diameter D_{h}. The hydraulic diameter is calculated as $4A/p$, where A is the cross section area of the channel and p is the perimeter of that cross section. For a circular channel this results in $D_{\text{h}} = 2r$, just as usual, but for other shapes the result are different.

The Nusselt number is less straightforward. It is defined as the relation given in equation 2.2 above, namely the ratio between the heat transfer coefficient and the thermal conductivity, multiplied by the hydraulic diameter, but when the heat transfer coefficient is unknown this helps little or nothing. The Nusselt number has no other clear physical representation, but its relation to other properties can be measured empirically. These empirical measurements show that the Nusselt number behaves differently depending on whether the fluid has a laminar or turbulent flow. For laminar flow, when all of the fluid is moving parallel to the channel only, Hausen's formula accurately estimates the Nusselt number as equation 2.3.

$$\mathrm{Nu_{lam}} = 3.656 + \frac{0.0668\,\mathrm{Re}\,\mathrm{Pr}\frac{D_\mathrm{h}}{L}}{1 + 0.04(\mathrm{Re}\,\mathrm{Pr}\frac{D_\mathrm{h}}{L})^{2/3}} \tag{2.3}$$

Equation 2.3 still gives hardly any explaining what properties that affect the heat transfer coefficient and how. The relation is empirically derived and just introduces more new variables. D_h is the hydraulic diameter as we already know. L is the length of the channel up to the point where the Nusselt number is calculated. This means the Nusselt number will vary close to the channel entry and then be close to 3.656 as soon as $L \gg D_\mathrm{h}$. Re is the Reynolds number and Pr is the Prandtl number, which are both explained below.

The Reynolds number shows whether the fluid flow is laminar of turbulent. If the Reynolds number is below 2300 the flow is often laminar. If the Reynolds number is above 4000 the flow is often turbulent. But exactly at what point the flow shifts from laminar to turbulent is difficult to predict. When the number is above 2300 the flow can easily go from laminar to turbulent if there is a disturbance like if the channel is physically shaken, but it can also happen if someone sneezes in the next room or if a butterfly flaps its wings on the other side of the earth. However, it generally does not go back to being laminar once it has become turbulent if the Reynolds number is not lowered well below 2300. If the flow is turbulent, Gnielinski's correlation (equation 2.4) accurately estimates the Nusselt number.

$$\mathrm{Nu_{turb}} = \frac{f}{8} \frac{\mathrm{Pr}(\mathrm{Re} - 1000)}{1 + 12.7(f/8)^{1/2}(\mathrm{Pr}^{2/3} - 1)} \tag{2.4}$$

For any value between 2300 and 4000 of the Reynolds number, the Nusselt number is interpolated linearly between the Nusselt number from Hausen's formula at Re = 2300 and Gnielinski's correlation at Re = 4000. This gives a more conservative estimation of the transition from laminar to turbulent flow than assuming a sharp boundary to turbulent at Re = 2300.

Equation 2.4 introduces a new variable, the Darcy-Weisbach friction factor f. This friction factor represents how the pressure drop in the channel is affected by the surfaces of the channel and has other formulas for its definition but for the application in this work it is accurately estimated for turbulent flow by the Petukhov equation in equation 2.5.

$$f = (0.79 \ln \mathrm{Re} - 1.64)^{-2} \tag{2.5}$$

The Reynolds number is calculated as shown in equation 2.6.

$$\mathrm{Re} = \frac{\rho}{\nu} u D_\mathrm{h} \tag{2.6}$$

The Prandtl number (first shown in Equation 2.4) is calculated as shown in equation 2.7.

$$\mathrm{Pr} = \frac{\nu c}{\lambda} \qquad (2.7)$$

In summary the heat transfer coefficient depends on the thermal conductivity λ, the density ρ, the specific heat capacity c, the dynamic viscosity ν and the flow velocity u of the coolant, as well of the hydraulic diameter D_h and the length L of the channel. However, it is not evident from the equations exactly how these parameters affect the heat transfer coefficient.
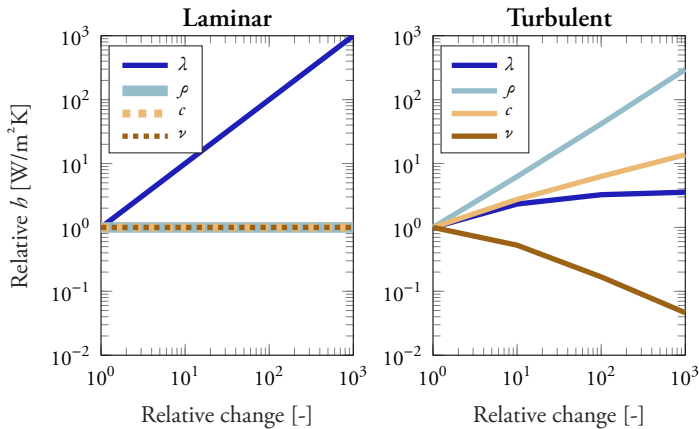

Figure 2.14: The effect of a change in different coolant properties at laminar flow and turbulent flow. The relative effect on the heat transfer coefficient is shown as functions of the relative change of the different properties.

For this reason, Figure 2.14 shows these dependencies for laminar and turbulent flows respectively, with the fluid properties varied by various orders of magnitude. The fluid velocity u is not included in these results. The results are calculated for a channel length of 120 mm (the length of the prototype machine).

For laminar flow the heat transfer coefficient is linearly dependent on the thermal conductivity λ, as described in Equation 2.2. Since the length of the channels is much greater than the diameter, the effects of the other properties are negligible.

For turbulent flow the thermal conductivity plays a smaller role, since the fluid is mixing all the time, effectively minimizing the temperature differences in the fluid while the heat does not have to be thermally conducted through the fluid. Instead the density and specific heat capacity play a bigger role, since these properties allow the fluid to consume more heat as it mixes. The viscosity also plays a role, since a lower value of the viscosity makes the fluid mix more, keeping a lower average temperature of the fluid in contact with the wall it is supposed to cool.

The coolant properties and their impact on the heat transfer coefficient provide a background to the coolant comparison in the next sections.

2.4.2 Coolant properties and cooling capabilities

Many different fluids can be used as coolants, both liquids and gases, but in order to give a comprehensible comparison this study focuses on three coolants: air, water and oil. The study also does not investigate blended or mixed coolants (like airflow with water droplet spray), or phase transition coolants (where the energy consumed in the evaporation process is used to dissipate more heat).

Table 2.3 shows the properties of air, water and oil, at 80 °C. The oil coolant properties are typical values for synthetic lubricant oils, but there is a big variety of oils that can be used as coolants. [71][72]

Table 2.3: Coolant properties at 80 °C

	Air	Water	Oil
Thermal conductivity λ (W/mK)	0.02	0.67	0.15
Density ρ (kg/m3)	1	1000	820
Specific heat capacity c (J/kgK)	1000	4160	2000
Dynamic viscosity ν (mPa·s)	0.02	0.35	20

The properties of water in the cooling channels used in the prototype will result in a Reynolds number of Re = $2u$. This means that the water flow velocity (u) needs to be more than 1150 m/s for the flow to become turbulent, which is far higher than any flow velocity that is practically reachable. For oil, the transition velocity is even higher (because of the higher viscosity), which means that for both water and oil, the flow will be laminar at all times. This means the thermal conductivity is the only important factor in comparison of the two liquid coolants in terms of generating a high heat transfer coefficient.

The airflow is several orders of magnitude lower than the densities of water and oil and thus the cooling capability is lower. The thermal conductivity and specific heat capacity are also lower, further decreasing the cooling capability. However, the low viscosity partly compensates the cooling capabilities of air.

Figure 2.15 shows how the heat transfer coefficient varies with the different coolant options and different flow rates in a cooling channel equivalent to those used in the prototype, as a function of the mass flow rate. The results show that water is superior as a coolant, considering its ability to generate a high heat transfer coefficient. At higher flow rates air almost reaches the heat transfer coefficient of oil, but here three things are important to
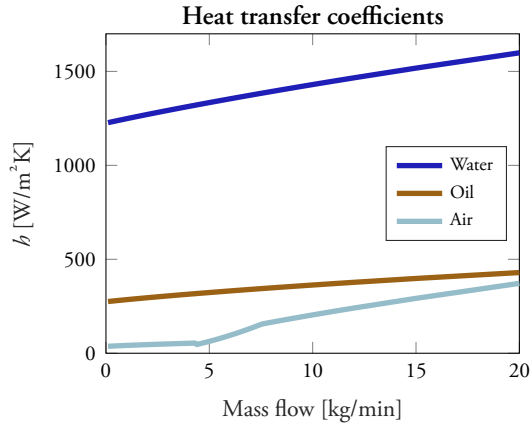
Heat transfer coefficients

Figure 2.15: The resulting heat transfer coefficients in a 120 mm long pipe with an inner radius of 1 mm, at 80 °C with various coolants and mass flows.

point out.

Firstly, the mass flow rate are compared in Figure 2.15. If the x-axis was instead volume flow rate, air would be on a completely different (smaller) scale, since it takes around 1000 liters of air to reach 1 kg, while water and oil only needs around 1 kg.

Secondly, because of the lower density and specific heat capacity of air, the air will be heated up much faster and soon be as hot as the windings it is supposed to cool. 1 liter of water can consume 4200 times as much heat as air before becoming as hot as 1 liter of air. For oil it is 1700 times as much.

Thirdly, air starts to behave differently for flow velocities of 100 m/s and higher. Air is treated as incompressible in normal conditions, but at around 0.3 times the velocity of sound in air, compressible behavior has to be taken into account. This has two consequences that are relevant to this work. The first one is that the models and equations that predict the properties of the air becomes less accurate and thus makes the models less reliable at higher velocities. The second is that it takes more pressure and thus more pipe friction losses to increase the air velocity, making further increase of the heat transfer coefficient more expensive. A mass flow rate of 10 kg/min corresponds to around 100 m/s and thus marks the limit of how fast the air can be blown through the cooling channels before further cooling capabilities becomes both more expensive and less accurate to model.

On top of this, it should be mentioned that the vaporization or boiling of water is not considered in the model. In normal pressure (1 atm = 101 300 Pa), it takes 418 kJ of energy to heat up 1 kg of water from 0 °C to 100 °C, and then another 2.26 MJ (over 5 times more) to make the same amount of water to convert from liquid into vapor while remaining constant at 100 °C. The boiling temperature of water varies with pressure and would thus not be

100 °C in a closed cooling channel, but this example still shows that neglecting boiling in the coolant model contributes to giving a conservative estimate of the capabilities of water as a coolant.

2.4.3   Coolant comparison and optimal hollow winding dimensions

For a direct comparison of the coolants, the cooling model of the licentiate thesis is adapted to the properties of water and oil, and used for a study across various dimensions of hollow conductors for direct cooling. There are many questions to answer about how to dimension hollow conductors for direct cooling. Some of these dimensions are the width of the conductors (assuming they are square shaped) and how large the inner channel for the coolant should be. Imagine a slot area in an electric machine, where some copper is sacrificed to give room for cooling channels. Is it more efficient and effective to have many small channels compared to fewer large channels? How much copper is optimally sacrificed?

Four example of the dimensions studied are shown in Figure 2.16. The two examples to the left have 2x2 mm conductors, so that 4 by 4 of them will fit inside a square with a side of 8 mm, while the two examples in the right have 4x4 mm conductors, fitting only 2 by 2 in the same space. The two examples in the top row have smaller inner channels, taking up only 8 % of the total cross section area, while the bottom row examples have larger channels, taking up 50 % of the total cross section area.
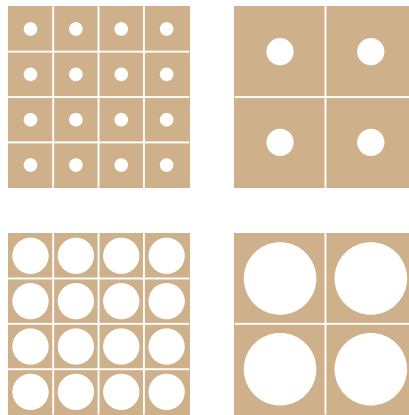

Figure 2.16: The corner cases of the study. Top row: 93 % fill factor. Bottom row: 50 % fill factor. Left column: 2 mm conductor width. Right column: 4 mm conductor width.

These examples define the extreme cases of the study, where all conductor widths from 2 mm to 4 mm are studied, along with all channel sizes from 8 % to 50 % of the total cross section area are studied. The same total current is applied in all cases, with a conductor current density of $J = 18$ A/mm^2 across the whole studied area (*conductor current density* means the average current though the cross section area of the whole conductor including the inner

channel). This means that the current through each conductor is the same, independent of the coolant channel is anywhere between 8 % or 50 % of the whole cross section area. It also means that the current through the 4x4 mm conductor is four times as large as the current though the 2x2 mm conductor (since its total cross section area per turn is four times larger). For all cases between 2 mm and 4 mm the current is scaled accordingly (and gradually even though they do not fit as neatly in the 8x8 mm square shown in the example).

In this work fill factor, ff, is defined as equation 2.8, where the whole rectangular copper cross section area is used, which for the hollow conductors used means $\frac{12-\pi}{12} = 0.74$. This means that when comparing the cooling performance of different winding dimensions with different inner channel radii, the effects of copper area loss due to rounded corners, varnish and packing factor are neglected.

$$\text{ff} = \frac{\text{Copper cross section area}}{\text{Copper cross section area} + \text{Channel cross section area}} \qquad (2.8)$$

The pipe friction losses p_{fric}, sometimes known as pumping power or just friction losses, are calculated as equation 2.9 and represents the power lost as the pressure of a fluid is decreased due to the friction against the pipe walls.

$$p_{\text{fric}} = dP \cdot \dot{V} \qquad (2.9)$$

All conductors are 120 mm long (as long as the active length of the reference machine). The coolant flow is parallel through all of the conductors. The results in Figure 2.17 show the pipe friction losses for all of these cases, for each coolant option respectively. The results are scaled to the total winding cross section area of the whole rotor of the prototype in this work so that the pipe friction losses and resistive losses can be compared to the output power of the prototype machine.
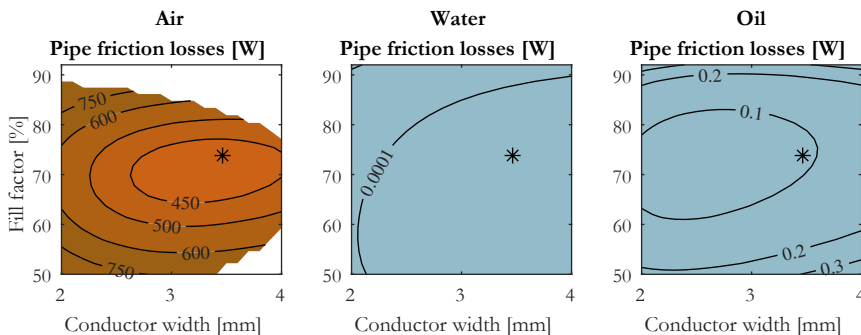

Figure 2.17: The resulting pipe friction power when keeping the copper below 160°C with an inlet coolant temperature of 60°C, for various conductor widths and fill factors, at a conductor current density of 18 A/mm².

The results show clearly how much more pipe friction losses it takes to pump the necessary amount of air, compared to water and oil. As stated earlier, the air has a lower heat transfer coefficient and heats up faster than the liquid coolants, resulting in pipe friction losses several orders of magnitude higher.

In Figure 2.17, the asterisk (*) marks the dimensions of the conductors used for the prototype. They are rectangular 4x3 mm conductors with an inner channel radius of 1 mm, which is equivalent to a square with a side of 3.46 mm ($\sqrt{12}$) and a fill factor of 74 %. These dimensions are close to the optimum for air cooling at this conductor length and current density. For larger cooling channels (lower fill factor, downwards in Figure 2.17), the resistive losses will increase, requiring more coolant to be pumped through the channels to keep the temperature below 160 °C. For smaller cooling channels (higher fill factor, upwards in the figure), the friction of the smaller channels will increase the pipe friction losses because the pressure drop increases. Likewise, for smaller conductors (left in the figure) the channels become smaller and the higher friction thus increase the pipe friction losses. For larger conductors (right in the figure), the fewer and larger channels will result in less surface area per channel cross section area, meaning there is less wall area for the heat to transfer through compared to the amount of air flowing in the channel. Both when the channels are too small (along the top of the figure) and when they are too large (in the bottom right), the inlet air velocity is too high for the model, resulting in white areas in the figure.

The dimensions of the conductors used in the prototype of this work seem to be close to a perfect trade-off between these effects. However, these are only the pipe friction losses and to draw good conclusions the resistive losses should be considered, since they are affected by the fill factor and temperatures.

In Figure 2.18 the total winding losses, calculated as the sum of the pipe friction losses and the resistive losses, are shown. For water and oil it is clear that the fill factor is the most important parameter. The higher thefill factor, the lower the resistive and pipe friction losses, with a slight tilt towards larger conductors, since smaller conductors means smaller channels and thus more friction, pressure drop and pipe friction losses. For air there is a similar trend, but the overall most efficient conductor dimensions are at odds with the air velocity limitation. For higher current densities the white area in the top right and bottom right of the figure will grow, disabling more and more dimensions from further cooling. This means optimizing the conductor dimensions for the most efficient cooling (for instance 3.5 mm width, 85 % fill factor), will lead to a less *effective* cooling design since it will not allow as high currents. On the other hand, optimizing the conductor dimensions for *effective* cooling (for instance 2 mm width, 65 % fill factor, which would allow much higher current densities before the air velocity is too high) would lead to a less efficient design. Also, this type of design would mean the cooling channels are narrower, leading to a higher risk of getting debris in the channels, disrupting the airflow potentially resulting in a failure of the whole machine due to overheating. With these aspects in mind, the

4x3 mm conductors with a 1 mm radius inner channel seem like a sound trade-off between the effects at play.
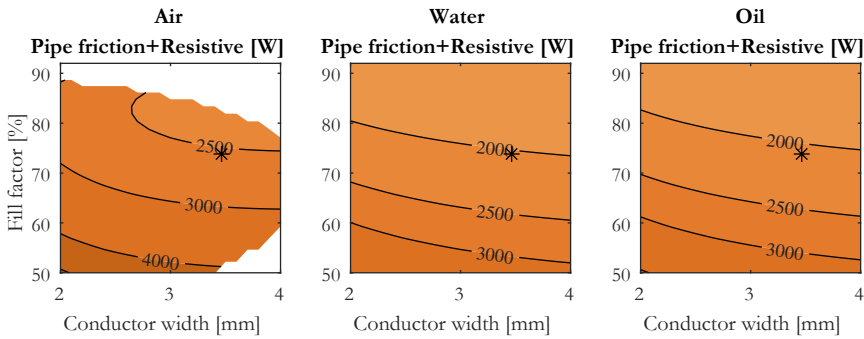

Figure 2.18: The total winding losses (resistive losses and pipe friction losses) that occur from keeping the copper below 160 °C with an inlet coolant temperature of 60 °C for various conductor widths and fill factor, at a conductor current density of 18 A/mm².

For clarity, an effective cooling system can achieve a higher maximum current density, while an efficient cooling system results in lower losses while working.

Parallel or serial coolant flow

The cases above show the results when all cooling channels are used in parallel. However, the cooling channels could also be used in series, so that the coolant flows along all the turns of a rotor pole, or even all the turns of all poles, between the inlet and the outlet. The prototype rotor has 6 turns on each of the 8 poles, resulting in 96 individual conductors. 96 is divisible by 1, 2, 3, 4, 6, 8, 12, 16, 24, 32, 48 and 96, so the coolant flow can be divided into any combination of parallel and serial flows of these numbers (like 96 parallel flows with 1 conductor in series for each flow, or 48 parallel flows with 2 conductors in series for each flow, or 1 parallel flow with all 96 conductors in series for each flow).

Figure 2.19 shows the total winding losses (pipe friction losses summed with the resistive losses) for all possible combinations of parallel and serial flows for each coolant respectively, for three different current densities. An x-axis value of 1 means there is one parallel flow (with all 96 conductors in series), and a value of 96 means there are 96 parallel flows (with 1 conductor in series per flow).

For air, it is clear that 96 parallels flows is the best design, since it is both more efficient and enables the highest current density. When air is lead through several conductors in series, the air temperature will become too hot to keep the copper below 160 °C if the airflow is not increased. When the air velocity reaches the maximum limit, higher currents can not be allowed. At J = 19 A/mm², only 96 parallel flows will require a low enough air velocity.
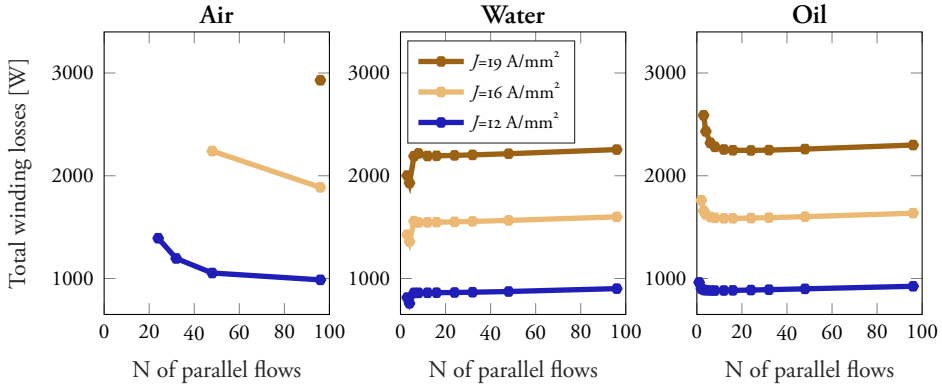
Figure 2.19: The resulting pipe friction losses that occur from keeping the copper below 160 °C with an inlet coolant tempera-
ture of 60 °C, summed with the resistive copper losses, for various numbers of parallel coolant flows and current
densities.

At $J = 16$ A/mm^2 48 and 36 parallel flows are also possible, but they require a higher flow
rate along a longer pipe, resulting in higher pipe friction losses, so 96 parallel flows are more
efficient through all current densities.

For oil, the behavior is the opposite. As the number of parallel flows is decreased from 96
to 48, the total winding losses (pipe friction losses + resistive losses) decrease, since a design
with more flows in series will result in a lower average copper temperature. This is because
the model is set up to find a flow that results in a peak copper temperature of 160 °C. This
means that the copper at the outlet is 160 °C, and then there is a temperature gradient
along the conductor all the way to the coolant inlet. Simply put, the longer the distance
from the inlet to the outlet, the larger the temperature difference, and this in turn results
in a lower average temperature, a lower average resistance (since the resistance varies with
the temperature) and a lower total resistive losses. However, the more conductors used in
series, the higher the coolant flow must be to keep the copper peak temperature below (or
at) 160 °C. When the number of parallel flows is sufficiently low, the pipe friction losses are
no longer negligible compared to the resistive losses and will increase significantly for each
step towards even lower numbers of parallel flows. Because of this, the optimal number of
parallel flows for oil is somewhere around 16 or 24 (which means having 4 or 6 conductors
in series between inlet and outlet of the coolant.

For water the behavior is similar to that of oil, but as the number of parallel flows is suffi-
ciently low, the velocity of the water becomes high enough for the flow to become turbulent
instead of laminar and that increases the heat transfer coefficient even more, resulting in
an even lower total power for very low numbers of parallel flows. This makes a number of
parallel flows around 4 a good choice for water.

With these numbers of parallel flows for each coolant respectively, the pipe friction losses

and the total winding losses (pipe friction losses + resistive losses) are calculated from the model at various current densities, with results shown in Figure 2.20. It shows the pipe friction losses are below 100 W for air cooling up to 14 A/mm^2, but the differences in the total winding losses (pipe friction losses + resistive losses) between air and water cooling is larger than 100 W at this current density. This is because water cooling has fewer parallel flows and thus a lower average copper resistance, resulting in a larger difference in total power than just the difference in pipe friction losses. The results show the maximum conductor current density with air is J = 19 A/mm^2.
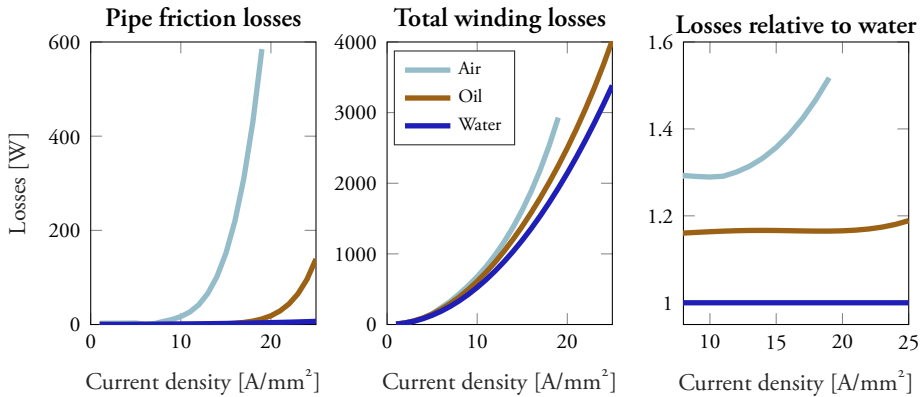

Figure 2.20: The resulting heat transfer coefficients in a 120 mm long pipe with an inner radius of 1 mm, at 80 ˚C with various coolants and mass flows.

Active control of coolant flow

For every degree (Kelvin or Celsius) that copper is cooled, the resistance is decreased by 0.393 % (see equation 2.10).

$$R_{T_2} = R_{T_1}(1 - 0.00393(T_1 - T_2))$$ (2.10)

The applied model is set up to find the least amount of coolant flow necessary to keep the copper below a certain limit (in this study 160 ˚C). This assumes the resulting coolant flow will produce the most efficient cooling case. But the results in the Figure 2.20 indicate that more power can be saved by increasing the coolant flow to further reduce the temperature, decreasing the resistive copper losses more than the pipe friction losses are increased. This points out three interesting aspects of active coolant flow control in direct cooled windings.

First, if the flow control is set up to minimize the total winding losses (resistive power + pipe friction losses) the losses of the machine can be reduced.

Second, the direct cooling conductors used in this work (4x3 mm with a 1 mm radius inner cooling channel) have 26 % less copper than solid 4x3 mm conductors. This means (according to equation 2.10) that if the direct cooled conductors are 66.6 ˚C cooler than the solid conductors the resistance is equal. In other words, for every operating point where the direct cooled windings can be kept more than 66.6 degrees cooler than the windings of an equivalent non-direct cooled machine, the resistive losses of the direct cooled machine will be lower than those of the non-direct cooled windings, despite having less copper.

Third, another important aspect of the winding temperature is the lifetime before failure. [73] has shown that with the degradation of winding insulation is accelerated with cycling temperatures compared to constant temperatures. This means the lifetime of the windings could potentially be prolonged by controlling the coolant flow rate in order to keep the winding temperature constant as current loads vary.

None of these aspects are further investigated in this work, but they still point out potential advantages with direct cooling and active coolant flow control. The nature of direct cooling means that the winding temperature will respond quicker to coolant flow changes compared to cooling concepts where the heat needs to travel farther between the heat source and the heat dissipation.

2.4.4   Practical aspects of coolant options

As already emphasized, not only theoretical aspects should be considered in the design of an electric machine prototype. Therefore some practical aspects of the coolant options have to be taken into account, in addition to the theoretical aspects presented in the previous section.

Water is difficult to use as a coolant for cooling in direct contact with the windings, because water is normally conductive. This means water short circuits the winding turns, potentially resulting in a complete breakdown of the machine. Water can still be used for direct cooling if either the water is kept in isolated tubes, or chemically clean water is used. Keeping the water in isolated tubes, preventing it from coming into direct contact with the windings adds a significant thermal resistance between the coolant and the windings, which is contra-productive since the purpose of direct cooling is to get rid of this thermal resistance (see Figure 2.9). This makes the cooling less effective. Water can be made non-conductive through distillation or de-ionization, but the water may get contaminated, which means the coolant water needs to be exchanged in order to keep the machine safe.

Oil can often be practical to use as a coolant, since it is already in use for other purposes in a vehicle, such as in the transmission, which is already implemented in electric machines [13]. If a filter is used to make sure no debris picked up in the transmission gets into the

narrow cooling channels, the transmission oil could be used for direct cooling. But this adds complexity to the oil circuit, adding places for potential leakage.

Both water and oil need to be kept in closed loop and a vehicle can only carry a finite amount onboard. In combustion engine cars, the water+glycol coolant is normally allowed to circulate in the engine until the coolant has reached 80-90 °C before it is started to be pumped through the radiator. This is partly because when the coolant temperature is lower than this, the temperature difference between the coolant in the radiator and the air flowing through it is too small to have any meaningful heat transfer. The operating temperature of the oil in an engine is also around 80-100 °C. These temperatures are higher than the 60 °C inlet temperatures used in the simulations in the previous section. This means the direct cooling in the electric machine will be less effective than in the simulations. The whole circuit of the coolant contains many parts where leaks may occur during the lifetime of a vehicle, and depending on where the leak appears, the coolant may itself pose a danger to the vehicle (if unclean water shortcuts the winding turns in the machine or if oil leaks where it might ignite).

Another complication for water and oil is the choice between serial and parallel coolant flows. Implementing only parallel flows is less efficient, and also introduces questions about how to lead the coolants through the machine. If the rotor is submerged in the coolant, what does that mean for mechanical losses due to friction? If serial flows are implemented, the conductors need to be wound in a way that ensures that the inner channel is intact in the bends, either by bending with a large radius, resulting in more bulky end turns, or by filling the channels with something that ensures the channel shape during the bending procedure. If direct cooling is implemented in the rotor (as in this work), this means the coolant should somehow be lead inside the windings of the spinning rotor. None of these tasks are impossible, but they significantly increase the complexity of the design.

With air, all of this is less complicated. The vehicle is surrounded by air that is cooler than 60 °C. This means all that is needed is a pump, fan or compressor that takes new cool air at the front of the vehicle, filters it and generates the pressure needed to push a sufficient amount through the machine windings (Generating this flow increases the air temperature, but depending on the exact setup and the ambient air temperature it should still be below 60 °C by the inlet to the windings). If all undesired paths for the air through the machine are sealed off, the airflow merely needs to be lead inside the housing on one side of the machine, for it to find its way through the channels. The used air does not need to be re-circulated, but can just be let out below or behind the vehicle. If the air leaks somewhere in the system, no particular harm is done so long as the winding temperature is monitored. If the ambient temperature around the vehicle is very hot, it could potentially be cooled by the air conditioning system, or be mixed with a mist of distilled water from the air conditioner in order to increase the cooling ability of the air (so long as the air humidity level and winding temperature ensures that no condensation takes place inside the windings). If the

outdoor temperature is very cold the used air from cooling could potentially even be used to heat up the passenger compartment of the vehicle, making efficient use of the machine losses.

2.4.5   Coolants conclusion

Water and oil are more effective and efficient coolants for direct cooling than air, but water and oil cooling also come with a number of practical complications which are easier to overcome if air is used. The losses of the air cooled hollow conductors are in general higher than those of non-direct cooled windings, but not necessarily for all current levels, since direct cooling can be used to reduce the resistive losses more than the pipe friction losses add, potentially even so much that it compensates for the lower amount of copper.

However, the most important feature is that direct cooling, even with air, can enable substantially higher continuous current levels (continuous current density of 19 A/mm^2, copper current density of 25.7 A/mm^2) than what can normally be reached without direct cooling. This ability is in itself enough to motivate the study of direct air cooled windings, despite its drawbacks compared to other coolant options, especially considering the practical aspects of implementing a direct air cooled prototype machine.

2.5   Chapter summary

This chapter presents the background of the work of this doctorate thesis in depth.

The aim of the work is to verify whether direct air cooling of the rotor windings of an electrically excited machine works in practice, and this chapter has focused on the theory and and background behind that aim. The specifications of the prototype machine are presented and the expected performance level is shown and compared to the performance of an IPM rotor. The rotor topologies have different strengths and the EE rotor is comparable in efficiency for many operating points, and significantly more efficient in some.

Different cooling concepts for electric traction machines from literature are presented and a selection of the concepts are compared, in order to point out the strengths and weaknesses of the concepts. It is shown that for rotor winding cooling, direct cooling is promising and potentially necessary in order to reach sufficient performance levels for EE rotor windings.

For direct cooling, different coolant alternatives are compared. It is shown that water and oil are more effective and efficient coolants than air in almost every aspect, except for the practical aspects - Especially for the cooling or EE rotor windings. For the studied coolant options, it is concluded that water is more efficient when used in series through the turns

of the rotor (or with a few parallel flows), air is more efficient and effective with a parallel flow through all turns, and oil is somewhere in between the two other coolants. It is also concluded that the lower fill factor of direct cooled hollow windings compared to solid conductors, does in general result in higher losses, but this can potentially be compensated by cooling the windings to a winding temperature that results in a lower total resistance than the solid conductors would have for carrying the same total current.

Chapter 3

Rotor design choices
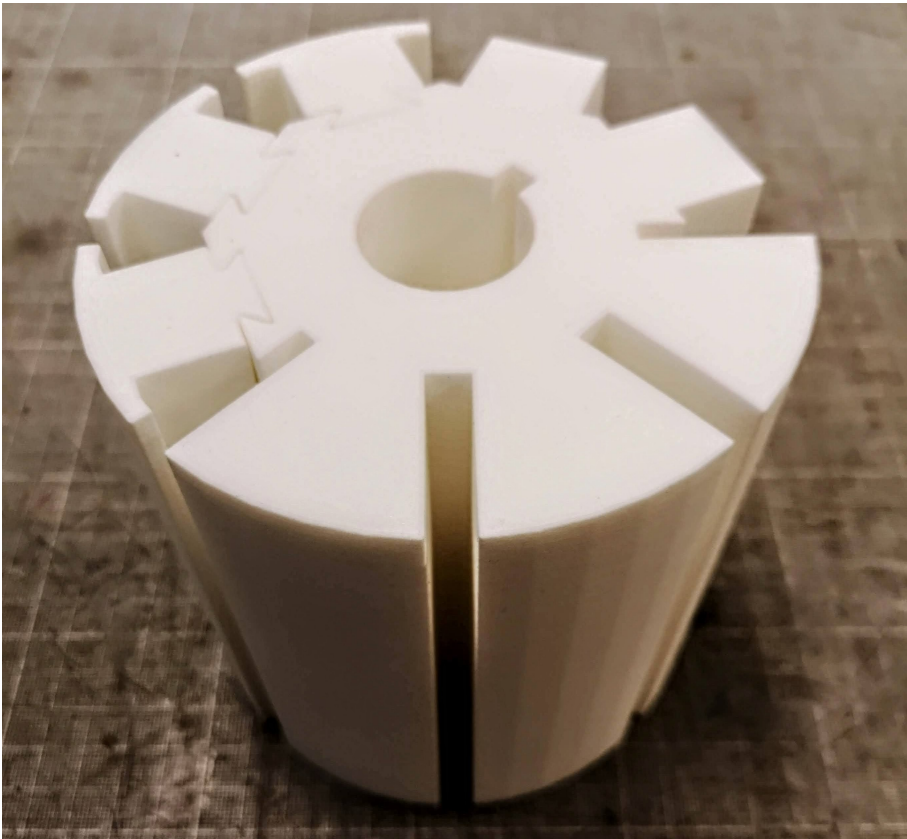
Figure 3.1: Previous page: Photo of 3D mock-up rotor for winding tests of rotor design options.

3.1 Chapter overview

The previous chapter explains the scope of the thesis and concludes that the concept of direct air cooled rotor windings for an EESM works well in theory, but needs to be verified in practice. This chapter takes the first steps in the realization of the prototype machine by looking into the rotor design with a focus on enabling manufacturing and assembly by penetrating the following:

- Whether the initial rotor design described in the previous chapter is suitable for the realization of the prototype.

- What adaptation can be done to the design in order to simplify the manufacturing and assembly of the prototype and the consequences of these adaptations on the electromagnetic performance of the machine.

- How the rotor can be wound in practice with the hollow conductors.

Parts of the content in this chapter are published in [62].

3.2 Rotor design options for prototype

A simple example of a rotor design is proposed in Chapter 2.2.1, as an intuitive outcome from the boundaries set for the design (24 mm tooth width, 12 mm yoke depth, 3 mm tooth edge, shown again in Figure 3.2). Being the initial design, this design is hereafter called rotor design option 1. Since the aim of this work is to verify the hollow winding direct cooling concept in practice, the practical aspects of the rotor design is crucial.
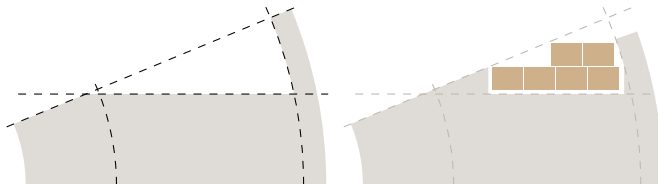

Figure 3.2: Left: Boundaries for slot area marked with dashed lines, and the resulting available slot area as white. Right: Possible winding layout with 6 turns of 4x3 mm windings within the boundaries.

3.2.1 Winding configuration

The practicality of the rotor design depends on the winding configuration and for this reason this needs to be determined as a first step.

There are four options for the winding configuration, shown in Figure 3.3. The first winding configuration option, *(a)*, means using one continuous hollow conductor through all turns of all poles in the rotor. The second configuration, *(b)*, means using one continuous conductor for each pole, and connecting the conductors between each pole. The third winding configuration, *(c)*, is basically a hairpin solution where a shorter conductor us bent in a U-shape, inserted in the slots and then bent on the other side to enable the connection between each hairpin. The fourth winding configuration, *(d)*, means every part of the conductor only goes through the length of the machine once.
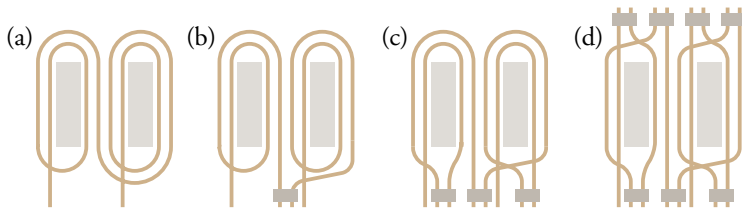
Figure 3.3: Four examples of winding configurations with 2 poles (light grey) and 2 turns per pole. From the left: (a) Continuous for all poles, (b) Continuous per pole, (c) Continuous per turn (hairpins), (d) Continuous per half turn. Soldering/connections marked with dark grey rectangles.

Option (a) has the advantage of enabling a continuous coolant flow through all poles. If water were used as a coolant this could be a good solution, since it has been shown in the previous chapter (see Chapter 2.4.3) that it is most efficient if pumped through at least two poles in series. However, this solution also comes with a practical disadvantage as it requires winding all turns of all poles either directly onto the rotor or inserted fully wound. Since the coolant of the prototype is air, there is no reason to handle this problem in this project.

Option (d) has the advantage of leaving the hollow conductor open in both ends, which is suitable for air cooling as it requires the maximum amount of parallel flows to optimize the cooling. However, this solution also needs six connections to be made on each side of each pole. It is not obvious how these connections could be made within the space available on each side of the machine without interrupting the air flow. It is also not obvious that the bends that would be necessary to connect the conductors would not impede the air flow inside the hollow conductors to an extent that would be counter-productive to the idea with open endings.

Options (c) and (d) both enable basically any rotor design, which is a clear advantage of these winding configurations. But they both require many connection to be made. If a machine with these winding configurations were to be mass-produced, this could be solved

not much harder than how solid hairpin-windings are industrially bent and welded together today. But with the equipment available for this project, implementing this amount of connections is preferably avoided. It is, as stated above, also not obvious that these solutions would simplify the inlets and outlets of the air flow.

Option (b) needs a minimum amount of connections without requiring handling one single conductor to be lead from pole to pole. It enables the windings to be wound either directly onto the poles of the machine, or to be prewound and installed on the poles and then connected between each pole. If there is a rotor design option that enables any of these winding methods without compromising too much on the manufacturability, the electromagnetic performance or the mechanical performance, this winding configuration is preferred over (c) or (d) for this prototype.

3.2.2   Practical evaluation of rotor design option 1

In order to assess the practical possibilities to wind a rotor with this design and winding configuration, a mock-up was milled out of plastic. The mock-up is 20 mm thick and 6 parts of these stacked together can represent a whole rotor with a length of 120 mm. However, only one piece was needed to realize that it is not possible to manually wind this rotor design option with 3x4 mm conductors.

Figure 3.4 shows a photo of the single piece mock-up where a prewound coil has been threaded onto one pole. This could be done since only one 20 mm piece of the mock-up was used, so that the coil could be turned at different angles so that one turn at a time could be stretched over the tooth tip into the slot opening. With a full length rotor mock-up this twisting and stretching would not be possible.
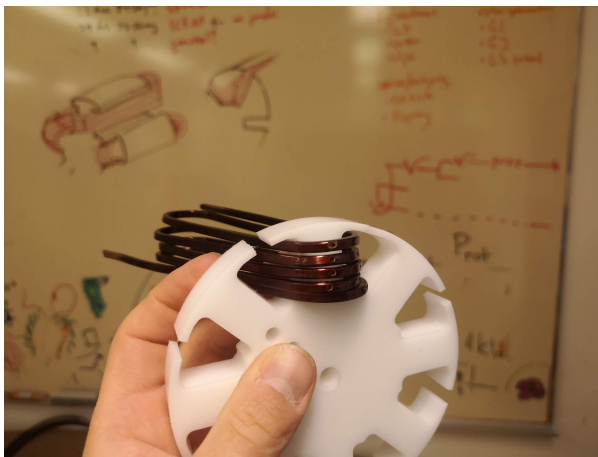

Figure 3.4: Plastic mock-up of the first design design option.

Winding the coil directly on the pole is not much easier. The stiffness of the conductors makes it difficult to get the conductor into the slot opening and then wound tightly along the length of the poles, turn after turn. The spring-back effect results in the coil adopting an oval shape, as shown in Figure 3.5.
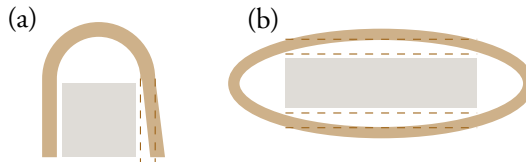

Figure 3.5: Top view of rotor pole (grey) and winding (bronze). (a) Illustration of spring-back effect from maximum bend (dashed lines) to final angle after spring-back. (b) Result of winding a conductor directly on a pole, due to spring-back: Maximal bend along tooth (dashed lines), resulting in an oval coil after spring-back.

The verdict is that the rotor design example from Chapter 2.2.1 is not possible with the winding configuration chosen, and thus other rotor design options are investigated.

3.2.3 Alternative rotor design options

With this in mind, some brainstorming resulted in three more rotor design options (Figure 3.6).

Figure 3.6: Early design sketches of rotor design options 1-3 to the left and option 4 to the right.

Rotor design option 2: Constant slot width

The second rotor design option is one with constant slot width. This geometry is shown to be the optimal geometry for electrically excited rotors in the licentiate thesis - Theoretically [9]. Practical aspects were not considered, such as the specific conductor dimensions chosen for this prototype, and even within the theoretical aspects the differences in machine dimensions, pole number, etc, may give a different result. The practical idea with rotor design option 2 is that is should be easier to thread the windings into the slot when the slot has a constant width compared to slipping it in along a narrow slot opening and then winding it tight to the tooth as design option 1 requires.

To achieve 6 turns, as in rotor design option 1, there are three winding layouts possible (Figure 3.7): Either a 1x6 layout (shown with a bright bronze colored line), a 4+2 layout (shown with a dashed bronze colored line) or a 2x3 layout (shown with a dark bronze colored line). The 1x6 layout makes some intrusion into the constrained area (12 mm rotor tooth and yoke width). The 2x3 layout makes a smaller intrusion and the 4+2 layout makes no intrusion at all. However, the 2x3 and 4+2 layouts remove some of the practical advantage with this rotor design, since they require the conductors to be wound into two layers around each pole. The 1x6 layout only needs one layer of conductors, maximizing the practical advantage of the idea and for this reason, this layout is chosen to represent rotor design option 2.
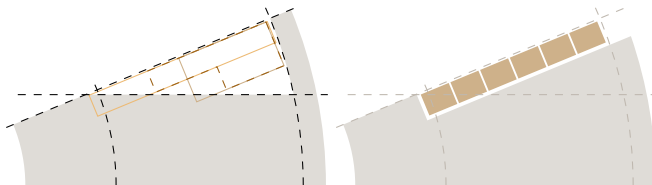

Figure 3.7: Left: Boundaries for slot area with three possible winding layouts for option 2. Right: The chosen winding layout for option 2.

The fact that the tooth is more narrow at the bottom means that the windings should be able to hold themselves in place once installed. Each turn from the bottom and out will have a wider tooth to enclose, keeping the turns from moving out into the airgap with the centrifugal force of higher speeds. For this reason, the lid suggested in the early sketch (Figure 3.6) is not needed.

Rotor design option 3: Separate tooth tips

The third rotor design option is an extrapolation of the first design, but with separate tooth tips. The practical idea is that this solution would enable installation of prewound coils.

This also enables an interesting possibility to use soft magnetic composite (SMC) material in the tooth tip and laminated steel in the tooth and yoke. SMC cores can produce less iron losses than laminated steel at higher speeds [74] but because of the inferior magnetic performance [75], SMC cores are mainly used when the design allows the advantages to overcome the disadvantages. Since rotor design option 3 uses a separate tooth tip, there is a potential that this tip could be made in SMC material and reduce the rotor core losses without affecting the magnetic performance of the machine. However the durability of the SMC is worse (around 60 MPa yield strength for Somaloy 3P series [76]) than that of laminated steel (460 MPa for SURA M235-35A [77]), so it is probable that the SMC tooth tip needs to be thicker than 3 mm, which would also affect the width of the rotor tooth, if 6 turns are still to be used. For this reason, the tooth tip constraint is increased to 5 mm in this rotor design, resulting in a tooth width of less than 12 mm. This can be seen in Figure 3.8 where the slot exceeds the 12 mm tooth width boundary.
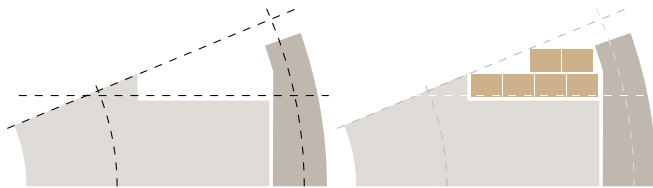

Figure 3.8: Left: Boundaries for slot area marked with dashed lines, and the resulting available slot area as white. Right: Possible winding configuration with 6 turns of 4x3 mm windings exceeding the boundaries due to the thicker SMC tooth tip.

The fastening of the tooth tips could potentially either be made with glue, screws into the laminated steel teeth or by wrapping the assembled rotor with a layer of durable tape. The tooth tip could also potentially be a whole ring (or rather a pipe) covering the entire slot opening into the airgap, although this option is not investigated further in this work.

Rotor design option 4: Dovetail joints

The fourth rotor design option separates the tooth from the yoke and connects it with a dovetail joint. A dovetail joint is basically a jigzaw-puzzle piece shaped joint that is slid on along the length of the joint. Since the rotor yoke is basically free from core losses, this solution allows the rotor yoke to be made from solid steel, lowering manufacturing costs and potentially increasing the magnetic performance (depending on the steel used), while the teeth are still in laminar steel in order to minimize the core losses close to the airgap. As shown in Figure 3.9 the slots in this design option do not exceed the boundaries, but the interface in the dovetail joint does potentially affect the magnetic flux.
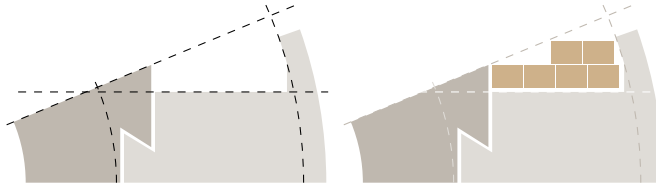
Figure 3.9: Left: Boundaries for slot area marked with dashed lines, and the resulting available slot area as white. Right: Possible winding configuration with 6 turns of 4x3 mm windings within the boundaries.

3.3   Electromagnetic evaluation of rotor design options

The main aspects to regard in the choice of a rotor design for the prototype are manufacturability and the practical assembly of the rotor. However, it is still relevant to know the electromagnetic performance of the different rotor design options since it may be considered in the final design choice. In order to enable this electromagnetic assessment a study has been made with 2D electromagnetic simulations that are run in ANSYS Electronics 2021 R1. The simulations are updated from the conference paper presenting the same study, [62].

3.3.1   Simulation setup

The 2D-models of the rotor design options used in the simulations are shown in Figure 3.10. If nothing else is mentioned, the material in the laminated core is M235-35A from Cogent SURA, and other dimensions as in Table 2.1. The mesh settings and cases used in the simulations are shown in Table 3.1, which were found to result in the most accurate results with minimum simulation time. Core losses and eddy currents effects are activated.

Table 3.1: Simulation setup for electromagnetic evaluation of rotor design options

Software	ANSYS Electronics 2021		
Mesh	General max length	5	mm
settings	Max length by airgaps	1.5	mm
	Rotor current	[25 65 130 162.5 195]	A
Simulated	Stator current amplitude	[25 75 150 187.5 225]	A
cases	Stator current angle	0.34:0.065:0.99	π rad
	Speeds	[3000 12 000]	rpm

Below, the relevant details of the simulation setup for each rotor design option are presented.
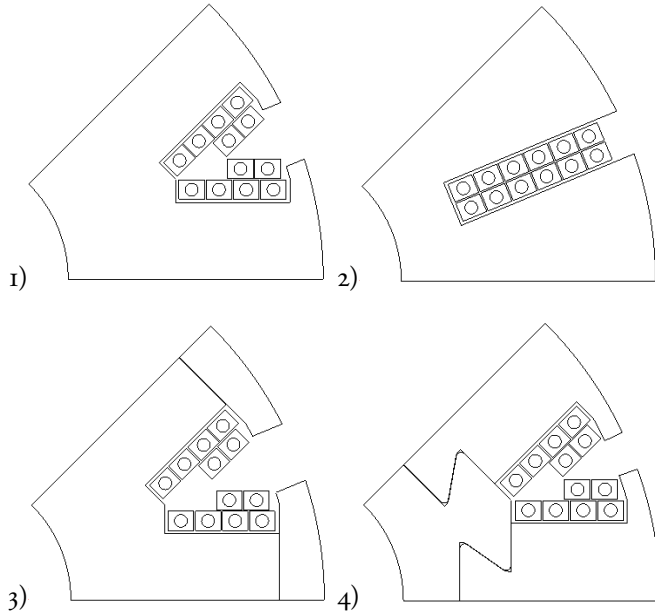
Figure 3.10: The rotor design options (1-4) as they are modeled in the simulations.

Option 1

The slot height is 18 mm, as a combination of four 4 mm conductors, two 0.5 mm insulation papers, and a 0.25 mm extra tolerance for each turns. The tooth tip is 3 mm thick, but no optimization is done on this thickness or the slot opening width for this study. The tooth width is 24 mm.

Option 2

The slot is 6 turns deep in one column, with isolation and tolerances similar to option 1. The outermost turn is situated at 3.5 mm depth from the airgap, in order to ensure some extra rotor tooth width to stop the outermost turn from being pulled into the airgap by the centrifugal force when rotating. This will put the outermost turn at the same distance from the airgap as that of rotor design option 1. The smallest tooth width is 17.3 mm.

Option 3

The tooth tip is modeled as an SMC-material (Somaloy 200 0.5Ken 600MPa). The tooth tip is 2 mm thicker than the tooth tip in option 1, to ensure some structural strength of the more fragile SMC-material. This results in the slot being deeper from the airgap, and with the same slot dimensions, while the teeth are thinner than in option 1. There is an airgap of 0.05 mm between the SMC tooth tip and laminated teeth, to simulate imperfect tolerances in the manufacturing and assembly. No account is being taken for fastening, like screws, etc. The tooth width is 20.8 mm due to the deeper slot.

Option 4

The dovetail shape is not optimized for structural strength, but a larger radius is added in the corners where the mechanical stress is concentrated during high speeds. In the middle if the dovetail, there is a 16 mm deep slit and a 6 mm diameter wedge that could be necessary to enable installation and tighten the mechanical fit. The dovetail interface is simulated with an airgap of 0.05 mm to account for tolerance imperfections. The rotor yoke is modeled as C45 steel. The slot size and tooth width are the same as in option 1.

3.3.2   Simulation results and discussion

The results of the simulations are post-processed and presented in the sections below.

Flux linkage

The resulting flux linkage (from excited rotor windings to a no-load stator winding) with various rotor currents are presented in Figure 3.11. As expected, *option 1* has the highest flux density per rotor current, since it has the least obstacles in the magnetic circuit (with the widest tooth width and no extra airgaps in material interfaces). *Option 2* follows option 1 at low excitation currents. At this point the narrower tooth width between the bottom of the slots is not saturated yet and the average tooth width is large enough to match the reluctance of design option 1. When the bottom of the tooth is saturated the derivative of the curve goes down significantly. *Option 3* has the lowest initial derivative of the curve, due to a combination of the extra airgap at the interface to the tooth tips, the narrower tooth width and the SMC tooth tips (which have a lower permeability than the laminated steel). It stays above the curve of option 2 after saturation begins, but the derivative is lower, so from 170 to 195 A the flux linkage barely increases at all. *Option 4* also has a lower flux linkage at low excitation currents, because of the extra airgap in the dovetail interface, and
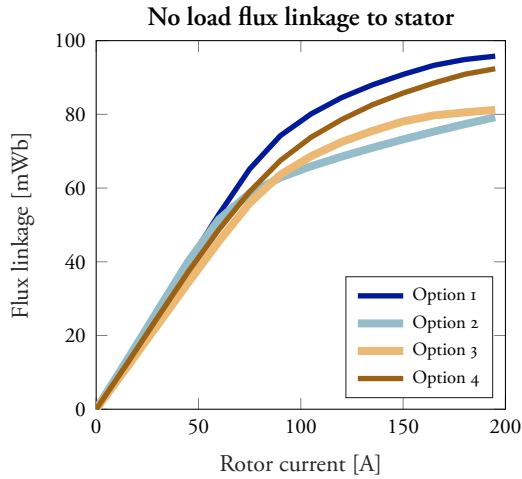
Figure 3.11: Simulation results of no-load flux linkage with various excitation currents, for all four rotor options.

because of the permeability of the solid steel in the yoke, but in the end reaches a higher flux linkage than both options 2 and 3 due the wider teeth.

All these curves could potentially change with a more extended design process, such as a mechanical analysis leading to a different tooth tip thickness of option 3 or a different dovetail design of option 4. The equivalent airgap thickness of the interfaces may also be different in practice, for instance depending on manufacturing method. Thus, no conclusions should be drawn from the exact values of the curves presented here, but design changes will not change the fundamental differences between the design options. The shapes of the curves can help to point out focus areas of a design process and also help understanding the other properties of the design options in the next sections.

Maximum torque and power

The maximum torques accessible at each rotational speed for the different rotors are presented to the left in Figure 3.12. This is the total torque from all poles, allowed by the current and voltage limitations (225 A for the stator, 195 A for the rotor, 360 V_{DC} supply - corresponding to 147 V_{RMS} for the stator windings). To the right in the same figure the maximum mechanical output power is shown.

It is no surprise that rotor design option 1 can produce the largest torque. Option 4 produces the second largest torque, but the torque starts dropping at a lower speed than the other options. All four options are similar in performance from 4000 rpm to 8000 rpm above which speed some differences are increasing - most prominent in the maximum power plot.
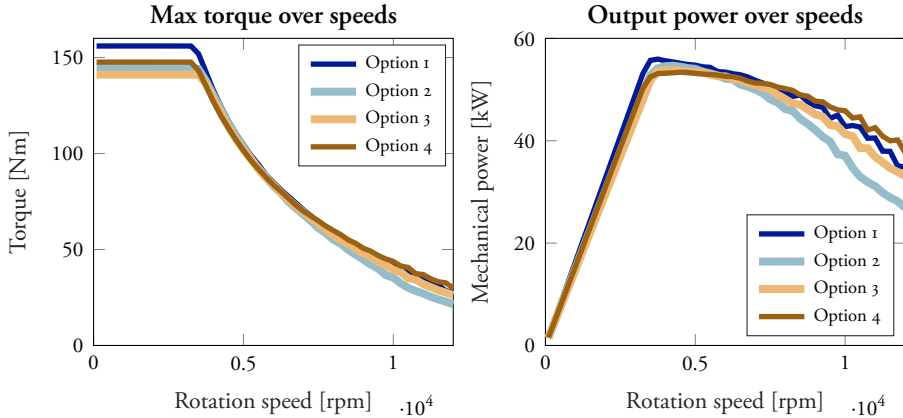
Figure 3.12: Left: Maximum torque for each speed up to 12 000 rpm for the four rotor design options. The voltage is limited to 360 V. Right: Corresponding maximum output power curves.

Above 8000 rpm option 4 produces more power than option 1. The ripples in the power curves come from the slight step-stair behavior seen at high speeds in the maximum torque plot, which is reinforced in the power curves.
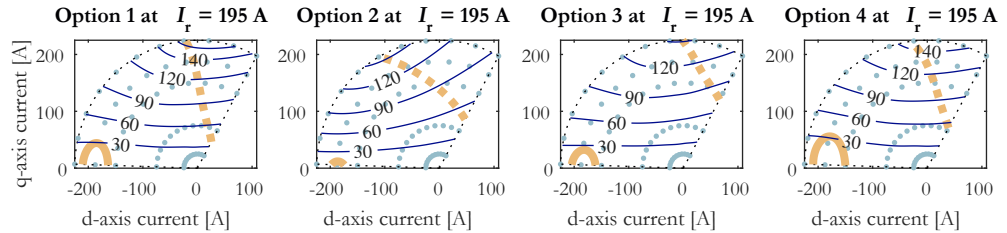

Figure 3.13: Torque map for different stator currents at 195 A rotor current, for all four rotor design options, with torque levels marked with marine blue lines and numbers. Dashed black lines mark the limits of the simulated stator currents. Beige lines mark an induced voltage equivalent of 360 V supply voltage at 3500 rpm (dashed) and 12 000 rpm (whole line). The bright blue dots mark the cases simulated, in-between which values are interpolated.

To understand the characteristics of the design options better, Figure 3.13 is shown. It shows the torque levels available with different stator currents at max rotor current, and the voltage levels that limit the available stator currents at different speeds. At this excitation level the iso-torque lines are almost horizontal, except for option 2, which tilts to the left. This shows that option 2 has a different reluctance torque compared to the other options. This is likely due to the very wide teeth close to the airgap, allowing more q-axis magnetic flux through than the other options. This difference in reluctance is also seen in the shape of the voltage curves, which are flat and wide, rather than tall and narrow. This means that despite left-tilt of the iso-torque lines, rotor design option 2 looses power faster when the speed increases, which is shown clearly in Figure 3.12. Referring to rotor design option 4 in Figure 3.13, the 3500 rpm voltage limit (dashed beige line) cuts off the peak torque area

more than for the other options, resulting in an earlier maximum power level. On the other hand, the 12 000 rpm voltage limit (beige line) reaches higher than the other options, allowing a higher maximum power at high speeds.

Efficiency and losses

The best efficiencies for all available operating points of the rotor design options are shown in Figure 3.14. The maximum torque of option 1 is shown as a dashed line with the other options as comparison. The general characteristics of the EE machine have already been explained in chapter 2.2.2, but there are some differences between the rotor design options here worth noting. All design options have a peak efficiency of over 96 %, which extends to high speeds and low power levels, where field weakening can be done with reduces rotor current instead of increased stator current. Rotor design option 3 even exceeds above 97 %, which might seem odd since it is superior to the other options in any other aspect studied yet.
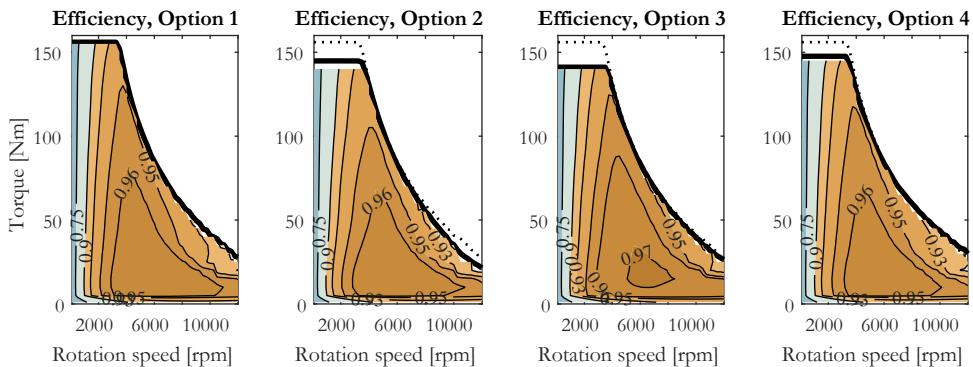

Figure 3.14: Efficiency for all operating points for all four rotor design options. The dashed line in options 2-4 mark the maximum torque of option 1 for comparison.

To explain this, the copper losses and core losses are presented separately in Figure 3.15, where the point of the peak efficiency of option 3 is marked with an asterisk (*). The SMC material in the tooth tip of option 3 can potentially decrease the core losses in some operating points, but the core losses in the rotor are small compared to the stator core losses, so the extent of this is marginal. The core losses figures show that option 3 does not have the lowest core losses in the marked operating point. Still, the SMC tooth tips are the reason for the higher efficiency, but not because of the core losses. Instead the copper losses are lower for option 3 in the marked operating point than for any other option. This is likely due to the thickness of the SMC tooth tips, which puts the rotor windings at a greater distance from the airgap than for the other options, resulting in less AC-losses in the DC rotor windings, due to less turbulent magnetic fluxes through and near the rotor
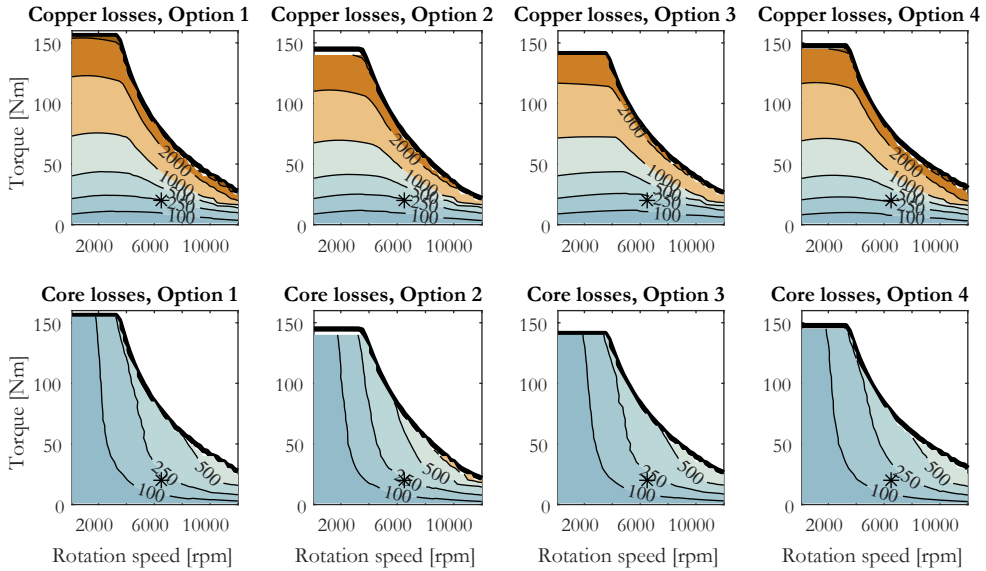
windings.

Figure 3.15: Excitation losses, consisting of resistive winding losses and the air pump power necessary to cool the windings continuously, for different excitation levels for all design options.

Conclusions on electromagnetic evaluation

The rotor design options studied are not optimized and may need further adaptions to be made due to practical aspects, so conclusions on the choice of rotor design should not be based on marginal differences which may come out different after further design work. However, it was never the purpose of this study to make up a basis for that choice, but rather to point out whether any of the proposed rotor design options should be avoided due to clear electromagnetic advantages, and to help pointing out potential focus areas in in a design process after a choice has been made on practical grounds.

Thus, the most important conclusion of this study is that there is no evident electromagnetic reason to avoid any of the proposed rotor design options. Rotor design option 1 is the overall best performing option, but it has already been ruled out due to practical reasons. Rotor design option 2 is perhaps the poorest performing option overall, but still not far behind and it does generate the highest peak power second to option 1. Option 3 stands out with its higher peak efficiency and option 4 stands out with its higher power output at high speeds. This means that the practical assessment of options 2-4 can be accepted without any reservations.

3.4   Practical evaluation of alternative rotor design options

The first rotor design option is evaluated practically in section 3.2.2. In order to test the possibilities to wind the alternative rotor design options (2-4), a rotor mock-up was produced. It was 3D printed in nylon plastic, with some poles representing each design option, shown in Figure 3.16. The tests were performed with non-hollow copper bars, with the same external dimensions as the hollow windings intended for final use.
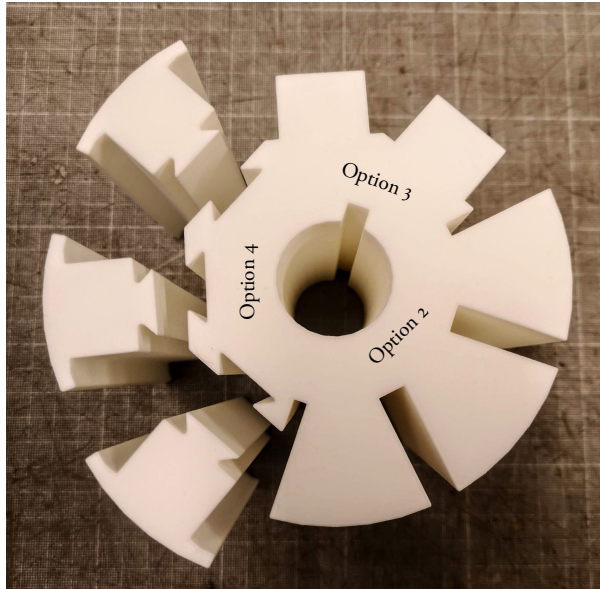

Figure 3.16: Photo of 3D printed rotor for winding tests, with rotor design options 2-4 marked.

3.4.1   Winding test of rotor option 2

Prior to testing, the hypothesis about option 2 was that the main advantage is also the biggest flaw: The angle between the slots. Once the windings are in place they will not slide off the tooth, since the tooth gets wider towards the airgap. But this also means the windings have to be applied directly on the rotor, and not on a separate fixture. The angle also means that the copper bar has to be twisted (bent in two directions) in the end winding to fit the next slot (illustrated in Figure 3.17a). How big these flaws are compared to the advantage, is the aim of the winding test of option 2.

The result of test winding option 2 shows another flaw. In addition to bending and twisting the copper bar, the bar has to be bent up to pass over the tooth (Figure 3.17b:3) before being slid down the slot into the intended position (Figure 3.17b:4). From there, the spring-back

effect result in the bar not staying in the intended position without force being continuously applied (Figure 3.17b:5). It is not impossible to hold the winding in place in the first slot while making the next end turn and thus securing it in the first slot, but it becomes a difficult and time-consuming task to accomplish with accuracy.
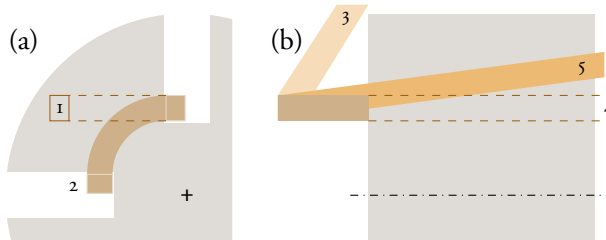

Figure 3.17: Illustration of winding tests with rotor design option 2 (here illustrated with a 4 pole rotor). (a) View from end of the machine, shaft center (+), natural bend angle for conductor (1), desired end winding after twist (2). (b) Side view, with conductor angle before being fit into next slot (3), maximum bend angle (4), angle after spring-back (5).
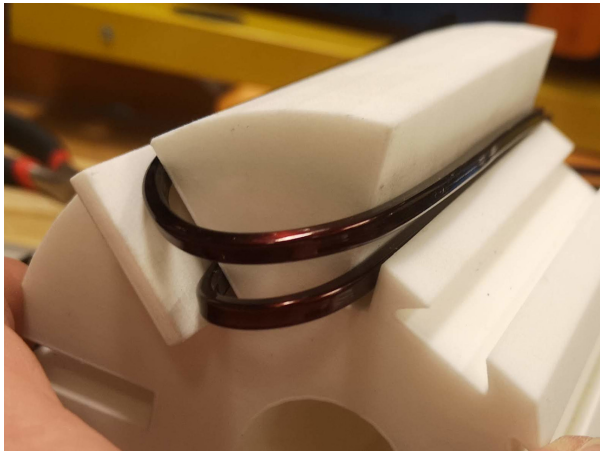

Figure 3.18: Photo of winding test with rotor design option 2, showing how the conductor will spring back up through the slot, rather than staying tight to the previous turn.

The conclusion is that the rotor design options that work with prewound coils are preferable for the prototype.

3.4.2   Prewound coils for rotor design options 3-4

With rotor design options 1 and 2 rejected in the practical assessment, design options 3 and 4 remain with the common denominator being prewound coils. The following section has two main objectives:

- Whether prewound coils can be made so that they fit on rotor teeth with rotor design option 3 or 4, without the slots being too cramped.

- How the method for making the prewound coils could be optimized regarding aspects like the space the end turns take up and the availability of the holes for air flow inlet and outlet.

Rotating coil fixture

A fixture for winding a coil was produced from 3D printed nylon plastic. The fixture is installed on a rod so that it can rotate freely as the conductor is reeled up on the fixture, as shown in the top photo of Figure 3.19. The coil will adopt an oval shape (as previously shown in Figure 3.5) and for this reason the sides of the fixture are concave (indented at most 2 mm from what would be a straight edge). When four turns are wound onto the fixture, two convex blocks are used to clamp the coil into the concave shape of the fixture, as seen in the lower photo. Because of the spring-back effect, when the clamp is released the coil will adopt a straight shape.
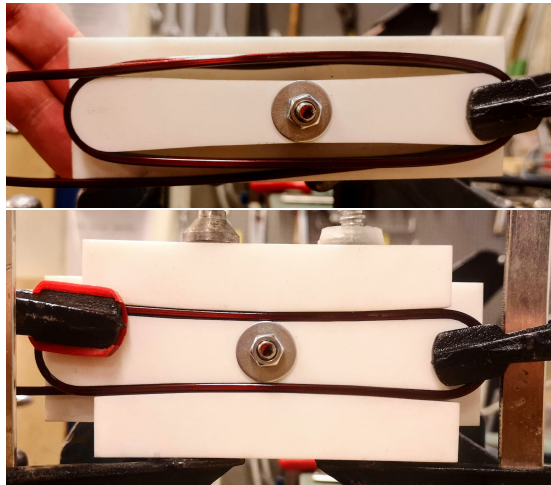

Figure 3.19: Photos of winding fixture. Above: Before pressing. Below: During pressing.

After winding the inner four turns, the two outer turns are wound with some extra distance in the end winding, and then clamped like the inner turns. The end result, with two coils with 6 turns applied on two option 4 teeth, is shown in Figure 3.20.

This test shows that rotor design options 3 and 4 work well with respect to applying the windings, since prewound coils could successfully be made and installed on the teeth of the mock-up rotor.
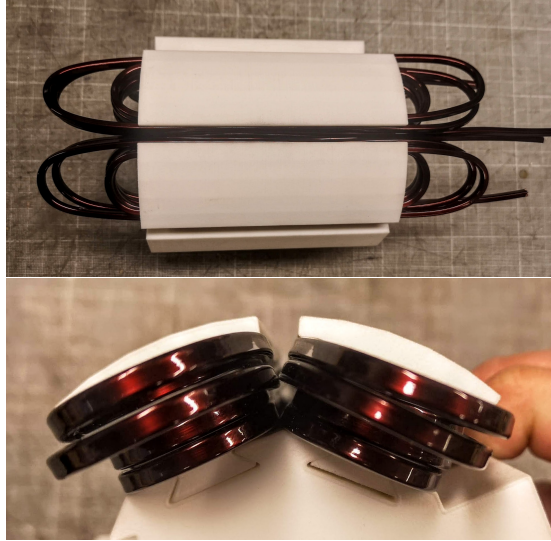
Figure 3.20: Photos of windings on teeth on rotor, completed for 2 poles.

This method of winding the coils leaves some questions. One is the space the end turns take up. The extra radius for the outer turns is exaggerated in this winding test, but there needs to be space enough between the layers for the air to get into the inlets of the inner layer of turns. Another question is how to make the holes for the air flow inlets and outlets. These holes can be drilled, but for the sake of the air-flow the drill holes need to be as far out as possible, where the angle of the end turns makes it hard to drill with accuracy. The inner turns would need to be drilled before the two turns of the outer layer are wound, which makes the procedure more complicated and time-consuming. It could work, but still calls for a better method for the production of the prewound coils.

90 degree fold method

Since no coolant flow is needed to flow through the end turns, there is no restriction on the turn radius of the end turns. This means the end turns can be folded close to 90 degrees. To accomplish this, some different methods were tested, with the aim to evaluate the aspects of a) minimizing the axial length of the end turns b) minimizing hotspot/resistance in the fold, c) ensuring a good opening for the coolant air into and out of the cooling channel, d) simplicity and time spent on procedure and e) the risk of copper debris getting into the channel, blocking the air flow.

Three main methods are tested, with some variations. The first method is bending the conductor 90 degrees and then drilling a hole to the inner channel. In the second method the hole is drilled first, and then conductor is bent 90 degrees. In the third method a small

cut can be made with a hacksaw prior to bending, making the bend result in the conductor tearing up and exposing the inner channel.

In the test, 6 samples were made to compare these aspects with these different methods for folding and producing the hole to the inner channel. The samples are listed below and shown in Figures 3.21-3.25.

1. Straight conductor for resistance reference.

2. Conductor where a 2 mm hole is drilled and then bent 90 degrees (Figure 3.22).

3. Conductor is bent 90 degrees and then drilled with 2 mm (Figure 3.23).

4. Conductor with a 0.5 mm deep cut made with a hacksaw and then bent 90 degrees (Figure 3.24).

5. Conductor with a 1.0 mm deep cut made with a hacksaw and then bent 90 degrees.

6. Conductor with a 1.5 mm deep cut made with a hacksaw and then bent 90 degrees.

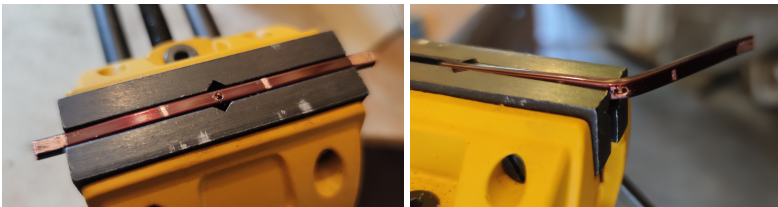**Figure 3.21:** All six test samples before bending.

Figure 3.22: Test sample 2 after being drilled (left) and while being bent (right).

To measure the resistance, a current of ~6 A was sent through the test sample and the voltage was measured between two points originally ~24 mm apart (before the bend, see Figure 3.21). The reference specimen measured a resistance of 56 μΩ. This resistance of the reference sample was adjusted to match the slight variations in length between the voltage measurement points, and then subtracted from the resistance measured for each sample, to conclude how much the resistance was increased due to the bending and hole-taking
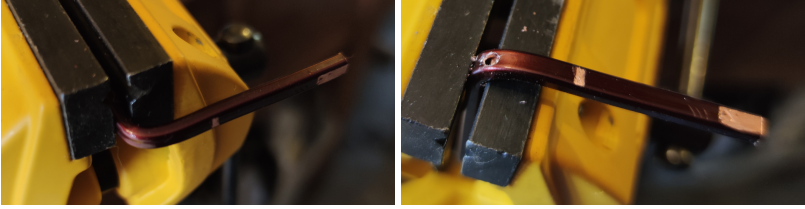
Figure 3.23: Test sample 3 after being bent (left) and after being drilled (right).

Figure 3.24: One of test samples 4-6 after being sawed (left) and after being bent (right).
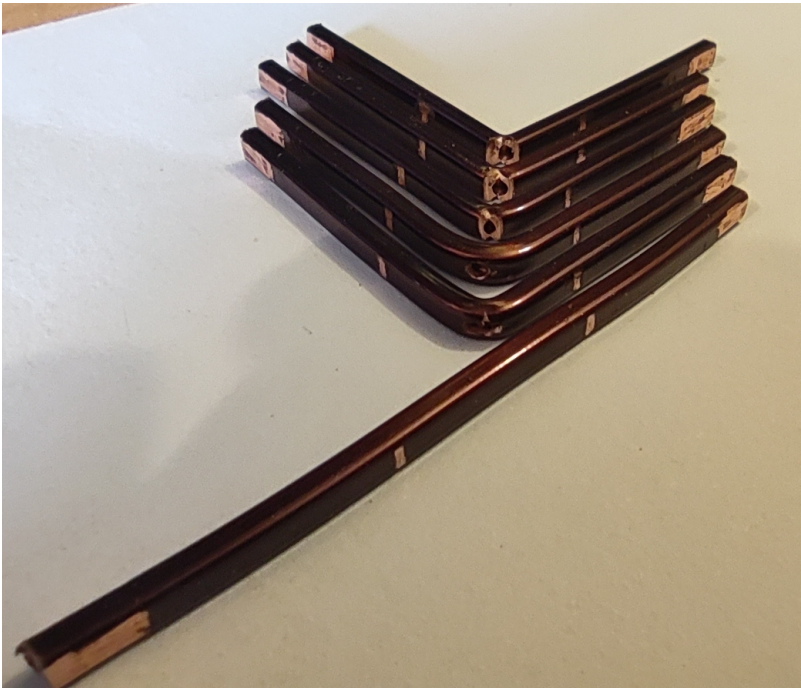

Figure 3.25: All six test samples finished, 1-6 from the bottom and up.

methods. The results from this, and the rest of the assessments, are shown in Table 3.2. All resistances were measured twice, with the mean value of both measurements shown in the table.

Table 3.2: Test scores of conductor bend test (Low score is better)

Nbr	Method	Axial L [mm]	δR [μΩ]	Hole [5-1]	Simplicity [5-1]	Debris [5-1]	Total score
1	Straight reference	-	0	-	-	-	-
2	Drill, bend	4.7	1.60	3.5	3.5	3	16
3	Bend, drill	6.3	0.44	5	4	3	19
4	Saw 0.5 mm, bend	4.4	1.79	2	1.5	1	11
5	Saw 1.0 mm, bend	3.7	3.89	1.5	1.5	1.5	12
6	Saw 1.5 mm, bend	3.3	5.10	1	1.5	3	14

After the bending procedures were done, the availability of the holes was assessed. A grade was set on a scale from 1 to 5 where 1 is the best grade. Tests 4-6 gets the best grades, because the holes are well aligned with the cooling channel. Test 6 gets the best grade because the hole is bigger. Test 3 gets a worse (higher) score because when drilling the hole after the bend makes it difficult to make the hole aligned with the channel, using handheld tools. Test 2 cracks up a bit in a similar manner to tests 4-6, but the conductor is also flattened by the bending of the conductor, closing the inner channel slightly. This results in a hole slightly less open to the channel than tests 4-6.

When assessing the simplicity of the procedures, this is done in the context in which the prototype was built, namely in the basement of the author's home, with only handheld tools available. In a professional workshop, or even in an established production line, the simplicity and outcome of the tested methods would likely be different. The purpose of this test is to establish which method is best for building the prototype, making the 192 bends needed in total for the eight coils of the rotor in a time efficient way while ensuring good inlets for the air and minimizing the risk for debris in the channels. With this background, drilling with the battery-powered handheld drill was more time consuming than using the hacksaw - especially for the already bent conductor in test 3.

The risk of debris getting into the cooling channel is assessed to be higher when drilling into the channel. Using the hacksaw to saw deep enough to reach into the inner channel will cause a similar risk. If the cut is not deep enough to reach the channel before the conductor is bent and the channel is instead cracking open, this risk is assessed to be significantly lower.

The important aspects around the choice of bending methods could be quantified and weighted in different ways. The one shown in Table 3.2 is just one way of showing that hacksawing a cut and then bending the conductor is a reasonable way to do it. If the cut is not too deep, the added resistance, and thus risk of a thermal hotspot, is not that much higher than any other way of bending the conductors and achieving an inlet for the air into the cooling channel. A drawback of this method is that it opens up an area of non-

isolated copper, which could potentially shortcut with the next turn in the coil. Another drawback is that it can be difficult to ensure that all bends are similar, so that one corner does not get more cracked and thus produces a more intense thermal hotspot. However, these drawbacks are not seen as serious enough to overshadow the advantages.

The conclusion of this test is to bend the coils with the method used for test specimens 4 and 5, with a cut of 0.7 mm into the conductor before bending it 90 degrees. With this, the cut will not be deep enough to reach the channel and risk debris in it, while it opens up the inlet hole slightly more than a 0.5 mm cut does.

Folding the 90 degree coil

With the method for bending the coils determined, a rig for the cuts was prepared by milling a 2.3 mm deep and 4 mm wide grove with a hand router, so that 0.7 mm of the 3 mm thick conductor protrudes above the surface. Hacksawing down to the surface of the rig will thus result in a 0.7 mm deep cut in the conductor.

Figure 3.26: Left: Rig for hacksawing a straight concuctor. The groove is 2.3 mm deep so that 0.7 mm of a conductor protrudes above the surface of the rig. Hacksawing down to the surface will thus result in a 0.7 mm deep cut. Right: A straight conductor 2 m conductor with 22 cuts made with the hacksaw, ready to be bent.

If two cuts are made in the conductor, 24 mm apart, and the conductor is then bent 90

degrees at the cuts, the inside distance between the two parallel conductors is 22.8-23.0 mm. I other words, the distance between the cuts needs to be 1.2 mm larger than the inside measurement of the final coil. The rotor tooth are 24 mm wide and 120 mm long and the coils need to be quite tight around them, but not too tight for insulation paper to fit in between. Some attempts were made with coils with just one or two turns before determining the suitable distances between the cuts to result in a full six turn coil. The inner four turns of the coils have a distance alternating between 26.2 mm and 124.2 mm, which is then increased in turns by 3.5 mm and 7 mm to make it to the final two outer turns. Markings were made on the rig the the 2.3 mm groove, so that the cuts could be made with accuracy and time-efficiency. With 8 rotor poles, and two spare coils, and 22 cuts per coils, over 200 cuts had to be made by hand. Figure 3.26 shows the rig for hacksaw cutting and a straight conductor with the 22 cuts made.

Figure 3.27: The rig for folding the coils, with a conductor with the first end turn completed.

The bends were made in a rig where a 3.2 mm wide groove was milled so that the conductor along the long sides of the coil would not bend when the 90 degree fold was made. Figure 3.27 shows the rig where the bends were made, and Figure 3.28 shows the procedure of folding a 6 turn coil.

With the completion of a coil produced with this method, it is concluded that prewound coils are indeed possible to make, with acceptable results regarding end turn length, resistance, repeatability and fitting on teeth after tests on the mock-up rotor.
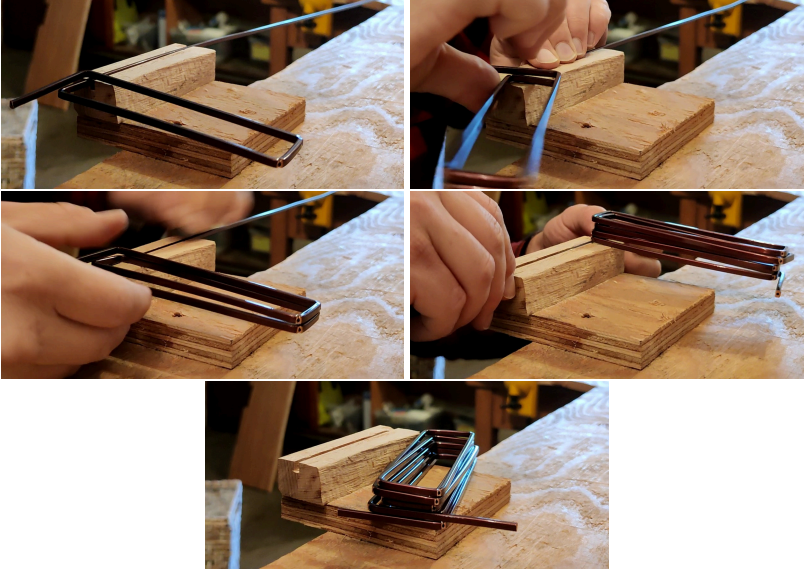
Figure 3.28: Folding a coil.

3.5 Choice of rotor design for prototype

Concerning winding the rotor, rotor design options 1 and 2 are both deemed to be time consuming and difficult to wind with accuracy. With another winding configuration, like hairpins or straight bars (Figure 3.3), these designs could be interesting, because of their electromagnetic performance (especially option 1) and simplicity. In lack of better methods of making all the connections between hairpins or straight bars, a rotor design option that allows prewound coils is preferred.

Rotor design option 3 has the electromagnetic advantage of being more efficient, due to its windings being farther from the airgap. The tooth tip thickness of 5 mm does not have a mechanical analysis behind it, but is rather chosen to make e conservative electromagnetic model. A design process of option 3 would involve balancing the mechanical and electromagnetic advantages of different tooth tip thicknesses. What is most unclear about design option 3 is how to fasten the tooth tips to the teeth. The attachment does not have to handle the weight of the tooth, but at least the weight of the tooth tips and the windings. A mechanical analysis would be needed to assess if this could be done with, for example, glue and screws. There is also an option to make a whole ring in which the core with teeth and windings is inserted into, but the material of such a ring would have to be core loss minimizing and be manufactured with sufficiently accurate dimensions not to worsen the electromagnetic performance by resulting in an interface airgap that is more than 0.05 mm thick.

Rotor design option 4 has the electromagnetic advantage of producing a higher power at high speeds than the other design options, with the dimensions used for the simulations. The dovetail joint may seem like a more complicated solution at first glance, but the material choice and manufacturing methods are deemed easier to realize than those of option 3. The dovetail joint can be milled from laminated steel and solid iron, or cut with more advanced methods. An equivalent interface airgap of 0.05 mm is small for such a complex shape, but also deemed necessary not only for the electromagnetic performance but also for the mechanical contact, to ensure maximal structural integrity.

Based on this reasoning, rotor design option 4 is chosen as the main candidate for the prototype rotor, although more work is needed on mechanical and manufacturing aspects before a prototype can be realized with this solution. Rotor design option 3 is seen as a fallback solution if the further studies on the dovetail joint result in negative results.

3.6   Chapter summary

This chapter penetrates how to realize the hollow direct air cooled windings for an EE rotor in a prototype. It is concluded that the basic designs of the first two design options are not practically executable without technologies like hairpin welding available. The hollow windings need to be prewound and the rotor thus needs to allow the installation of pre-wound coils. This can be done by either having a separate tooth top which is screwed into the rotor teeth, or with a dovetail joint at the bottom of the teeth.

The electromagnetic implications of these designs are studied and evaluated, ensuring none of the design options needs to be disregarded due to particularly poor electromagnetic performance. Two methods for producing prewound coils with the hollow conductors are tested. Finally, the rotor design option with dovetail joints is chosen as the main candidate for the prototype, based on the assessed possibilities in realizing the design in practice compared to the design option with separate tooth tips. However more work is needed on the realization of this concept.
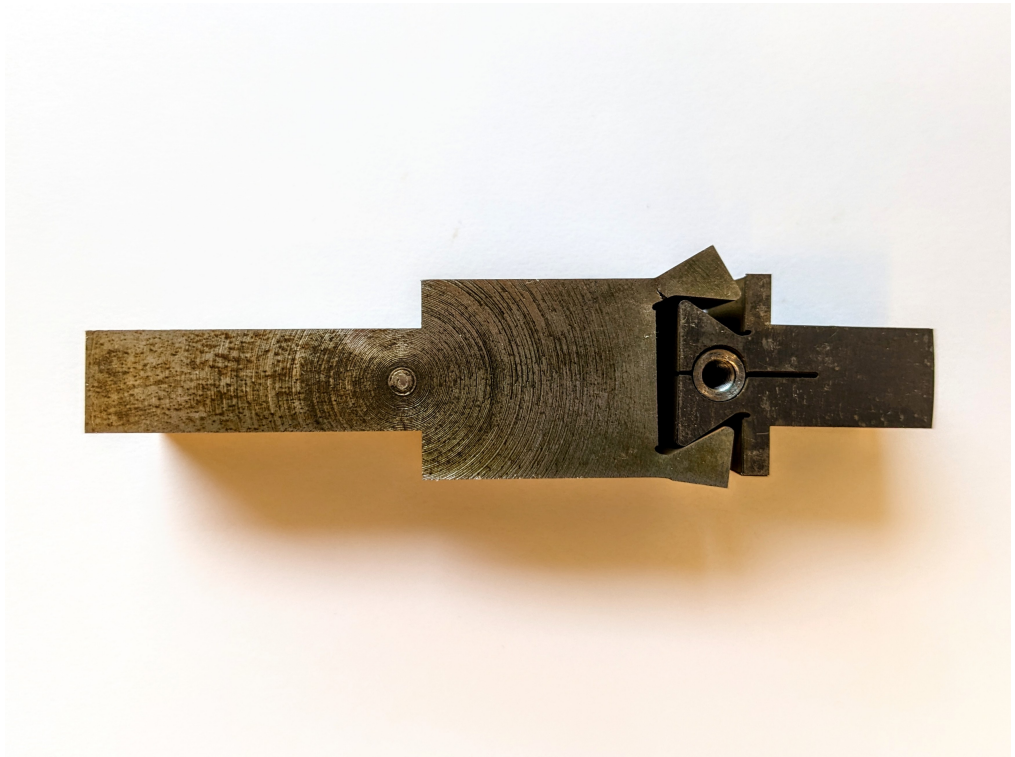
Chapter 4

Design of the dovetail joint

Figure 4.1: Previous page: Photo of dovetail joint after tensile test.

4.1 Chapter overview

The previous chapter concludes that the most promising design to realize and verify the hollow direct air cooled windings in an EE rotor, is a design with dovetail joints for the rotor teeth which allow the installation of prewound coils for the rotor windings. This chapter deals with the design of these dovetail joints, focusing on the following:

- What work has been done earlier on dovetail joints in electric machines or other applications that can contribute to this work.

- Ensuring the dovetail joints have the structural strength to handle the forces involved in meeting the target maximum speed of the prototype, set to 10 000 rpm, with simulations validated by experiments.

- What is necessary in the design to ensure that the parts are manufacturable and possible to assemble in practice.

- Minimizing any negative effects of the dovetail joint on the electromagnetic performance of the prototype.

Parts of the content of this chapter are published in [63]. The experimental tensile tests are performed by Jonas Engqvist, from which the experimental data analyzed in this chapter are taken from.

4.2 Dovetails in literature

Dovetails have been used for a long time in large generator rotors [78], but they operate with lower rotational speeds and larger dimensions than machines for EV traction. [79] presents the work on a small hybrid excited machine (PM+EE) where the coil and pole piece are inserted into a whole axial-wise, completely enclosed by the rotor core. This is an interesting solution to a similar problem, but not really contributing to the dovetail joint solution on the prototype of this work.

However, on other applications than electric machines, like in compressor or turbine blades, there is a lot of experience on dovetails. Making dovetail joints durable for rotational loads is a challenge, but the biggest and most important challenge is avoiding fretting fatigue.

This is the most common cause of early failures on compressor blades [80]. It often begins around the lower contact point and the risk increases with higher friction coefficient. Cracks often occur due to residual stresses in the joints, but can also originate around corrosion pits or dents. In multi-layer fir-tree dovetail joints, fretting often happens in the lower layers [81]. It is difficult making 3D simulations converge, so most theoretical knowledge is based on 2D simulations.

The loads on a the dovetail joints on the turbines of an aircraft engine compressor, are cyclic, multi-axial and non-proportional [82]. Repeated plastic deformation is the main cause of fretting cracks. Small cracks can cause small crumbles of material to loosen and cause more fretting where it gets stuck. Imperfections in the surface will increase the risk of cyclic plasticity, and thus increase the risk of fretting cracks to appear. This can be seen as the cracks often do not originate where the stress is expected to be higher. If vibrations are inducing fractures in the dovetail, then crowning the dovetail joint is an effective countermeasure [83].

The key lesson is that plastic deformation should be avoided. The prototype machine has a target maximum speed of 10 000 rpm. With a mass of each pole at $m = 725$ g, centered at a radius of $r = 55$ mm, the centrifugal force of each pole is around 44 kN following $F = m\omega^2 r$. The dovetail joints each have to withstand a force of 44 kN, not just once without breaking but throughout the whole lifetime of the electric machine without fretting fractures appearing and spreading. This means that the dovetail design must handle 44 kN without letting any area in the parts reach the tensile strength of the material. Other experiences of vibrations or imperfections will have to be left to future works, in order to limit the design efforts of the dovetail joint.

4.3 Experimental validation of mechanical simulations

The tool for the design of the dovetail joints are 2D FE-simulations in ANSYS Mechanics. In order to ensure that the contact problem in the simulation software is set up to be sufficiently accurate for the purpose, a tensile test with a test geometry of a dovetail joint is performed, which can then be compared to the simulation results.

4.3.1 Test setup

The test geometry with a simple 45 degree dovetail joint shown in Figure 4.2. At the beginning of the design process the materials of the final prototype were not yet decided, and the friction coefficient in the contact region was also unknown. The first experimental tests were therefore made in aluminium and simulated with frictionless contact regions. If the

material properties of aluminium will lead to accurate simulation results, then the material properties of other metals can be expected to be accurate. A frictionless contact will create conservative simulation results, so any friction in the real dovetail will be able to handle more forces than the simulations predict.
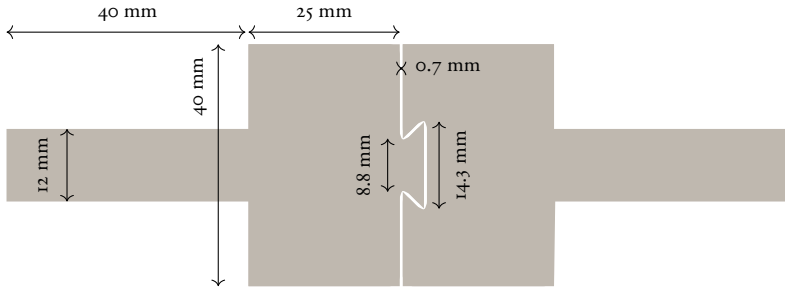

Figure 4.2: The 45 degree dovetail geometry for experimental and simulated tensile tests. The thickness is 20 mm.

The experimental validation is made with three test samples made out of water jet cut aluminium (grade 6082 T6). Aluminium is used for the samples since they can be manufactured quicker and cheaper than if they are made from the steel materials intended for the machine prototype. The tensile test was run in displacement control mode using a constant displacement rate of 2 mm/min measured on the grips of the machine. As the displacement was measured on the machine grips this measurement includes the compliance of the system. Therefore, the tests were supplemented with video footage from which the gap between the two parts in the dovetail joint could be extracted (marked in Figure 4.4).

The finite element based tensile test simulations were run in ANSYS Mechanical on a 2D model of the geometry where the 12 mm wide gripping areas were omitted. The 6082 T6 aluminium is simulated with bi-linear elasticity as presented in Table 4.1. A plane stress assumption is used rather than plane strain, since the thickness of the test samples (20 mm) is smaller than their length (50 mm between the gripping areas).

Table 4.1: Mechanical simulation setup for first tensile test

Material properties	Young's modulus	69.2	GPa
	Poisson's ratio	0.33	
	Yield strength	260	MPa
	Plastic tangent	1.2	GPa
Simulation settings	Software	ANSYS Mechanics	
	Formulation	Normal Lagrange	
	Contact type	Frictionless	
	Mesh @ contact	0.1	mm
	2D Behavior	Plane Stress	

4.3.2   Results and discussion

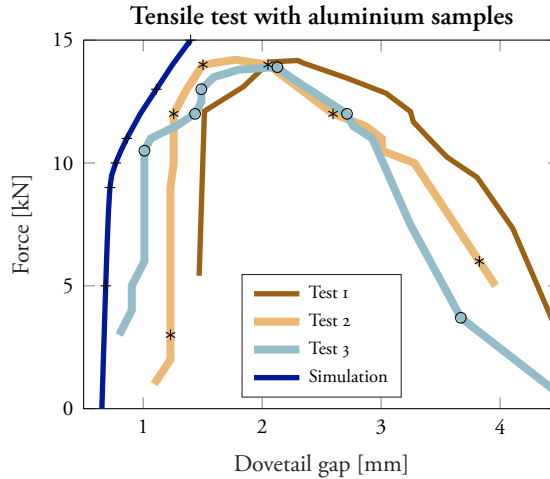**Tensile test with aluminium samples**

Figure 4.3: Force-dovetail gap curves from three tensile tests with identical parts along with simulation results. An example of the dovetail gap is shown in the top right photo of Figure 4.4 (which shows photos from the points marked with '*', 'o' and '+'.)

The force-gap results of the tensile tests and simulations are shown in Fig 4.3. The gap in the dovetail joint (shown with an example to the top right in Figure 4.4) is measured with the help of video footage for the experiments, and by comparing the displacement on each side of the gap for the simulations.

The curves from the three experimental tests show similar characteristics. The dovetail gap is different at the initial position for the different tests, due to manufacturing imperfections of the test samples.

There is a clear linear region in all three curves, ranging up to 10-12 kN, but due to the low resolution of the video footage, it is not possible to accurately derive the slope of the linear region (The lines are vertical when the change of the gap size is less than one pixel in the footage). The curve from the simulations bend off to the nonlinear region earlier, at 9-10 kN. This could be an indication that the simulations will give a conservative picture of the durability of the dovetail joint. It could also depend on the design or material properties not matching between simulation and experiments, or the fact that the simulations are modelling a frictionless contact in the dovetail, which will result in a conservative curve (reaching the nonlinear region earlier). A final potential explanation is that aluminium has slightly different stress-strain curves for tension compared to compression, while the simulation model assumes the curves are the same (as the lower tension curve).

In the nonlinear region, the maximum strength is around 14 kN for all three experimental tests. The bi-linear material properties used in the simulations, models the elastic and plastic

behavior of the material, but not the following characteristics (of hardening, necking and fracture). This means after the clear bend into the nonlinear region, the simulations are unable to accurately predict the development in the dovetail joint. However, the aim of the simulations is not to find out the ultimate strength of the dovetail joints, but to make a design that avoids fretting fatigue by avoiding any region to reach the yield strength of the material. Therefore this issue is acceptable.
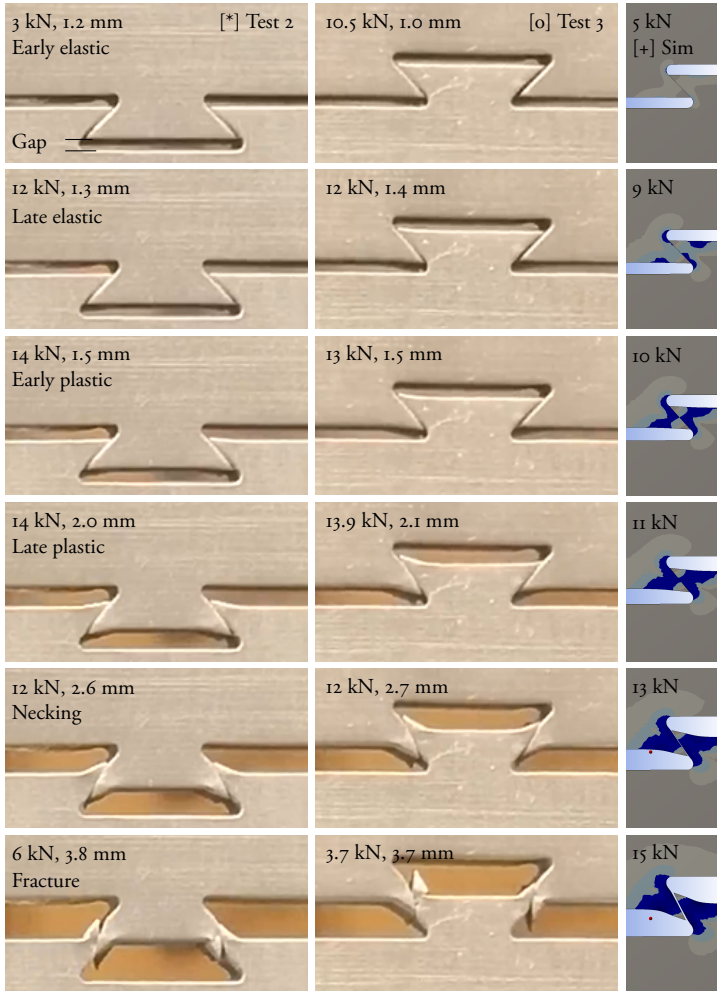

Figure 4.4: Photos of tensile tests 2 (left), 3 (middle) and screenshots of simulation (right), at different occasions (marked with 'o', '*' and '+' in Figure 4.3) and with gap size given for each test photo. The occasions are somewhat synced throughout the tensile process (early and late elastic deformation, early and late plastic deformation, necking and fracturing), but not entirely. The colors in the right column: gray to light blue is elastic deformation, dark blue is plastic deformation.

Photos of the experimental tensile tests along with screenshots from the simulations are shown in Figure 4.4. The marks (* o +) in Figure 4.3 show what steps the pictures are from.

The pictures shown are somewhat synced throughout the tensile process, with early and late elastic deformation, early and late plastic deformation and finally necking and fracture.

The photos from the experimental tensile tests show that the failure occurs due to large deformations of the dovetail flanks. In test 3, some asymmetric behavior is seen. This might stem from asymmetric gripping of imperfections due to the water jet cutting. This could explain the differences in the characteristics in the plastic region between the different tests. The videos of both test 1 (for which the footage is not of the same quality, hence it it not shown here) and 3 show this asymmetric gripping, and they both have similar dents in the force-gap curves before the ultimate strength, while the grip in test 2 is more symmetric and have a more even transition from linear to nonlinear behavior.

The images from the simulations give a better understanding of the transition from elastic to plastic behavior. The results of this is supplemented with the data in Figure 4.5, which shows the equivalent stress (von Mises) of three different points in the dovetail joint. The color turns from grey to dark blue at the yield stress level. The first areas reach the yield strength early (at this point, too tiny to see in the screenshot), at around 4 kN, by the *neck* of the joint. The opposite side, the *base*, reaches yield stress levels slightly later. As the force continues to increase, the stress in these areas ceases to increase, as they *yield* for the stress and behave plastic. The areas that have reached yield strength grow in size as the force increase (see 9 kN in the rightmost column of Figure 4.4), until the areas originating in the neck and the base finally meet and overlap at the midpoint (see 10 kN in Figure 4.4). At this point, the stress is almost equal over the whole dovetail joint, and it is not until this happens that the whole dovetail joint starts to behave plastic and enter the nonlinear part of the curve in Figure 4.3.
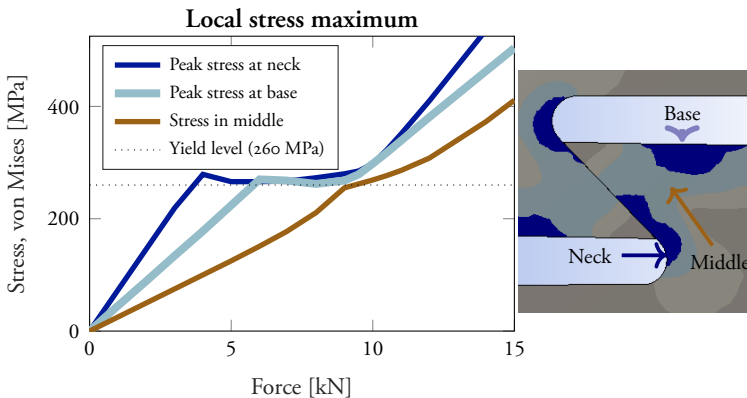

Figure 4.5: Local maximum of stress concentration around the neck and base and local stress at the midpoint (approximately pointed out to the right) for different force loads, from the tensile test simulations.

4.3.3   Conclusions from first experimental validation

Two major conclusions are drawn from this experimental validation. The first is that the simulations gives a reasonably accurate representation of the dovetail joint transition into nonlinear behavior. The simulated results are conservative compared to the experimental results, which could be due to a number of reasons in the simulation settings, such as the frictionless contact setting, the plane stress assumption or using the same yield strength for tension and compression.

The other conclusion is that the prediction of the nonlinear behavior is not as important as that of regional plastic behavior. Unfortunately it is difficult to derive from the experimental data when the points in the material with the highest stress enter plastic behavior, so it can not be said with certainty that the simulations predict this procedure accurately. But with the level of accuracy in the simulations for the rest of the procedure, some confidence can be held in the accuracy of the simulations when it comes to predicting regional plastic behavior.

4.4   Dovetail rotor design

In the design of the dovetails of the prototype, there are several aspects to consider, discussed in detail below. For dovetail joints, *tenon* is used to describe the protruding part of the joint that of the rotor teeth (sometimes know as *male* joint) and *mortise* is used to describe the socket or groove in the rotor core into which the tenon is inserted (sometimes known as *female* joint).

4.4.1   Manufacturing constraints on the dovetail design

The manufacturing method of the dovetail rotor (which is described in more depth in Chapter 5.3) sets some constraints on the design. The rotor teeth are manufactured in two steps: Laser cutting and bonding of the laminated steel with a rectangular bottom of the teeth, and electric discharge machining of the dovetail joint. These steps are shown in blue and bronze respectively in Figure 4.6. The rotor slot depth is 18 mm (as stated in Chapter 3.3.1), while the rectangular laser cut teeth has a length of 26 mm. This leaves 8 mm for the dovetail male joint, which makes up a definite limitation for the depth of the dovetail joint.

The bronze colored shape for the final design of the dovetail in Figure 4.6 shows another important feature added to enable the practical assembly of the joint. In the middle of the joint there is a slit and a circular hole. The hole is for the insertion of a wedge. The wedge
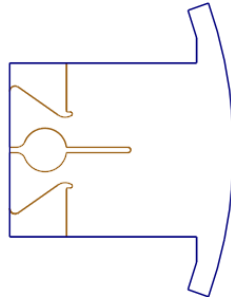
Figure 4.6: Sketch of the shape the laser cut tooth in blue, and the shape to be cut with electric discharge machining in bronze (with the final design of the dovetail tenon joint).

is an almost cylindrical, slightly conical, piece of metal which is hammered into the hole in the joint, in order to press the sides of the joint firmly into the mortise joint. The hole is also slightly conical, with a smaller radius towards the middle of the length of the rotor. However, all modelling is done with the largest radius. During assembly, the parts (teeth and core) need to be able to slide into each other, while after the assembly the fit should be perfect with no space for axial wriggle. This design fulfills that with minimal obstruction of the magnetic flux. This feature is seen in some, but not all, of the upcoming examples from the design process of the dovetail, since the idea was introduced in the middle of the design process. Also seen in the examples is a gap at the bottom of the joint, which is also needed to allow assembly and fixation of the joint.

4.4.2   Angle of the dovetail joint

The angle of the dovetail joint flanks is an important aspect. The prototype has 8 poles, which results in an angle of 45 degrees between each pole, or 22.5 degrees from each pole to the midpoint between the poles. This means there are two important necks to consider in the dovetail rotor, shown in blue and red in Figure 4.7.

The stress concentration in the joint is affected by the angle of the contact region, but also, following from the angle, the width of the necks and the length of the contact region. In order to find the optimal balance of these factors, a series of simulations were run. The cases run are shown in Figure 4.8, where the tip of the dovetail joint is limited to a 22.5 degree angle, ensuring a certain neck width, while the dovetail angle is varied, thus deciding the depth of the joint with respect to the 22.5 degree line.

The results are shown in Figure 4.9. The results show that the length of the contact region gets longer as the angle is increased. The stress peak by the neck of the tenon joint has an inverted dependency, and grows with a smaller angle. The stress peak by the neck of the mortise joint on the other hand is higher both for high and low angles, with a minimum
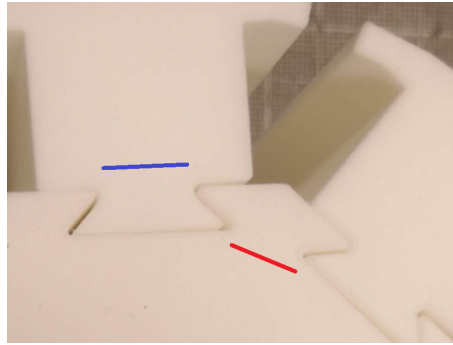
Figure 4.7: The tenon joint neck marked with blue and the mortise joint neck marked with red, on the plastic 3D mock-up of the dovetail rotor.
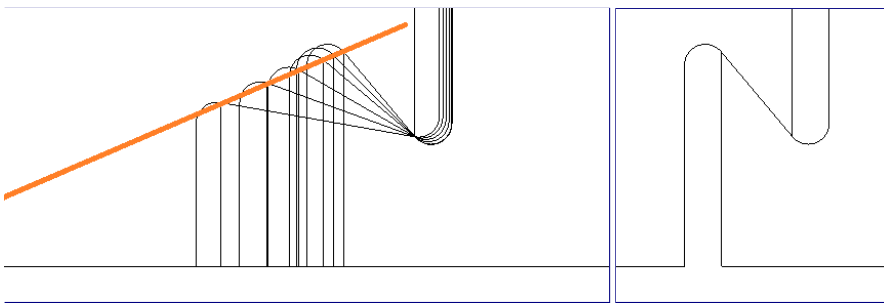

Figure 4.8: Left: All the investigated dovetail angle designs overlapping (with 40, 45, 50, 60, 70, 80 degrees measured from from a vertical line). The depth of the dovetails is constrained by a 22.5 degree line (from horizontal, marked in orange), in order to give room for 8 poles. The radius of the tenon and mortise necks are 7 mm. Right: Only the 40 degree design.

value somewhere between 50 and 60 degrees. For this reason the final design has a dovetail angle of 55 degrees.

4.4.3 Radius of the dovetail joint necks

The radius of the necks is another important aspects to consider. The stress tends to concentrate around sharp corners and the larger the radius of the corner, the less the stress concentrates. Increasing the radius of the neck may decrease the length of the contact region, if made as shown in example b compared to a in Figure 4.10. However, it can be made as example c, where the radius is extended into the core instead. Both examples b and c reduce the stress peak at the mortise neck by around 20 % compared to example a (For example c, the longer contact region length is compensated by the narrower neck width).

A larger radius at the neck can be seen as if it spreads out the stress over the larger radius, instead of concentrating it over the smaller radius. This can be seen in Figure 4.11 where
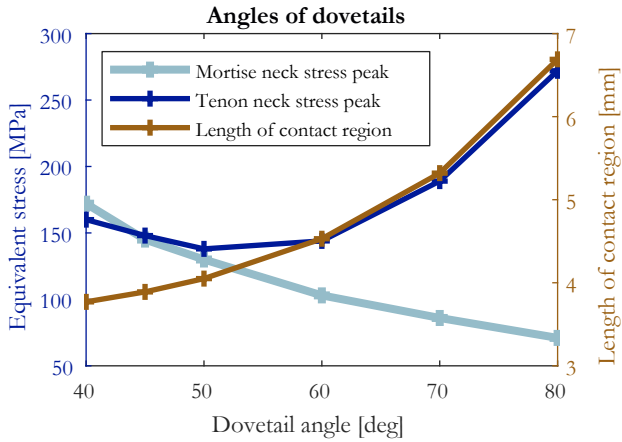
Figure 4.9: Simulation results from different angles of dovetails, showing equivalent peak stress at the tenon and mortise dovetail neck (black, left y axis), and the length of the contact region (grey, right y axis).
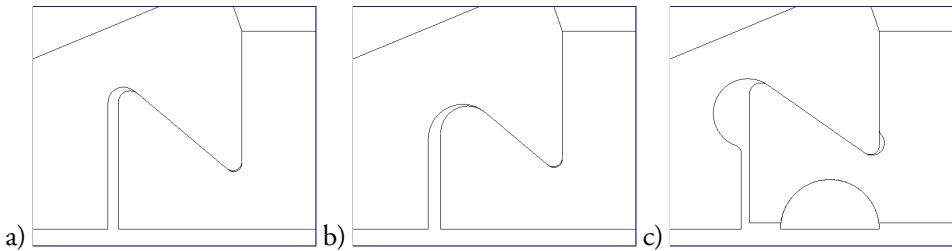

Figure 4.10: Three examples of dovetail designs with different approaches to the radius of the mortise neck (to the left).

example *b* has a larger radius than example *a*. The bottom row shows the stress concentration, where the blue area (with a lower stress concentration than the black area) stretches farther to the right. The stress peak of example *b* is 4 % lower than the stress peak of example a. After the stress concentration of example *b* have subsided, the radius switches to at sharper radius, instead of continuing the large radius like in example *a*. However, if the larger radius is not extended far enough, like in example *c*, the stress concentration will increase again. In example *c*, the black region of higher stress concentration stretches into the smaller radius, and in this area the stress peak of is 14 % higher than the peak in example *a*.

The optimal distribution of stress concentration would be found by gradually decreasing the radius, like an elliptic curve, but that type of curve could not be handled by the software of the cutting machine that was used to cut the dovetail shape, so for this reason the shape consists of two different radii in examples *b* and *c*, as well as in the final design.
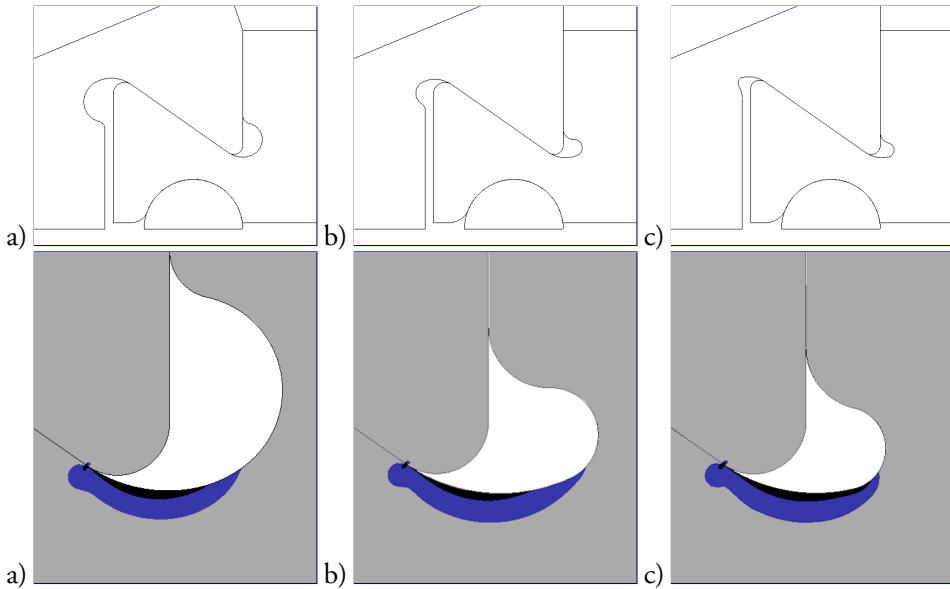
Figure 4.11: Top row: Three examples of dovetail designs with different approaches to the extended radii of the necks. Bottom row: Zoom in of the stress concentration by the tenon neck of the three examples. Black regions marks high stress concentration and blue marks a lower stress concentration (75 % and 50 % of the yield strength of the material respectively, at a force equivalent to 10 000 rpm).

4.5   Experimental verification of final dovetail design

All the aspects discussed in the previous section are implemented in the final design, while minimizing the electromagnetic impacts of the design. The resulting design is shown in Figure 4.12, with both the stress distribution and the magnetic flux density. The materials were determined to be, for the rotor teeth M235-35A laminated steel from Cogent SURA, and for the rotor yoke C45 solid steel. Since the magnetic flux is relatively stable in the rotor yoke, the core can be solid without having any notable impact on the core losses.

4.5.1   Test setup

In order to verify the final dovetail joint design another tensile test sample was made, shown in Figure 4.13. The tensile test was set up as in the experimental validation of the FE-simulations, described in section 4.3.1. The test sample has a thickness of 23 mm.

The tensile test results are compared with mechanical simulations. The simulation model has a mesh size of 20 μm and the contact region is modeled as both frictionless and frictional. The frictionless model gives a very conservative result, while the results of a model
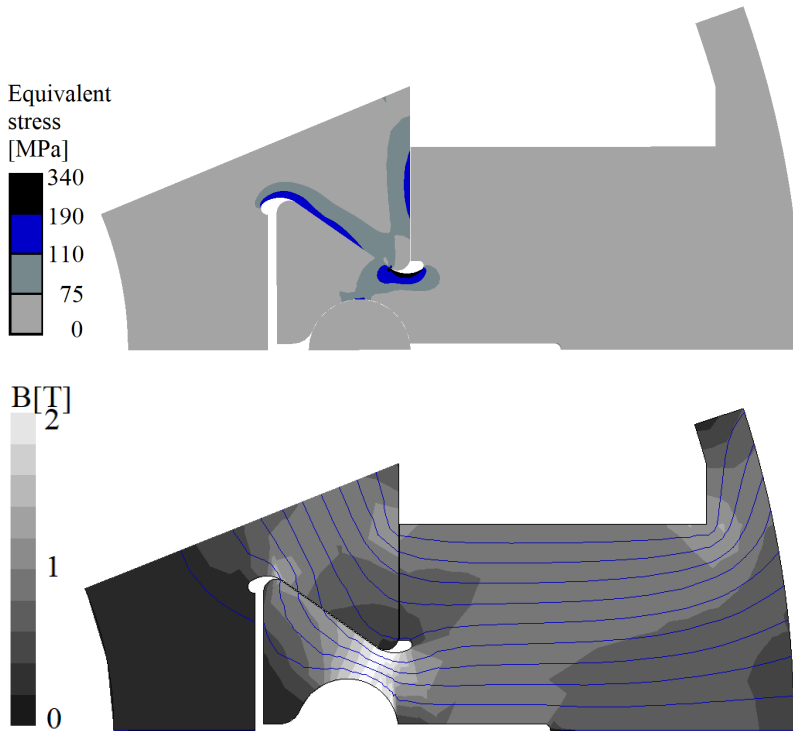
Figure 4.12: Top: Equivalent stress at 44 kN of the final dovetail design (The yield strength of the tooth steel is 460 MPa, of the core steel 400 MPa). Bottom: Magnetic flux (density and lines) at 50 A excitation current.

Figure 4.13: Photo of the tensile test sample with the final design and material choice.

with a friction factor are more accurate, if the friction factor is known. Finding this friction factor is thus an important result of the verification. The most important parameters for the simulations are shown in Table 4.2.

Table 4.2: Mechanical simulation setup for dovetail design process

Laminated steel teeth properties	Young's modulus	190	GPa
	Poisson's ratio	0.292	
	Yield strength	460	MPa
Solid steel core properties	Young's modulus	206	GPa
	Poisson's ratio	0.290	
	Yield strength	400	MPa
Simulation settings	Software	ANSYS Mechanics 2021	
	Formulation	Normal Lagrange	
	Contact type	Frictional, 0.35	
	Mesh @ contact	0.02	mm
	2D Behavior	Plane Stress	

4.5.2   Results and discussion

The force-gap results of the second tensile test are shown in Figure 4.14, along with the results of simulations on the corresponding model.
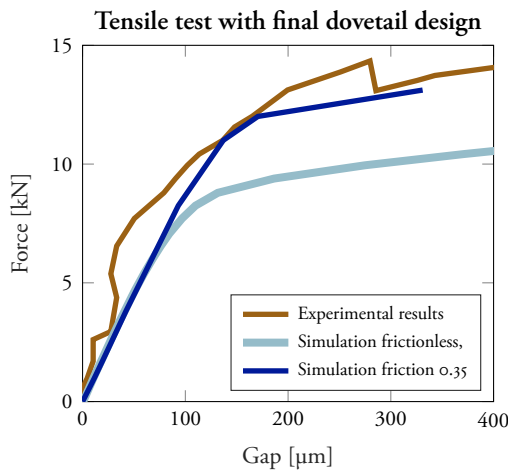

Figure 4.14: Experimental and simulation results of tensile test with final dovetail design. The simulation results are shown with a frictionless contact model and with the friction factor set to 0.35.

The experimental results from the tensile test show some interesting behavior. In the linear region (especially up to just over 5 kN), the measured gap varies as the applied force is increased, while in the nonlinear region the force varies as the gap increases (just before 300 μm). Intuitively this does not seem correct, but this tensile test is, just like the first one, run with a constantly increasing displacement of the whole setup. The total elongation

of the experiment is distributed over the machine grips as well as the test sample itself, while the gap shown in the Figure 4.14 only measure the gap at the tenon base. So in the beginning of the experiment, the test sample and the grips may not yet have stabilized under the increasing load, and when one part stretches or slides slightly, the distribution of the displacement may change so that the gap in the dovetail joint decreases temporarily. Later, in the plastic region, some strain may have built up some friction in the dovetail joint. When that friction is overcome and the joint slides slightly, suddenly less force is needed to produce the requested displacement of the whole set.

None of this is included in the simulation model, but a constant friction coefficient of 0.35 seems to give an accurate representation of the behavior of the sample, excluding movement in the grips of the machine and temporarily higher friction in the joint.

Table 4.3: Simulation results on the final dovetail design for the prototype

Plane	Stress peak at 44 kN	328	MPa
stress	Equivalent speed at yield strength	11 900	rpm
Plane	Stress peak at 44 kN	293	MPa
strain	Equivalent speed at yield strength	12 600	rpm

Simulation results on the final dovetail design as it appears in the prototype with the friction factor settled from the tensile test, are shown in Table 4.3. The simulations are made both with a plane stress assumption and a plane strain assumption, since the axial thickness of the rotor (120 mm) is three times larger than the radial length (39.75 mm from the inner to the outer radius of the rotor) and it is therefore not trivial which assumption is more accurate. Both assumption render roughly the same results, with a 9 % higher peak stress for the plane stress case. This means that, no matter if plane stress or plane strain is the more accurate assumption, the results from the plane stress assumption will ensure the conclusions are on the conservative side.

4.6 Discussion on the dovetail joint design

The results consider only the pure centrifugal force, but no forces from braking or accelerating, from the torque in the airgap, from the magnetic flux or from torque ripple. As an example the maximum torque of the machine (around 140 Nm) equates to 17.5 Nm per pole, which in turn equates to less than 300 N at the radius of the rotor, which is small compared to the centrifugal force at 10 000 rpm (44 kN). Without further studies it can not be ruled out that other forces can add up and pose a threat to the mechanical durability. This also applies to other factors which are not considered such as vibrations, pre-tension

in the material, thermal impacts and expansion, axial stress or strain through 3D analysis, or fatigue from repetition (such as through a stress over number of cycles curve).

No safety factors are considered in the design process. Safety factors varies with the application in a range from around 1.5 to 10. The smallest safety factors are used in aeroplanes despite the fact that a failure is likely to have fatal outcomes. A higher safety factor would mean the aeroplane could not fly at all. Instead frequent ocular inspection is applied, so that any sign of material fatigue can be spotted and parts be replaced. Still a safety factor is needed, because the real material is not as clean as it is modelled in the software. On a microscopic level the material properties may vary throughout the material, so that plastic deformation is reached in a microscopic region, eventually leading to fretting cracks which can then spread.

The maximum speed of the rotor with the proposed dovetail design, without a safety factor, is 11 900 rpm. Any factor higher than 1.19 means the design does not meet the target maximum speed of 10 000 rpm for long term use. In other words, more work is needed on the mechanical analysis, the dovetail design, material choices and choice of safety factor along with a practice for how to ensure failures are avoided (such as ocular inspection or replacement of dovetail poles after a certain time or distance driven).

For the testing of the prototype the risk of a failure at 10 000 rpm is small. However, since they can not be completely ruled out, the experimental test plan should be set up so that high speed tests are run as late in the process as possible. This means that in case of a failure due to the mechanical load from high speeds, as much as possible of the test plan is already completed.

Another aspect of dovetail jointed rotor teeth is that they make the rotor modular and should enable easier exchange of rotor coils in case of failure so that the machine can continue to service after such a failure. At the end of the lifetime of the machine, the materials of the rotor should also be easier to recycle. This approach is not in the scope of the work, but can still be seen as an outcome of the work that has been done on dovetail joints.

4.7   Chapter summary

This chapter deals with the design of the dovetail joints for an EE rotor with hollow direct air cooled windings. Works from literature are presented that have some relevance for dovetail joints in the rotor of electric machines for EV traction. The most important lessons are how fractures are initiated in the dovetails of generator turbine blades, stressing that no regions in the joint should go into plastic deformation repeatedly.

In order to understand the behavior of dovetail joints under load, as well as the simulations

of such contact problems, an initial tensile test with a simple aluminium dovetail joint is made and compared to simulation results. Using the experience from this study, adapting for constraints set by the manufacturing method and minimizing the effects on the electromagnetic performance, the design of the final dovetail joint is determined.

An experimental tensile test is made with the final design of the dovetail, with the correct materials. This test is used to validate the design and to calibrate the friction factor of the simulations, so that a better prediction of the maximum speed limit of the rotor can be made with respect to the centrifugal forces. It is concluded that the joint can handle the forces from the testing of the prototype. However more work is needed to set a definite maximum speed taking a safety factor into account for more long term use than laboratory testing.

This concludes the design work of the rotor and thus the process moves on to the manufacturing and assembly of the prototype.
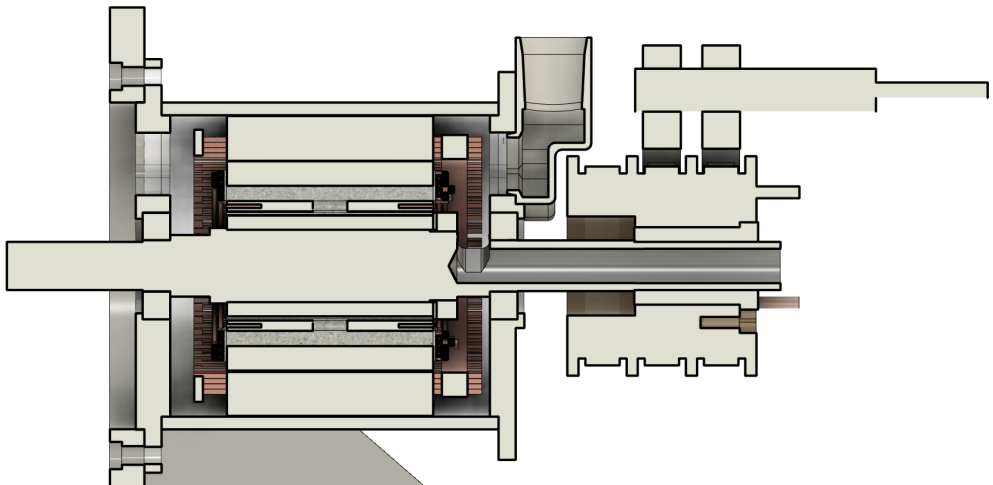
Chapter 5

Prototype machine preparation

Figure 5.1: Previous page: Radial cross section of the 3D CAD model of the prototype.

5.1   Chapter overview

The previous chapter concludes that there is a design for the dovetail rotor which should enable the assembly of the rotor with hollow direct cooled prewound coils, without compromising the electromagnetic or mechanical performance of the machine too much. This chapter deals with the manufacturing and assembly of that rotor and the rest of the prototype, focusing on the following:

- The manufacturing of the dovetail rotor parts, including finding a satisfying machining method for the dovetail joint and ensuring the joints fits in practice.

- The assembly of the rotor, including ensuring the final assembled rotor following the measurements of the CAD drawings, as well as the inter-connection of the coils.

- The design of the housing and shaft, with a focus on enclosing and leading the coolant air flow, as well as the current supply.

Most sections and subsections of this chapter begin with a CAD cross section highlighting the parts relevant for the section whenever applicable. These figures are not necessarily referenced in the text.

5.2   Stator

The stator was provided by Chalmers University. The core is made by laminated steel (B35AV1900) and the windings are copper hairpins. The exact design and dimensions are protected by immaterial rights and are not published in this thesis. For this reason, the figures of the stator design shown in this chapter are only similar representations of the actual stator design.

5.3   Rotor manufacturing

The rotor design process is described in the previous chapters. This section describes the process of manufacturing and assembling the prototype rotor.
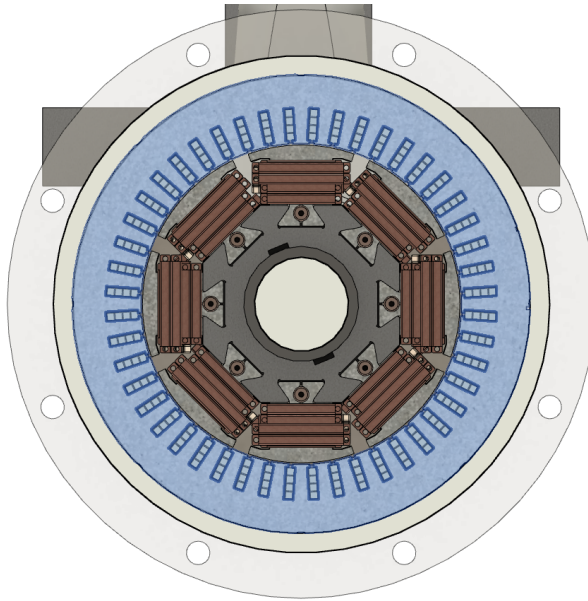
Figure 5.2: Transverse cross section of the 3D CAD model of the prototype, highlighting the stator in blue.
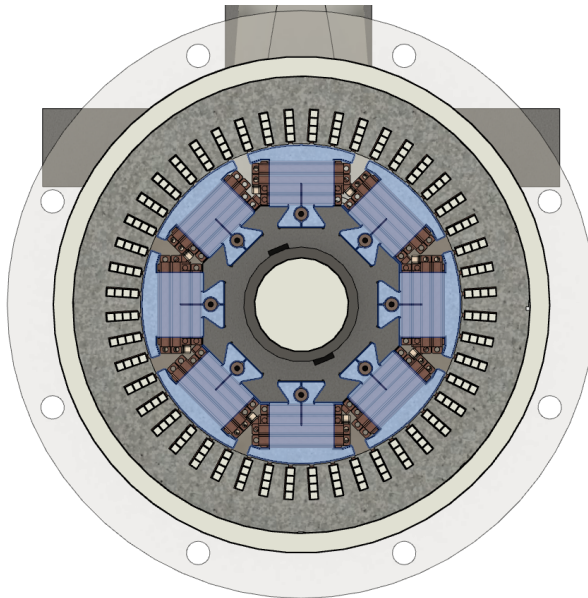
5.3.1 Rotor teeth manufacturing

Figure 5.3: Transverse cross section of the 3D CAD model of the prototype, highlighting the rotor teeth in blue.

The rotor teeth are made from laminated steel (M35-35A from Tata Steel, previously known

as Cogent SURA). For this custom order the steel sheets were laser cut individually, stacked and then bonded with backlack varnish by Bevi AB into 120 mm teeth. However, stacking and and bonding the teeth like this will result in a quite uneven surface, where the steel sheets may protrude as much as 0.1 mm from the average surface. This is not good enough for the dovetail joints for at least two reasons. First, it will be very difficult to slide the dovetail joint parts into each other when the sheets are not perfectly aligned, especially as the sheets are perpendicular to the joint. Second, when the individual sheets protrude unevenly only the most protruding sheets will have physical contact with the core part in the dovetail joint and thus take more of the stress. For this reason, the laminated teeth were made with a rectangular bottom. To the left in Figure 5.4, a photo of a stacked laminated tooth is shown (but shorter than the final 120 mm teeth). To the right, the laser cut shape is shown in blue, from which the bronze colored dovetail shape is to be cut with another method. 11 teeth were delivered from Bevi in order to have some spares when practicing the dovetail joint machining (the prototype machine has 8 poles). One of these teeth broke in half before further machining was made, a part of which is seen in the photo.
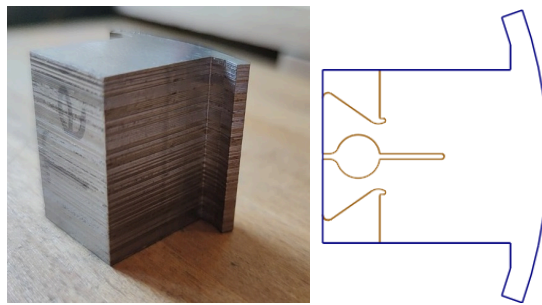

Figure 5.4: Left: Photo of a bonded rotor tooth before wire discharge cutting. The sample in the photo is shorter than the 120 mm tooth for the prototype. Right: Sketch of the shape the bonded tooth in blue, and the shape to be cut with electric discharge machining in bronze.

The dovetail shape was then cut from the rectangular shaped bottom of the teeth with electric discharge machining. It was not known to the author if laminated steel could in fact be cut with electric discharge machining, since this cutting method relies on a large current being conducted through the material and the laminated sheets are electrically isolated from each other. However, a test at the Linova workshop outside Malmö proved it is possible. The rest of the machining was then performed by the Arema Mekano workshop in Malmö. The machining was made in three iterations, resulting in a finer and finer surface and after the final cut all deviations are smaller than 0.01 mm.

After this work was performed a study has been published confirming electric discharge machining can be used to cut electrical steel along with investigating the electromagnetic effects on the material [84].

While learning how to cut the laminated steel, one test sample was first cut in one of the

Figure 5.5: Left: Photo of the first test of electric discharge machining of shorter tooth sample. Right: Photo of the first electric discharge machined tooth of full length, fit into a test sample of a mortise joint in aluminium.

shorter parts of the eleventh tooth (see left in Figure 5.5), then nine in the final teeth. To test the tolerances of the joint, a test mortise joint was cut in an aluminium block and fit to the first cut entire tooth (see right in Figure 5.5). Because of the fixture holding the teeth while machining, the depth of the slit for the wedge was reduced from 18 mm to 12 mm. The wedges have a threaded hole in the thicker end, so that a screw can be put there to help pull the wedge out for when the joint is to be disassembled (Such a screw can be be seen in the wedge in Figure 5.5).

5.3.2 Rotor yoke manufacturing

The material of the rotor yoke is solid SS-1672 Steel (internationally known as C45E). After ensuring the tolerances work, the dovetail joints were cut with electric discharge machining.

With the rotor yoke and all the rotor teeth machined, the dovetail joints proved to function very well. The teeth could be slid easily in and out of the joints, and with no play in the fit whatsoever. With a wedge on the dovetail tenon hole, the joint is then impossible to slide out with hand force and when the wedge is removed the tooth is easily slid off again.

5.3.3 Rotor coils

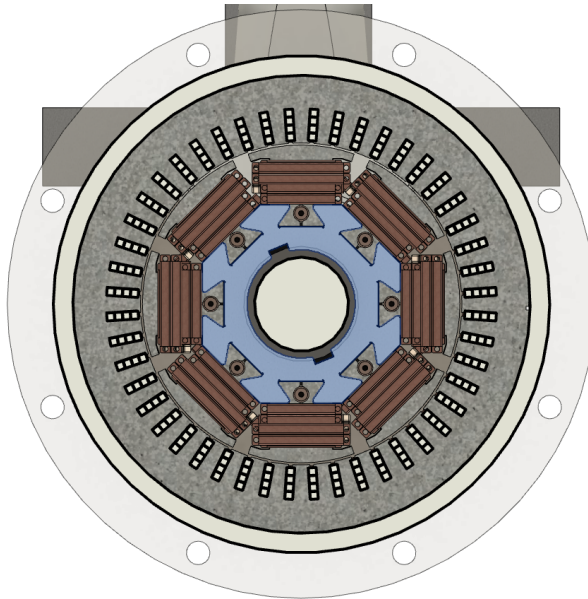The procedure of winding the prewound coils is presented in Chapter 3.4.2.

Figure 5.6: Transverse cross section of the 3D CAD model of the prototype, highlighting the rotor yoke in blue.

5.3.4 Insulation paper

Nomex 410 was used for insulation between the coils and the rotor core. Ordinary printer paper was used to test out the suitable dimensions of the insulation paper before cutting, folding and installing it on all none rotor teeth with coils. See Figure 5.9 for photos of the test with printer paper and then all teeth with Nomex 410 in place.

5.4 Rotor assembly

With the coils folded and installed on the teeth along with the insulation papers, the teeth could be assembled on the yoke. This test was successful in the way that all the teeth could be slid into the dovetail joints, with just a little bit of fiddling to make room for the coils and the insulation paper. The assembled rotor is shown in Figure 5.11. Note the rust on the outside of the rotor teeth, which occurred on the laminated steel after the electric discharge machining during which the teeth are submerged in water.
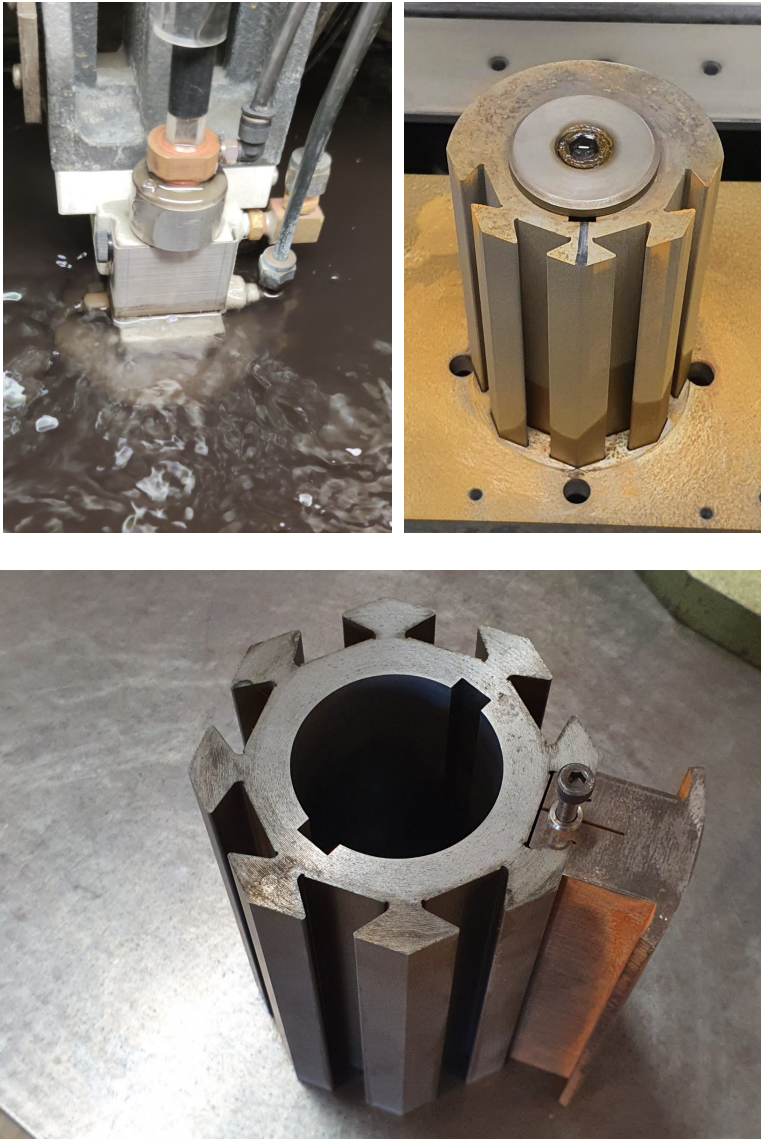
Figure 5.7: Top left: Rotor yoke submerged in water while electric discharge machining is performed. Top right: Rotor yoke after having four dovetail joints machined. Bottom: Finished rotor yoke with one tooth in place.

5.4.1 Post-assembly processing

However, the assembly was not successful in one important way. The rotor did not fit inside the stator. Somehow the diameter of the rotor had become too large. The radius of the rotor should be 60.75 mm (leaving 0.75 mm airgap to the stator with an inner radius of 61.50 mm) according to the design, but the actual radius of each teeth varied between 60.75
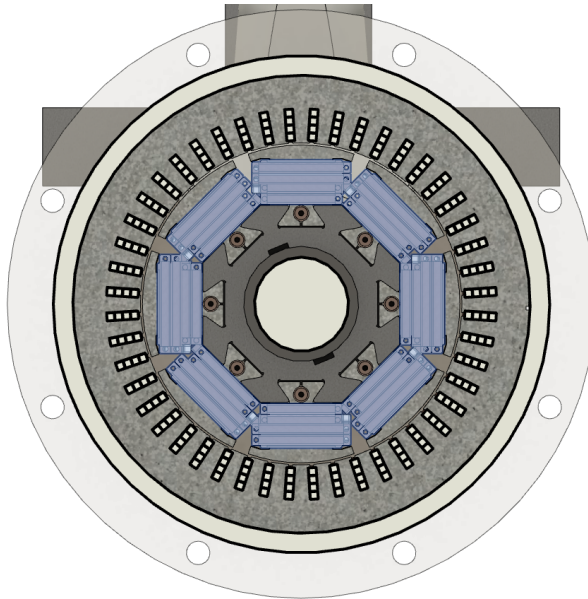
Figure 5.8: Transverse cross section of the 3D CAD model of the prototype, highlighting the coils in the blue.
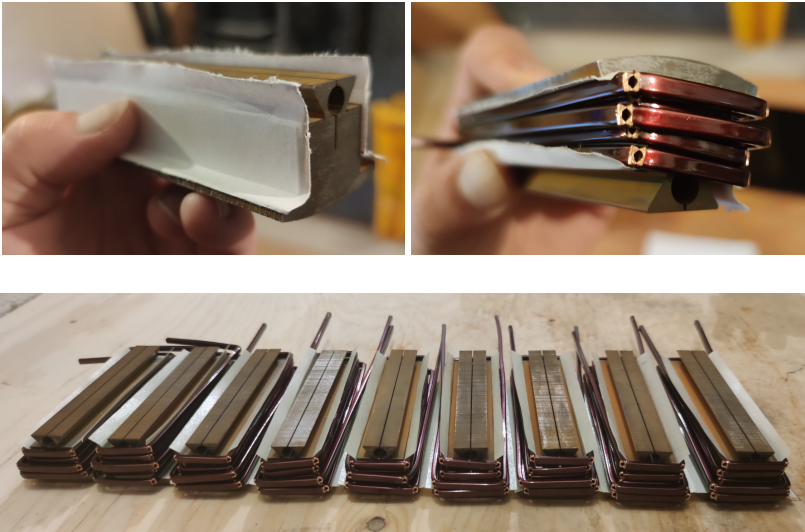

Figure 5.9: Top left and right: Insulation paper fitting test with ordinary printer paper, with and without coil. Bottom: Nine teeth with coils and Nomex 410 insulation paper.

and 61.80 mm which means several teeth had a radius larger than the stator inner radius. The radii of the teeth in both ends of the rotor were measured from the inside of the shaft hole, adding the shaft radius to the measurement, with the results shown in Table 5.1.

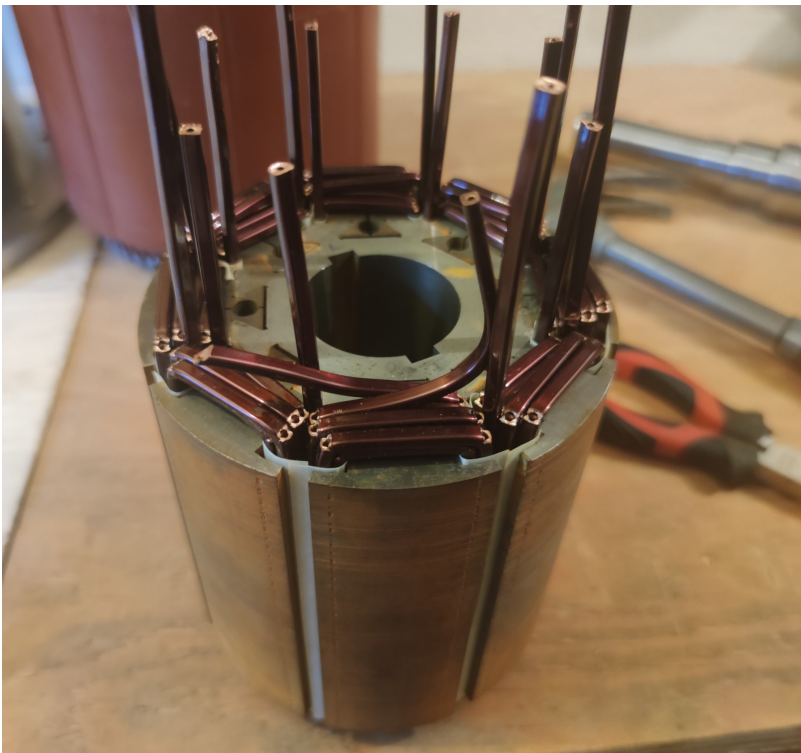**Figure 5.10:** All motor parts ready to be assembled.

Figure 5.11: All eight rotor teeth assembled on the yoke.

Most of the poles have one radius that is clearly larger than the other, marked with bold font. It seems in the process of working out the tolerances and details around wire discharging the tooth, the target radius was omitted. The aim was to make the dovetail joints work with tight precision, which resulted in all measured radii but one being larger than the target 60.75 mm.

Table 5.1: Radius of teeth after first assembly

Tooth number	Radius side a [mm]	Radius side b [mm]
1	60.98	**61.60**
2	**61.80**	60.98
3	61.15	61.10
4	61.10	**61.65**
5	**61.50**	61.00
6	60.95	**61.20**
5	**61.40**	60.75
6	**61.40**	60.95

The fact that most of the teeth has one radius that is clearly larger than the other side has another cause. The dovetail joints of this side with the larger radius looks clearly different (See Figure 5.12. This is most noticeable from the width of the airgap at the bottom of the dovetail joint, pointed out in the figure, but is also seen at the amount of material beneath the wedge hole, and radius at the edges of the tenon joint. The fixture and the cutting program of the electric discharge machining of the rotor teeth assumes the bonded uncut teeth are straight. This assumption, intuitive as it may seem, proved to be untrue.

One theory of the cause if this is that the bonded teeth may be slightly twisted or curved, as shown in an exaggerated sketch to the right in Figure 5.13. When the electric discharge machine cuts along a vertical line (like the dashed bronze lines) the shape cut from the original piece will be different at the bottom and the top of the tooth, as illustrated in Figure 5.14 and shown in Figure 5.15.

To solve this, the rotor was re-assembled without the coils, and with all teeth facing the same way (having all the correct dovetails on the same side), and the wedges inserted in the holes on one side, before leaving it to HMJ Centerless Slip AB to grind the whole rotor to the right radius of 60.75 mm. After this the rotor fit perfectly in the stator, see Figure 5.16.

The procedure of assembling the rotor with the wedges, grinding it to the right radius, then removing the wedges and disassembling the rotor again resulted in most teeth braking in at least one place. On top of this, some of the outermost laminate sheets were damaged or fell off from its tooth. At this point, it was important to keep all the tooth parts in the correct order, since the holes for the wedges are slightly conically shaped. The teeth from this step in the process are shown in Figure 5.17. It can be noted that the rust remains only on one end of two teeth, while the rest was removed in the grinding process.

The coils were put back on the teeth, including the insulating paper, and slid back onto the
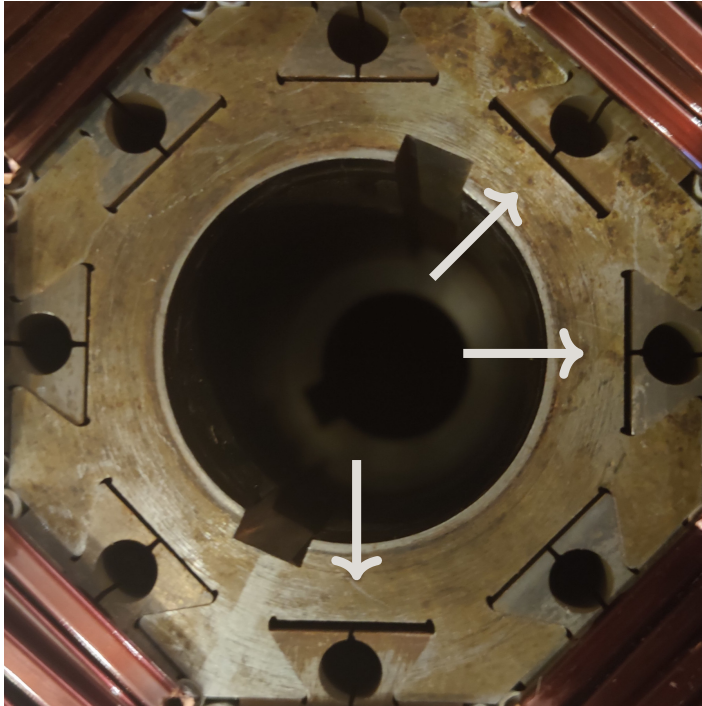
Figure 5.12: All eight rotor teeth assembled on the yoke, zoomed in on the dovetail joints. Three dovetails are different (pointed out with arrows).
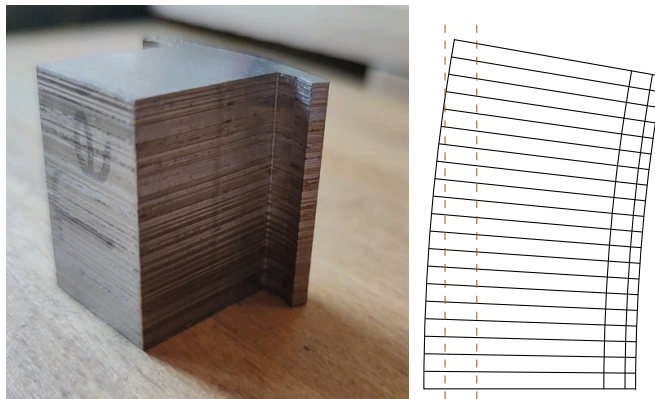

Figure 5.13: Left: Bonded rotor tooth before electric discharge machining. Right: Sketch of exaggerated twisted tooth and straight vertical cutting line (for comparison).

rotor yoke. The broken teeth proved to be just a minor inconvenience, as it was slightly difficult to perfectly align the pieces when sliding the teeth back on, but once a whole tooth was in place with the wedges in place, it was hard to make out where the breaks were. Since some teeth had lost up to three lamination sheets, the length of the teeth varies between
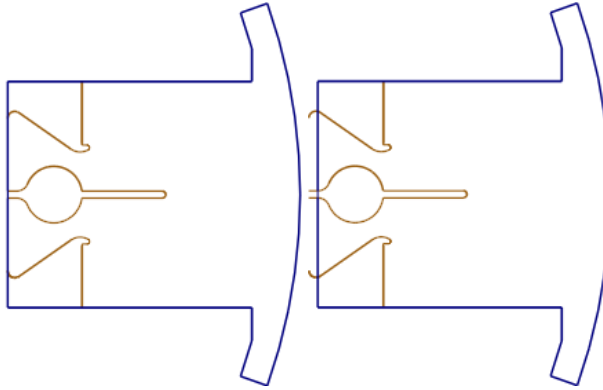
Figure 5.14: Sketch of how the shape cut by the electric discharge machine at the bottom of a twisted tooth (left) and the top of a twisted tooth (right), marked in bronze color. Original shape of bonded tooth in blue.
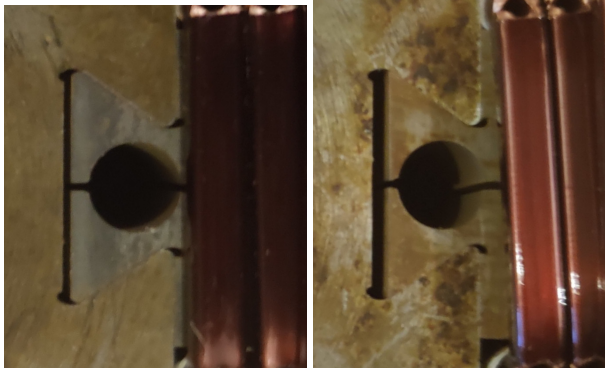

Figure 5.15: Left: Photo of a correctly cut end of a dovetail joint. Right: Photo of a distorted dovetail joint.

119 and 120 mm in the final assembly.

5.4.2   Rotor coil interconnections

The next step was interconnecting the coils so that the current could flow in series through all poles. In order to get the polarities of the rotor poles right, the interconnection has to be done the right way. If every other coil have their ends folded outwards, and the rest folded inwards, this is achieved.

A test was made to interconnect the coils with cable ties, but they could not be tightened hard enough to not risk poor contact between the conductors. Another attempt was made with soldering, but the mass and heat conductivity of copper made it impossible to heat up the material enough to melt the solder wire with the soldering iron available in the basement of the author's home. In industry, copper conductors can be welded together,

Figure 5.16: Grinded rotor inside stator.

Figure 5.17: Broken rotor teeth before remounting. All pieces are marked and kept in the right order and direction.

but that equipment was not available in the basement either. The remaining option, screw
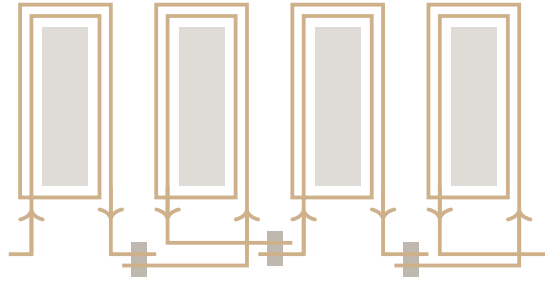
Figure 5.18: Sketch of interconnection of four poles, where the ends of every other coil folds outwards (first and third from the left), and every other coil folds inwards (second and fourth), in order to get the polarity right.
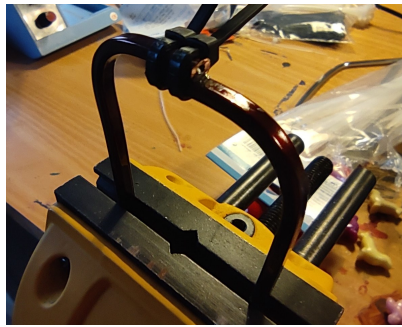
terminals, was thus chosen.

Figure 5.19: Coil interconnection test with two cable ties.

Normally, screw terminals are used so that the current is lead through the screws and the busbar to the next conductor, as shown in the left of Figure 5.20. Since the prototype is designed for high current this approach is not appropriate. Instead, the insulating varnish is grinded off from one side of the ends of each conductor, so that the conductors can overlap and have direct contact. Therefore, a larger size of screw terminals has to be used (see the right in the same figure), which will make the rotor winding end turns bulkier, but they can still be kept within the end turn length of the stator hairpin windings.
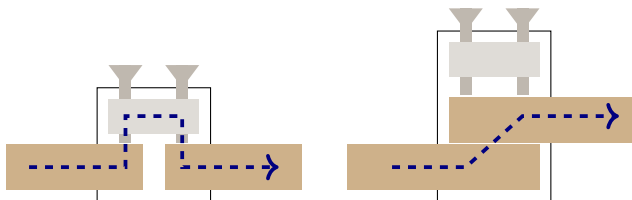

Figure 5.20: Current flow in different ways to use screw terminals (Left: Normal use of screw terminals. Right: The way screw terminals are used in the prototype, to ensure better electric contact.)

The procedure of connecting the screw terminals is shown in Figure 5.20. The ends of

each coil need to be bent in place on the rotor, which includes sawing a groove manually to *roughly* 0.7 mm before bending, in order for the inlets for the air to be opened during bending. The varnish on the ends of the conductors also needs to be grinded off on one side to enable contact between the terminals. An extra bend is often needed on the conductors to align them before they are sufficiently straightened to fit into the screw terminal as they are re-bent. This process is critical, as bending back and forth too many times will cause the end of the conductor breaking off completely, thus ruining the whole coil. The procedure was successful for the prototype, but this step in the process is still deemed as a huge drawback of the coil bending method chosen, as the difficulty of this step was not anticipated.

Finally, the connections to the supply cables coming through the shaft were made. The space for all the screw terminals, radially inside the stator end turns and axially inside the end plates, was cramped just barely enough to fit everything. Four supply cables were used inside the shaft, two in parallel for the current into the rotor, and two in parallel for the current out of the rotor, since this allows a higher total amount of copper than if just two slightly larger cables were used.
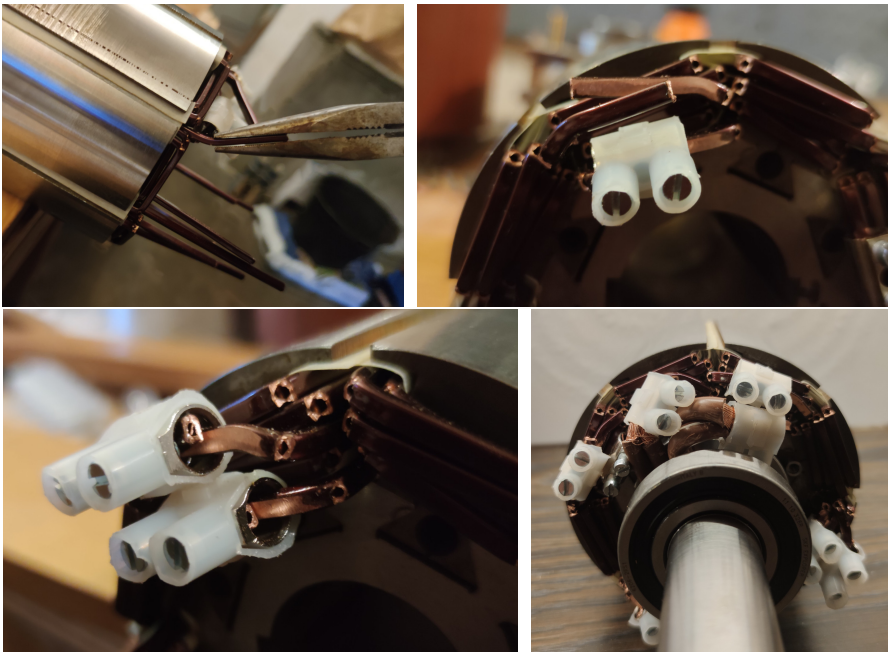

Figure 5.21: Procedure of applying screw terminals. Top left: Bending the ends. Top right: aligning the grinded ends. Bottom Fitting the ends into a screw terminal. Bottom right: All the screw terminals in place, including the connections to the supply cables through the shaft.
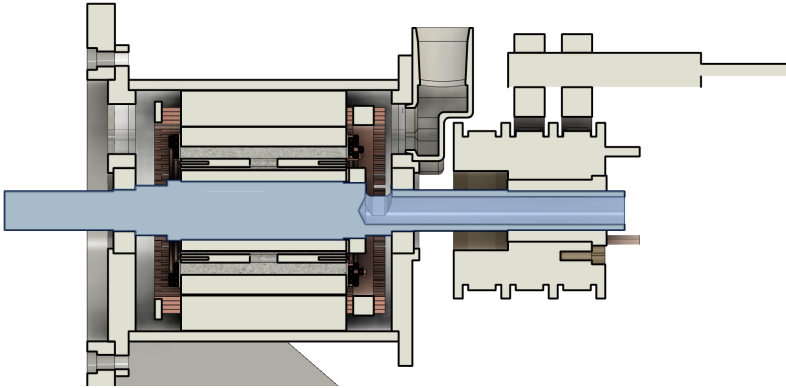
Figure 5.22: Radial cross section of the 3D CAD model of the prototype, highlighting the shaft in blue.

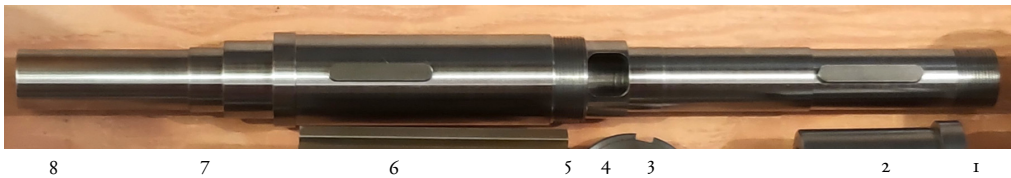## 5.5 Shaft design and manufacturing

Figure 5.23: Photo of the shaft. The numbers from the right mark the position of: 1. Threads for the slipring clip nut. 2. slipring socket with wedge in place. 3. Cool side bearing. 4. Hole to the inner cable channel. 5. Threads for the rotor yoke clip nut. 6. Rotor yoke with wedge in place. 7. Hot side bearing. 8. Shaft end to connect to the other machine on the test bench.

The shaft is shown in Figure 5.23, with numbers marking the axial position of different parts explained in the caption. The shaft was lathed out of a 45 mm diameter 2172 grade steel cylinder. The cool side of the shaft was made hollow for the current supply cables. Wedge slots were milled for the rotor core and the slipring socket. One side of the rotor core has an increased radius of 1 mm to stop the rotor core from axial movement in that direction. On the other side, the shaft is threaded for a lock-nut which will stop the rotor from axial movement in the other direction. The shaft is also threaded at the slipring end, for a lock-nut to keep the slipring socket from axial movement. The radius of the shaft is gradually increased in steps from both ends, so that all parts can be assembled without unnecessary friction and then obtain its correct position as it stops when the radius gets too large. All outer parts were heated up to around 80 °C before being assembled on the cool shaft.

5.6 Housing and mounting preparations

5.6.1 House
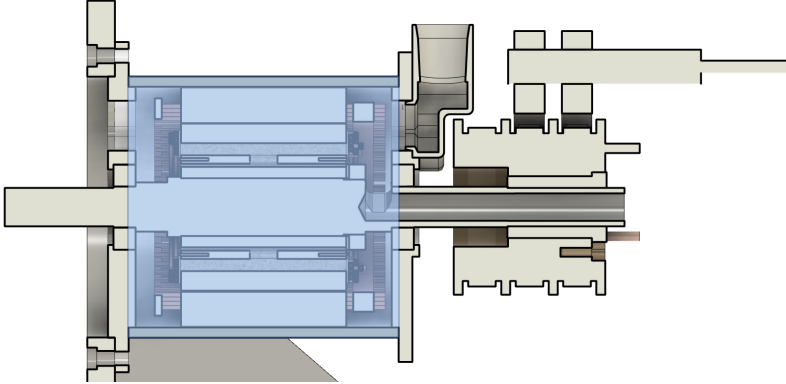
Figure 5.24: Radial cross section of the 3D CAD model of the prototype, highlighting the housing in blue.

The house was designed to be produced from a standard D190x10 aluminium pipe (with an outer diameter of 190 mm and a thickness of 10 mm), only needed to be cut to the right length and lathed on the inside to fit the stator diameter, see Figure 5.25.

Figure 5.25: Photo of the housing.

Since the tests were planned to be run in no-load modes with little torque involved, the stator was only thermally fit inside the housing by heating the housing in a kitchen oven and sliding the stator in and measured into the correct position. If torque tests are to be run on the machine in the future, a hole can be drilled in the housing that is then threaded on the inside to fit a bolt which will stop the stator from rotating inside the housing.

5.6.2   End plates
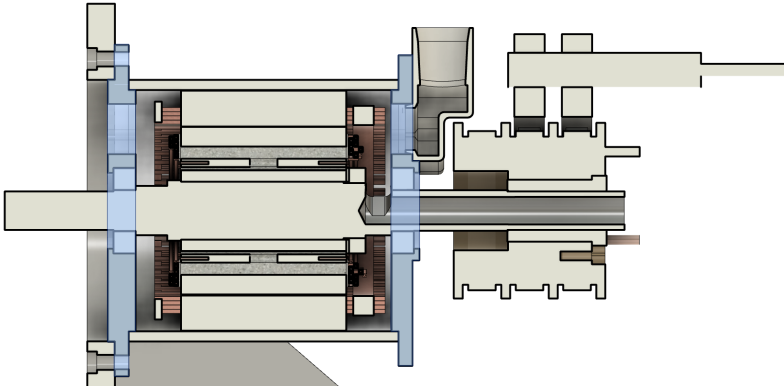
Figure 5.26: Radial cross section of the 3D CAD model of the prototype, highlighting the end plates in blue.

The two end plates are similar but not identical, as can be seen in Figure 5.27. Both has a rim that fits inside the inner diameter of the house, slots for the bearings and a 90 degree opening window for air inlet and outlet. The window is also aligned with the end terminals of the stator windings.
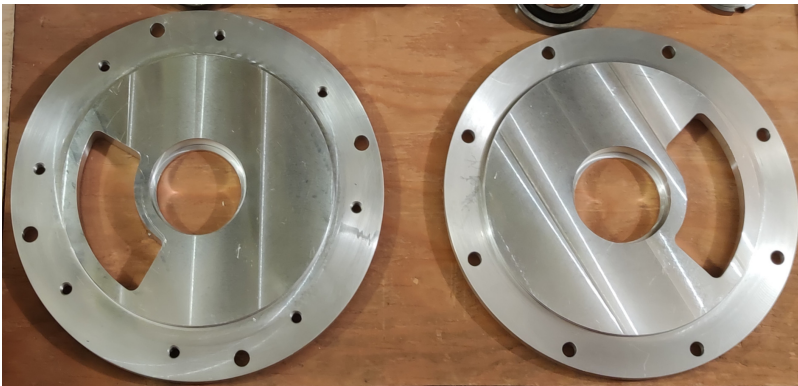
Figure 5.27: Photo of the end plates, with the hit side end plate to the left and cool side end plate to the right.

The end plates are held together with eight M8 threaded rods when mounted. For this reason the end plates have eight holes outside of the house rim. On the cool side, the holes have a diameter of 8 mm, so that the M8 rods fit though the holes and can then be tightened with nuts. The hot side end plate has an extra rim on the outside for mounting it to the motor bench, which does not leave any room for nuts for the M8 rods. For this reason, the M8 rod holes on the hot side are threaded so that the rods will be screwed directly into the end plates on this side, see Figure 5.28. The hot side end plate also has a larger diameter to fit four extra holes for the bolts that will fasten it to the mount.
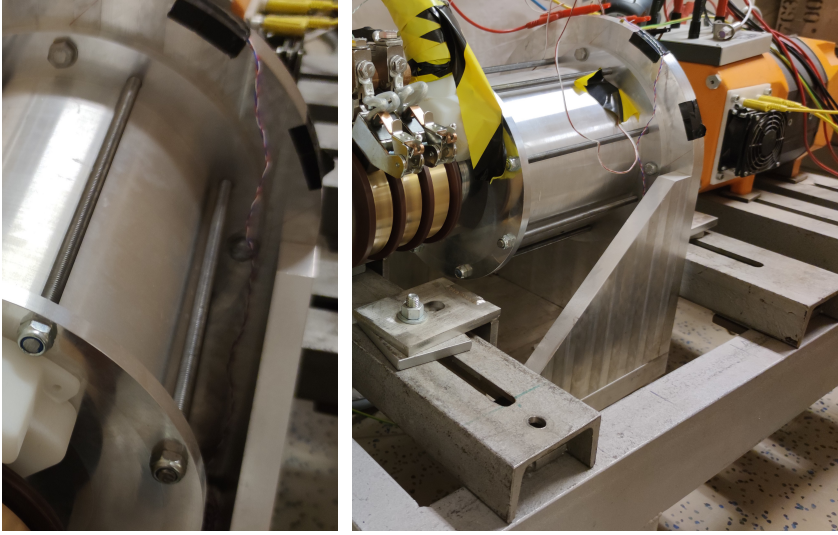
Figure 5.28: Left: Photo of the threaded rods holding the end plates together. On the far side, the holes in the end plate is threaded on the inside, while the close side tighten the rod with nuts. Right: Photo of the prototype fixed into the motor bench.

5.6.3 Bearings
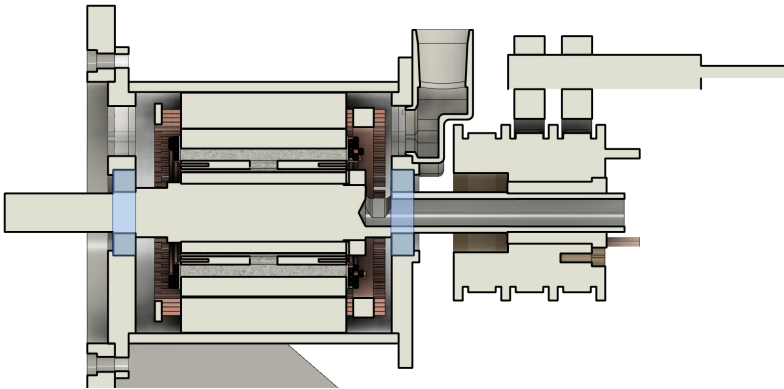
Figure 5.29: Radial section of the 3D CAD model of the prototype, highlighting the bearings in blue.

The bearings used are 6206 from SKF. On the cool side the bearings are sealed to enclose the coolant air on the pressurized side (SKF 6206 2RS1, with a maximum speed of 8000 rpm). On the hot side the bearings are open (SKF 6206, with a maximum speed of 15 000 rpm). The maximum speed of the cool side bearing means the prototype will not be able to run at the target speed of 10 000 rpm in this implementation.
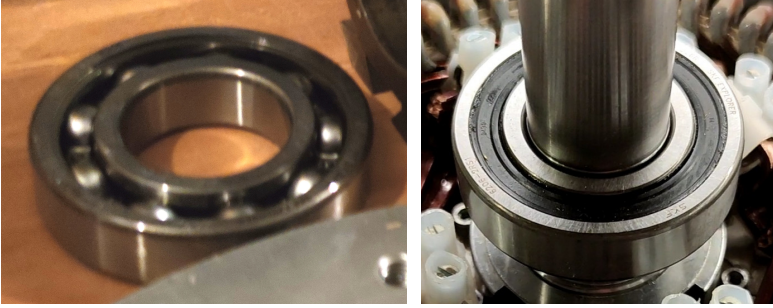
Figure 5.30: Photo of the bearings. To the left: Hot side bearings without sealing. To the right: Cool side bearings with sealing.

5.6.4   Motor bench mount

The motor bench has a mount used for the machine previously tested in the bench, for which the design of one of the housing end plates was adapted, so that the prototype could be properly mounted into the bench. The design of the mount and the hot side end plate ensures that the prototype is centered and aligned with the shaft of the other machine on the motor bench.
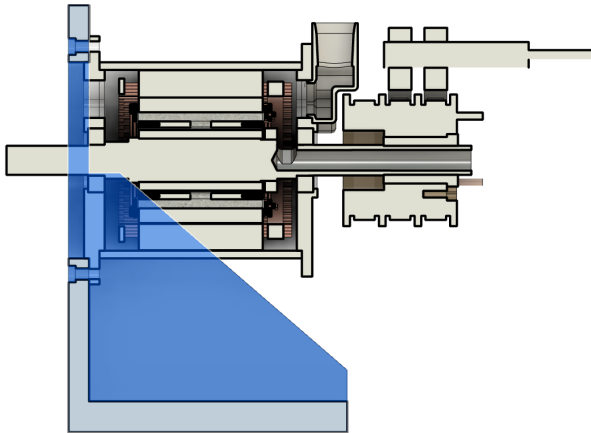
Figure 5.31: Radial cross section of the 3D CAD model of the prototype, highlighting the mount in blue.

5.7   Auxiliary system parts

To complete the design of the prototype systems some auxiliary system parts are needed. These parts involve the current supply and air supply.

5.7.1   Current supply and sliprings

To lead the current onto the rotations rotor, sliprings are used. The option to use a rotating transformer was rejected for the prototype due to the magnitude of the current per turn in the rotor. Conventionally, an EE rotor has thinner conductors, more turns and fewer amperes per turn, which leads to lower losses in the rectifying diodes needed on the rotating side of the rotating transformer. It may be possible to design and manufacture a custom made rotating transformer for the prototype, which could handle the rectifying H-bridge losses, but in order to save time for all the other parts that were designed and manufactured for the prototype, a slipring solution was chosen.
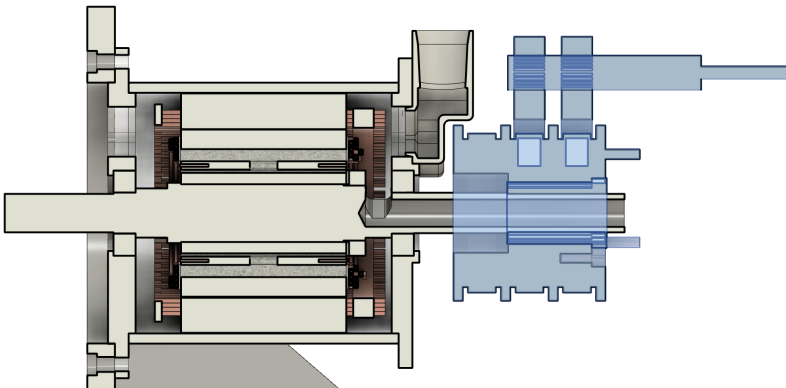
Figure 5.32: Radial cross section of the 3D CAD model of the prototype, highlighting the sliprings and brushes in blue.

Carbex provided the sliprings used for the prototype. The brushes can supply a current of 72 A each, which means a total excitation of 144 A is possible with two brushes. The maximum peripheral speed allowed by the silver graphite brushes (CA50) is 40 m/s, which means around 6600 rpm with a ring diameter of 115 mm.

The sliprings used are not optimized for the prototype or high speed applications, which is seen on the large diameter (smaller diameter means lower peripheral speed, which means that a higher rotational speed can be allowed) and the three rings (while only two are needed). The limitations in speed and current means that extensive testing at high currents and speeds with the prototype comes with a risk for failures. Note that the sliprings used are thus bulkier than necessary, and not representative for the size of sliprings optimized for the application.

In order to install the sliprings on the shaft a socket is needed (seen to the top left in Figure 5.33 and inside the slipring frame in the CAD cross section in Figure 5.32), since the slipring frame has a larger inner diameter than the shaft. This socket is designed to protrude the slipring frame so when a clip nut is put on the threaded end of the shaft, it will tighten against the socket alone. A bolt with washers is then used to keep the slipring
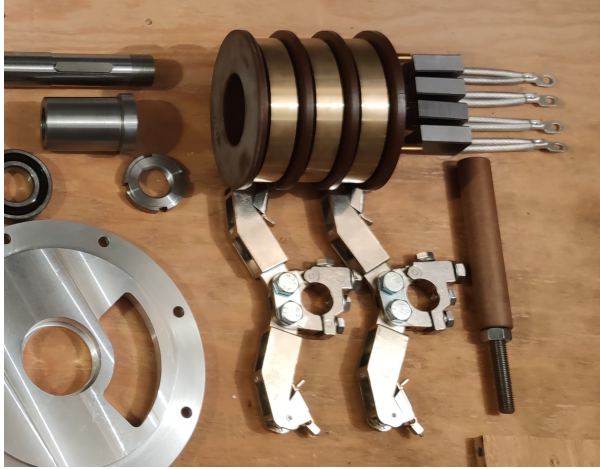
Figure 5.33: Photo of the parts for the sliprings before installation. To the top left the socket for the sliprings and the clip nut is seen. In the top center is the sip ring frame, and in the bottom center are the two brush holders. To the top right are the four silver graphite brushes and to the bottom right is the socket for the brush holders.

frame in place on the socket. Between the socket and the shaft there is wedge to keep the socket from rotating on the shaft. In order so secure the slipring frame from rotating on the socket, another arrangement with the clip nut could be made. However, this is not realized on the prototype, as the torque is assumed to be sufficiently small for the bolt with washers to secure the frame from rotating.
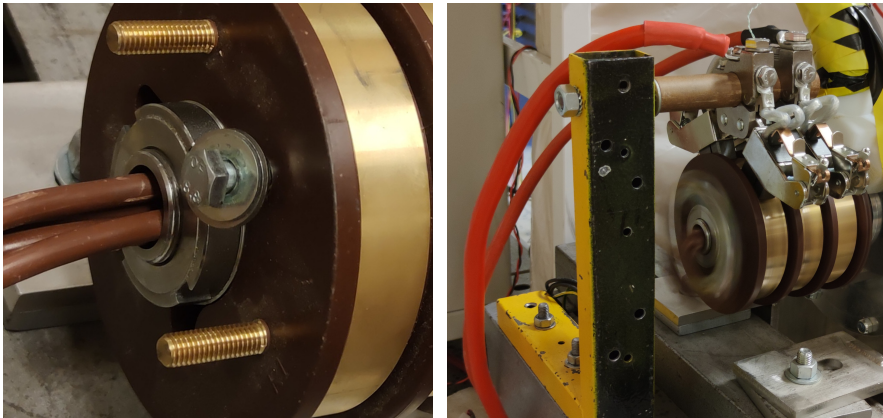
Figure 5.34: Left: Photo of the clip nut fastening the sliprings to the shaft. Right: Photo of the sliprings and brushes installed and rotating.

Figure 5.34 (left) also shows the cables inside the hollow shaft, which were connected to the protruding copper bolts shown in the same photo. This connects the slip rings to the cables and on to the windings inside the housing via the hollow shaft.

5.7.2 Air supply

A Kärcher WD 6P Premium and a Biltema WD1202/30 is used for the air supply. Both are vacuum cleaners with the ability to blow the air instead of sucking. The Kärcher also has the ability to adjust the power, while the Biltema uses a constant (but higher) power.

The air supply units were used in turn and not at the same time. A vacuum cleaner provides a pressure drop with respect to the surrounding air, which is what produces the air flow. Thus, coupling the air supply units in parallel does not increase the total airflow, just as coupling two batteries in parallel over the same load does not increase the total current or power (Note: Ohm's law analogy).

Cooling channels

The main idea was to only let the coolant pass through the hollow windings, and not through the airgap or the prism shaped gaps between the coil (The term prisms is taken from [50]). This could be achieved by filling the prisms with some paste able to withstand expected temperatures and covering the airgap with a silicon ring able to slide along the stator teeth edges as the rotor rotates, or by installing a sleeve that covers the airgap and stator end turns, much like the oil-cooled stator in this work [45] (but covering the stator instead of the rotor).

However, the choice was made to leave all the channels unobstructed, allowing the air to pass through the hollow windings, the prism gaps between the coils and the airgap. The reason for this choice was a combination of time-saving (to maximize the time in the lab before being closed for renovation) and an opportunity to learn about the cooling effect in the other air channels with the help of a further developed thermal cooling model. If the cooling effects when allowing the air through all channels are not satisfactory, then barriers can be added later, easier than it would be to remove the channel barriers (especially the paste in the prisms) later in order to test that option as well.

Air inlet nozzle

To steer the air flow from the hose from the air supply into the cool side inside the housing, a special nozzle was designed and 3D printed in nylon plastic (using Selective Laser Sintering), shown in Figure 5.36. It is designed to fit as tightly as possible between the cool end plate and the sliprings, always maintaining at least 800 mm^2 internal cross section area for the air flow (roughly corresponding to the total cross section area of all the cooling channels, as well as that of the hoses and pipes used for the air supply). Figure 5.37 shows the nozzle installed on the prototype.
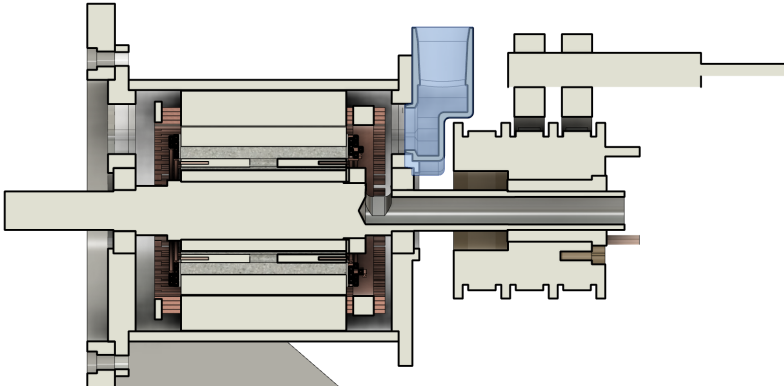
Figure 5.35: Radial cross section of the 3D CAD model of the prototype, highlighting the air inlet nozzle in blue.
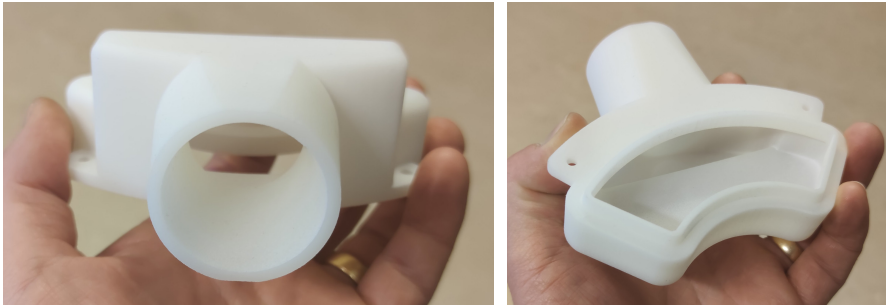
Figure 5.36: Photos of the air inlet nozzle.

5.8   Final prototype assembly

The prototype was assembled, mounted on the motor bench and made ready for testing. Apart from the complications mentioned earlier in the chapter, the assembly went according to plan. Figure 5.38 shows all the parts before assembly, Figure 5.39 shows the machine inside the housing without one end plate. After the assembly the machine was ready for testing.

5.9   Chapter summary

This chapter focuses the work around the manufacturing and assembly of the prototype, making it ready for testing in order to verify the concept of hollow direct air cooled windings for an EE rotor.

It has been concluded that electric discharge machining can be used to machine the dovetail joints, not only in the solid rotor yoke but also in the laminated teeth, with a final
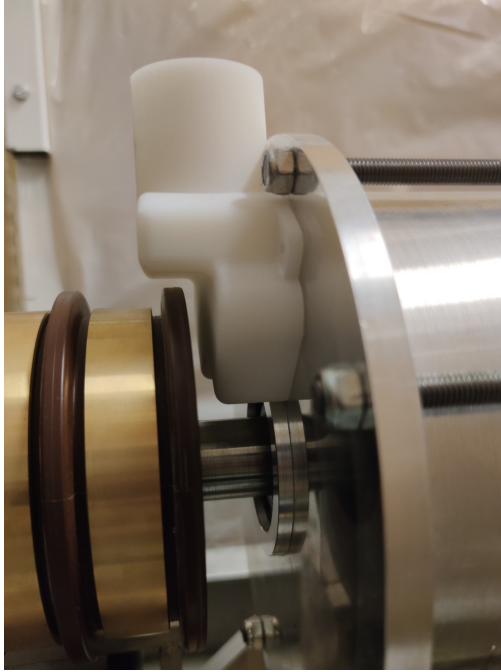
Figure 5.37: The air inlet nozzle installed on the prototype.

Figure 5.38: Photo of all the parts for the prototype before assembly.

tolerance meaning no deviations in the surface are larger than 0.01 mm. The tolerances mean the joints do fit in practice with no visible or mechanically noticeable play. Due to imperfections in the laminated teeth before electric discharge machining, the radii of the rotor poles were too large after assembly so the rotor had to be ground to the correct radius to fit inside the stator with the correct airgap length. With this, it can be concluded that
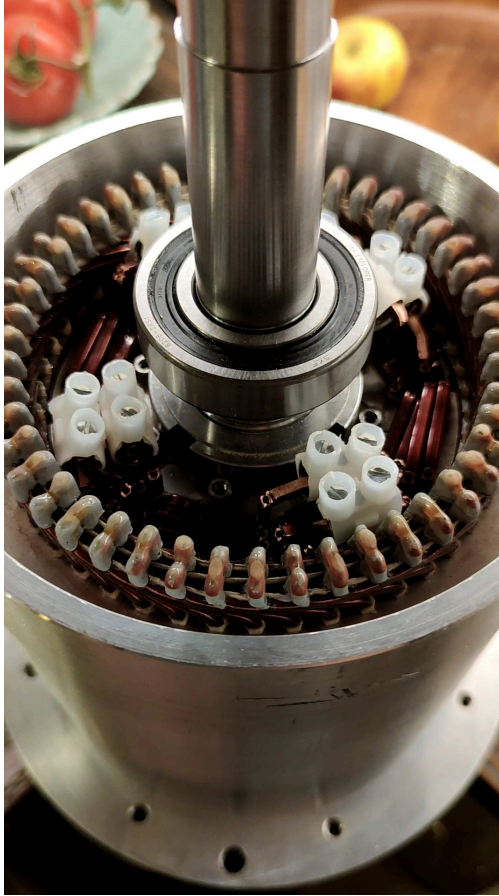
Figure 5.39: Photo of the prototype with the machine in place inside the housing, before the cool side end plate is attached.

the dovetail rotor is a successful way to enable the assembly of a EE rotor with prewound coils with hollow windings for direct air cooling.

The housing is designed, manufactured and assembled, successfully enabling the inlet of coolant air through a 3D printed nozzle. The current is supplied to the rotor with sliprings, which are not optimized for the application and thus limits the testing range to some extent. The design and manufacturing of these and other parts of the whole prototype are presented in order to explain their effect on the experimental test plan and the thermal behavior of the prototype.

Chapter 6

Testing and assessment
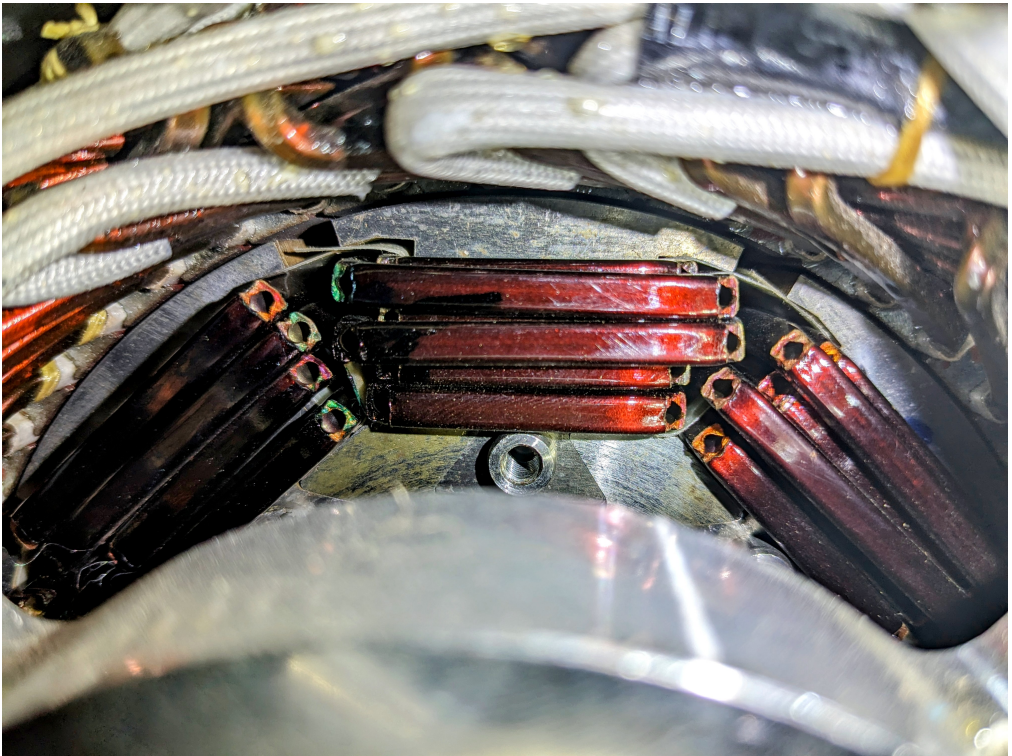
Figure 6.1: Previous page: Photo of the rotor end turns of the prototype.

6.1 Chapter overview

The previous chapter presents the manufacturing and assembly of the prototype, making it ready for experimental testing. This chapter presents the experimental setup, the results from the measurements and the thermal model with focus on the following:

- Presenting the experimental setup and the delimitations in the setup affecting the test plan.

- Measuring the electromagnetic characterization of the prototype and compare it to simulated values.

- Thermal characterization of the prototype in order to verify the cooling concept and the complete design.

- Using a steady-state thermal model of the prototype to further evaluate the cooling concept, how it is actually cooling the machine and what the thermal limits are.

- Using the new thermal limits estimated by the model to evaluate how the performance of the prototype would compare to the initial design and assumed limits that were presented in Chapter 2.

6.2 Experimental setup

A sketch of the experimental setup is shown in Figure 6.2. The machine, with all its parts including housing, shaft and slip rings, are presented in the previous chapter. The experimental setup adds the current supply (marked in bronze color in the sketch), the air supply system (marked in marine blue color in the sketch), the motor bench with a DC machine (marked in pink color in the sketch) and the data logging system (marked in bright bronze color in the sketch).

Below the different parts of the experimental setup are explained along with the delimitations affecting the test plan.
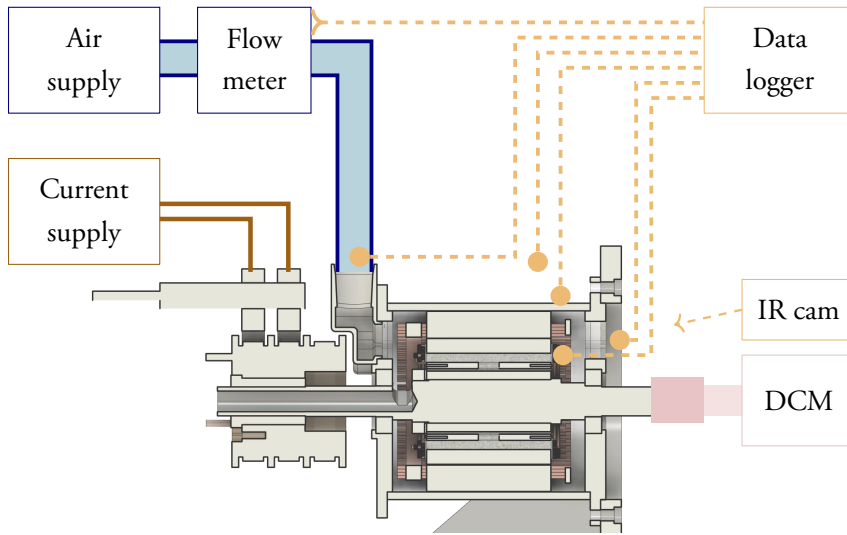
Figure 6.2: Sketch of the experimental setup, cased on the CAD model of the prototype with various sensors as well as the current and air supplies marked.

6.2.1 Motor bench

The prototype was installed in a motor bench as shown in Figure 6.3, sharing shaft with a DC machine (orange colored in the photo) and a surface mounted PM synchronous machine (grey colored behind a steel plate in the photo, and not shown in Figure 6.2). The permanent magnetized synchronous machine (PMSM) was only used to measure the induces voltage in order to derive the rotational speed of the shaft. The DC machine was used as the driving machine, setting the rotational speed of the shaft by applying an armature DC voltage and a magnetization current to reach the desired speed. This means that tests were only made in constant speed conditions (as opposed to dynamic), and due to the limitations of the DC machine the maximum speed available was 1000 rpm. This is far below the mechanical limitation of the dovetail joints of 11 900 rpm divided by a factor of safety (with an undetermined value below 10), which means that the mechanical limitation of the dovetail joints does not affect the test plan of safety measures.

6.2.2 Current supply converter

The current is supplied to the slip ring brushes with a high current buck converter, converting a three phase alternating current to a desired direct current up to 300 A. This means this part of the current supply does not limit the testing. The effects of the current ripple from the buck converter are neglected and not studied further in this work.
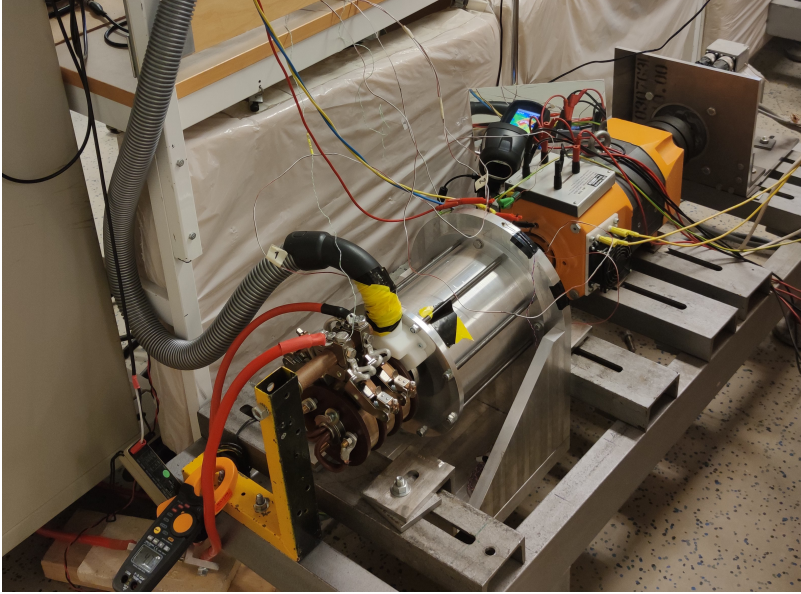
Figure 6.3: Photo of the prototype in the motor bench, ready for testing. From the bottom left the prototype machine is seen before a DC machine and a PMSM.

The provider of the slip rings states that the slip rings and brushes can handle up to 144 A and 6600 rpm continuously, but that cooling or monitoring of the slip ring temperatures can make higher currents achievable. This means that the predicted maximum 195 A rotor current is above the limitations from the slip rings, but that higher current can be tested with cooling or monitoring of the temperatures. The speed limitation does not affect the testing as the driving machine in the motor bench already has a lower limit.

6.2.3 Air supply

The coolant air was supplied with two different vacuum cleaners, one with a (higher) fixed power setting, and one with variable (but lower) power. This means that an airflow rate of up to 1800 l/min can be obtained for constant flow and 500-1100 l/min for variable flow.

6.2.4 Data monitoring and collection

Data is gathered in two ways during the experiments. Some data is continuously logged in a data logger, while other data are recorded manually with a timestamp to enable syncing with the logged data.

Data logger

The data logger stores the data from the input channels every three seconds. (Up to) eight channels were used for thermocouple sensors. One channel was used for the airflow rate and one for air pressure, both measured with a Sonoair VPF-R200-M100 flow meter (see Figure 6.4) installed along a metal pipe between the air supply and the prototype. The thermocouple sensors were used to measure the temperatures of the airflow inlet and outlet, ambient air, housing, windings (in two spots), slip ring brushes and the motor bench mount.

Figure 6.4: Photo of the Sonoair VPF R200 flow meter used to measure the airflow rate.

Manual data logging

Some data was logged by manually writing it down in a spreadsheet with timestamps, so that logger data and manual readings can be synchronized. Among the data logged in this way are the temperature readings from the IR camera and the optical temperature sensor, along with the rotational speed of the shaft (by noting the voltage set to the DC machine) and the current set from the current supply. Manual measurements of winding or slip ring voltages were also logged this way.

6.2.5   Test plan

The prototype was tested only in no-load conditions, meaning no stator current was applied. There are three reasons for this choice. Firstly, the stator was originally designed to be used with a water jacket housing, which was not available for the experimental tests,

so the practical thermal limitations of the stator in operation were unknown. Secondly, the aim of the tests is to verify the rotor cooling concept, which can be done without the stator operating. Thirdly, the lab was only available for experiments for three weeks due to renovations, significantly limiting the extent of tests that could be performed.

The parameters set for the different cases in the experiments were thus rotational speed, excitation current and coolant airflow rate.

A test plan was set up in order to use the time available in the lab efficiently, trying to delay the risks of a thermal or mechanical failure of the prototype as much as possible:

1. **Electromagnetic properties** such as flux linkage, impedance and core losses with rotor currents up to 140 A and rotational speed up to 1000 rpm for shorter times.

2. **Initial thermal characterization** at various airflow rates, **no rotational speed** and rotor current up to 140 A, in transient and post-transient time-spans.

3. **Continued thermal characterization** at various airflow rates, **various rotational speeds up to 1000 rpm** and rotor current up to 140 A.

4. **Extended thermal characterization** at various airflow rates, various rotational speeds up to 1000 rpm and **rotor currents above 140 A** to the extent deemed possible with regard to temperatures.

5. **If possible, more deepened thermal characterization** to improve the understanding of the thermal behavior within possible ranges.

6.3   Electromagnetic characterization

The first part of the test plan was to measure the fundamental electromagnetic performance, which is presented below and compared with simulation results.

6.3.1   Flux linkage

The no-load flux linkage was measured by setting the DC-machine to rotate the shaft in a stable speed and applying a current through the prototype machine rotor windings. The induced voltage in the stator windings was measured and integrated into a value for the flux linkage. The results are shown in Figure 6.5, compared to simulation results.

Two sets of simulation results are shown. One where the dovetail joint is modeled as a 0.05 mm airgap and one with the airgap set to just 0.01 mm. The 0.05 mm airgap was

expected from the tolerances of the machining of the dovetail joints, but the results in Figure 6.5 indicate that the physical contact between the rotor teeth and the rotor yoke is better than expected.
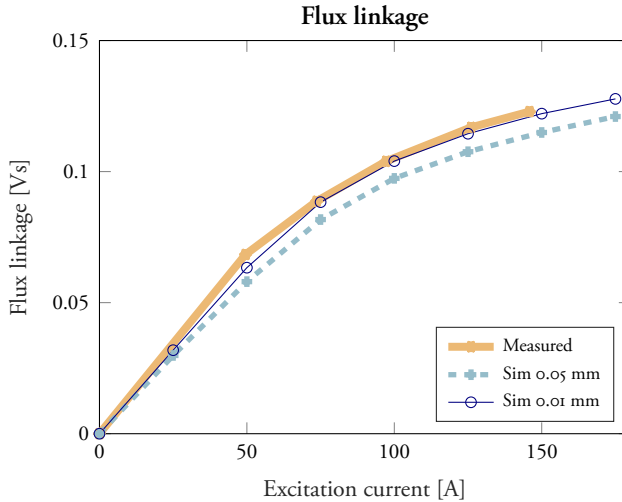
Figure 6.5: Induced voltage at various excitation currents.

The waveform of the measured induced voltage is compared to the simulated induced voltage in Figure 6.6. In both waveforms, six humps are seen for each half period, corresponding to the stator teeth of each rotor pole. However, the humps are distributed differently, with the humps of the measured waveform are shifted slightly along the wave compared to the symmetric shape of the simulated waveform. The cause for this shift is unknown.

6.3.2 Winding impedance

An LCR-meter was used to measure the impedance of the rotor and stator windings. The reason for this is mainly to check that the windings are functional by comparing the measurements to calculated of simulated values. The results are shown in Figure 6.7.

The rotor winding impedance was measured in two ways, one measurements of the windings only, and one including the slip rings. The results of the winding impedance measurement show that the total DC resistance of the rotor windings, including the cables inside the shaft, is 36.3 mΩ. The theoretically calculated resistance of the windings, including the extra resistance for the bends, but not for the screw terminals or the shaft cables, is 32.2 mΩ. The measured resistance is 13 % higher than the calculated value, which is reasonable considering the omitted screw terminals and shaft cables. This suggests that the rotor windings are intact after the prototype assembly.

Induced voltage in phase A at 200 rpm and 75 A rotor current
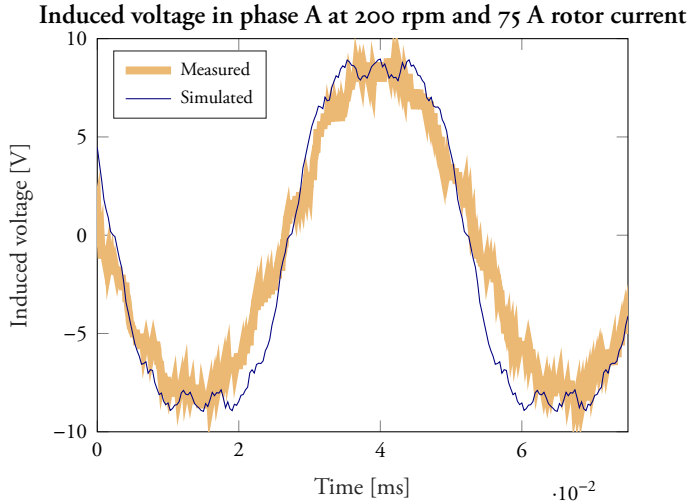
Figure 6.6: Induced voltage wave, measured and simulated.

Rotor impedance   ### Stator impedance
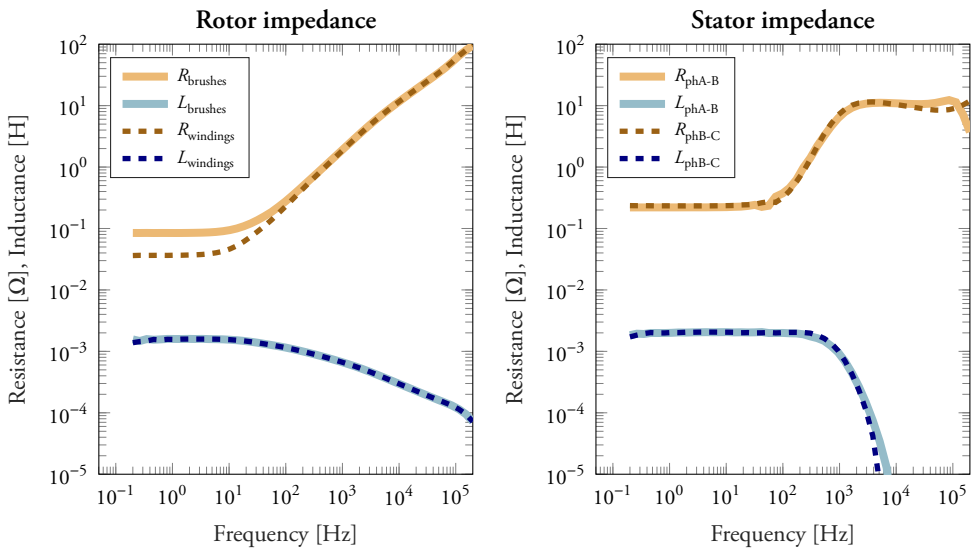
Figure 6.7: Impedance of the rotor and stator windings measured with an LCR-meter.

The second rotor impedance measurement was taken from one slip ring brush-holder to the other. The DC resistance was measured to 84.6 mΩ, namely about 48 mΩ higher than the winding resistance, leaving 24 mΩ per each brush and slip ring.

The stator winding impedance was measured between two phases in three sets (a-b, b-c, c-a). The c-a measurements turned out to be faulty, but since the a-b and b-c measurements

are good and match more or less precisely, there is no reason to expect the fault was in the windings, but rather in the measurement itself. The measured resistance over two phases is about 230 mΩ, leaving 115 mΩ per phase.

The resistance was also measured by repeatedly measuring the winding voltage while applying 140 A of rotor current. 18 measurements were made during 2 hours and 25 minutes with an average voltage of 6.46 V (minimum reading 6.35 V, maximum 6.53 V). This corresponds to a resistance of 46.1 mΩ. The average winding hotspot temperature (measured with the IR camera and converted) was 101.4 ˚C, which explains the higher resistance compared to the LCR data.

During the same time period the machine was repeatedly accelerated to 1000 rpm. The voltage between the brush holders of the slip rings was measured to an average of 7.08 V while rotating (out of 8 readings) while the average voltage at standstill was 7.14 V (out of 7 readings). This was done in an attempt to better understand the load of the slip rings. The conclusion is that resistance over the slip rings does not increase while rotating, at least not in the speed range tested (up to 1000 rpm).

6.3.3 Core losses

The core losses were measured by setting the DC machine to rotate the shaft at a constant speed and then measuring the current through the DC machine. The nominal flux linkage of the DC machine can then be used to derive the torque required to keep the system rotating at the current speed. This was repeated for a series of speeds without applying any current in the prototype machine. The set of torques derived in this way corresponds to the total losses of the DC machine and the mechanical losses of the whole motor bench.
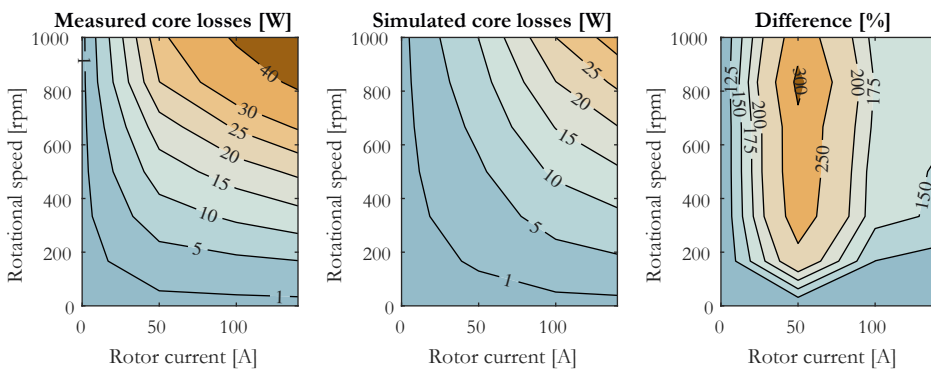
Figure 6.8: Core losses, measured experimentally, simulated and compared.

This procedure was repeated with various rotor currents applied in the prototype machine, increasing the core losses with increased magnetization. The results are shown in Figure 6.8

and compared to the simulated core losses. The comparison shows that the measured core losses are up to three times as high as the simulated core losses. There could be several reasons for this discrepancy. One is that the accuracy of the loss models in the materials used in ANSYS are not validated. It can also be the case that the grinding of the rotor teeth has caused short circuits between laminated steel sheets. This is not analyzed further in this work, as the effect on the cooling concept is limited. The comparison between the EE prototype machine and the equivalent IPM machine is based on the core loss models which may be inaccurate, but at least both machines use the same core loss models and the comparison is therefor deemed to be fair nonetheless.

6.4   Thermal characterization

The main part of the experimental tests consists of thermal characterization, in order to verify the cooling concept of the prototype. This makes up step 2-5 of the test plan. It starts with a focus on understanding how to measure and interpret the temperatures of the prototype in step 2, before going on to the continued, extended and deepened testing of steps 3-5. However, the presentation of the results in the following sections are not necessarily following the test plan steps, but rather focusing on one aspect of the testing at a time.

6.4.1   Measuring temperatures of the prototype machine

Measuring the temperatures of a rotor is not easy. There are different ways to do it, but none of them are perfect.

Measuring the hotspots of the windings are the single most important data to verify the function of the cooling concept. The hotspot is expected to be located on the end turns on the outlet side of the machine, either by the bends where the resistance is higher, or in the middle of the end turns farthest from the direct cooling channels.

Thermocouple sensors

Thermocouple sensors are accurate and easy to use with automatic data logging. But the cable makes them impossible to use during rotation if there is not either a special set of slip rings for transferring the signal from the rotating shaft, or a wireless solution. Both of these technologies exist, but were not possible to implement on this prototype given the timespan during which the lab was available before renovation.

Thus, thermocouple sensors were used for stationary measurements like inlet and outlet coolant air temperature (Figure 6.9), housing and ambient temperature, plus some initial winding temperature measurements before rotational speeds were introduced (Figure 6.10). It is difficult to measure the outlet air temperature, since the sensor must be placed so that the air is blended as little as possible with the surrounding air. Other actions in the outlet window will interfere with the sensor output, such as moving the winding sensors or measuring the winding temperature with the optical sensor.
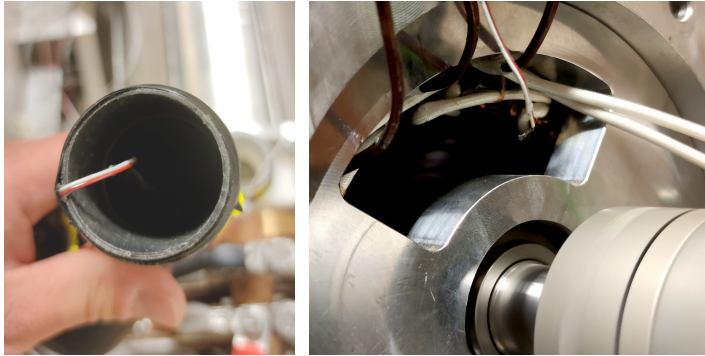
Figure 6.9: Photos of the thermocouple sensors in by the coolant air inlet (left) and outlet (right).

Figure 6.10: Photo of two thermocouple sensor inserted between the turns of two different coils.

Optical temperature sensor

The optical sensor Optris CT measures the temperature by using frequencies that makes accurate temperature readings on metals with an error of around ±2.5 °C in the temperature range used (of up to 100 °C). The sensor is pointed at the target and needs to be close to the target (but not touching). This means the sensor needs to be inside the outlet window, and in there it is quite bulky and makes out a significant obstacle for the outlet air. In

other words, it can not be placed to constantly monitor the winding temperature. Instead it was held in hand to search around along the winding end turns for a highest reading of the winding temperature. This must be done quickly, in just a few seconds, to not cause too much increase in the winding temperature due to the outlet air being hindered by the sensor and the hand (See example in Figure 6.21). This means that despite the technical accuracy of the sensor, the accuracy of the data depends on location of the measurement (and whether a real hotspot is found). Furthermore, this has to be done quickly, so that the winding temperature is not distorted by the hindered coolant air.

Figure 6.11: Photos of the optical temperature sensor with the sensor in the top photo and the electronic box in the bottom photo.

For that reason, the optical sensor was used mostly to calibrate the IR camera readings, mainly in a non-cooled case (with no outlet air to be hindered) but also occasionally in other cooled cases (more about this in the section about temperature sensor correlation below). The readings from the optical sensor were noted manually from the display of the electronic box. In the data presented in this chapter, the results from the optical sensor are generally shown as a ring or 'o'.

IR camera

A Flir i7 camera was used for most of the temperature readings. It is a very practical instrument that can be used quickly during both stationary and rotating tests (either by stopping the machine and reading quickly, or even during rotation). But its accuracy is highly dependent on the reflectivity and emissivity of the object measured. Metals like copper are generally too reflective for IR cameras to read the temperature accurately.

The insulation varnish reduces the reflectivity. To test the reflectivity of the varnish on the windings used in the prototype windings an experiment was performed (Figure 6.12). A piece of copper winding was held in boiling water and then brought up for a reading with the IR camera.

Figure 6.12: Photo of the equipment used for the IR camera test.

Figure 6.13 shows the results from the IR camera. The left figure shows how the temperature from inside the water boiler is around 100 °C but the winding piece freshly brought out of the boiling water is significantly cooler (and the stainless steel of the boiler turns out blue on the color scale despite being too hot to touch). The IR camera looks for the hotspot within the square in the middle of the total view and in the left photo the hotspot if the winding piece is only 82.2 °C, since it reflects too much of the cooler room and emits too little of its own heat (However the black electrical tape on the plier show a higher and more accurate temperature). It is concluded that the temperature readings on the copper windings can not be trusted to be accurate.

The IR camera was fixed in a stand pointing at the outlet window of the prototype. An attempt was made to paint the end turn of a winding or to put a strip of black tape, but the painted or taped parts only proved to show lower temperatures than the unmarked windings when the winding temperature was increased, thus not attracting the hotspot search of the IR camera.
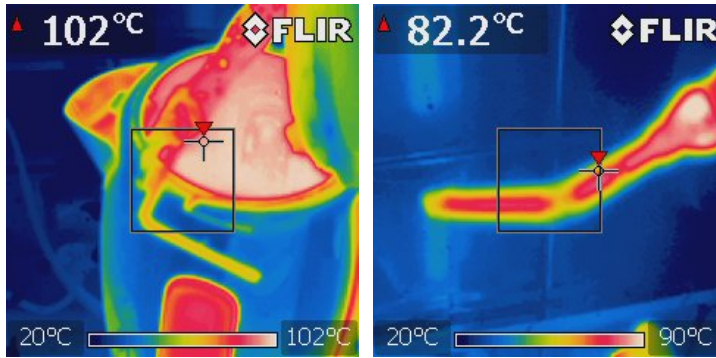
Figure 6.13: Photos from the IR camera during the test.

A photo from the IR camera is shown in Figure 6.14, along with a normal photo from the same angle. Exactly when the IR camera photo was taken, the thermocouple sensors measured 83.8 °C and 73.7 °C respectively (The thermocouple sensors had been removed when the normal photo was taken and are thus not visible).
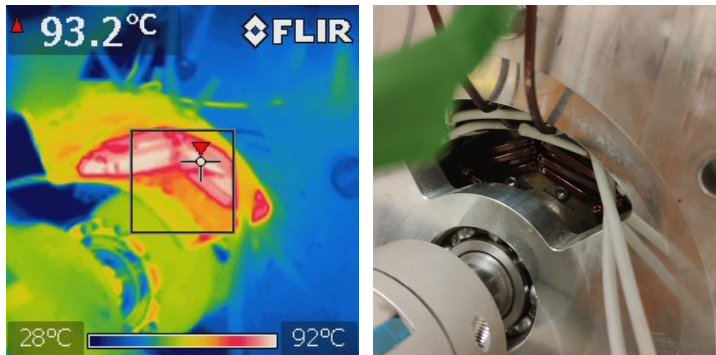
Figure 6.14: Left: Photo from the IR camera pointed at the windings. Right: Normal photo from the same angle.

Figure 6.15 shows a photo of the IR camera mounted in the position used in Figure 6.14.

Temperature sensor correlation

Due to the uncertainty in IR camera measurements, as many measurements as possible were supplemented with measurements with either thermocouple sensors or the optical sensor. In this way it was possible to achieve at least additional temperature besides the IR camera readings for a number of different occasions. All of these cases are shown in Figure 6.16, with the value of the IR camera reading along the x-axis and the other sensor on the y-axis. This means that the IR camera vs IR camera-series forms a 1:1 line.
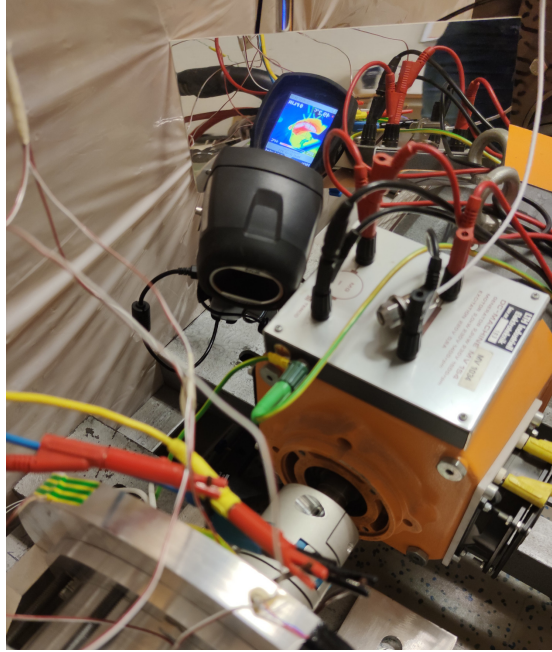
Figure 6.15: Photo of the IR camera pointing at the outlet window of the prototype, with a mirror to ease readings.
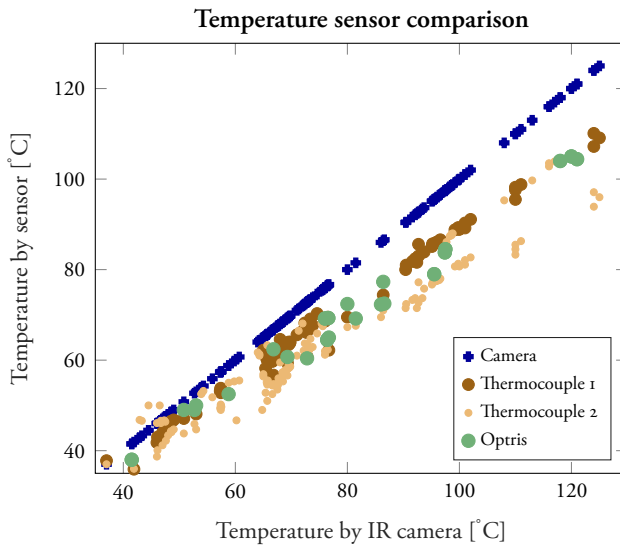
Figure 6.16: All IR camera readings of the winding temperature where one or more sensors were used simultaneously.

There is some distribution of in the data from the other sensors, but all but a few show a lower temperature than the simultaneous IR camera reading. Especially the readings from

thermocouple sensor 2 are more scattered and also generally lower than the others. This indicates this sensor had either worse contact with the windings, or that this particular coil had better cooling due to individual differences between the coils. The readings from this sensor are thus less interesting, both because they are more more distributed and also generally lower. The goal of the winding temperature measurements is to find the winding hotspot.

For this reason, the data from thermocouple sensor 2 are omitted, while the data from thermocouple sensor 1 and the optical sensors are used to make one first degree curve fit each. The resulting lines are shown in Figure 6.17.
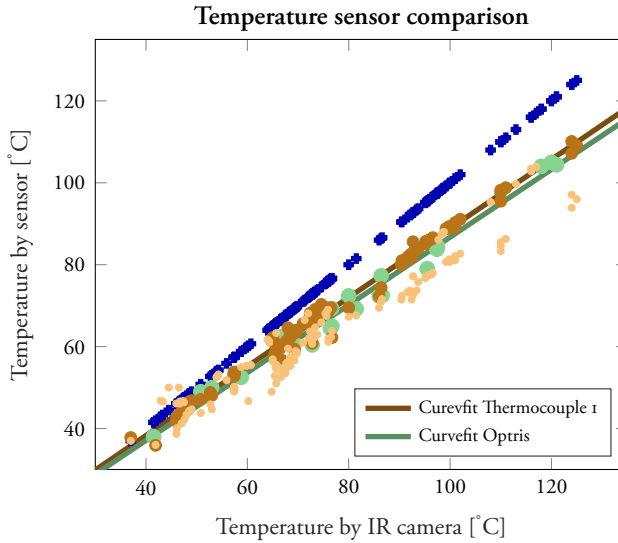
Figure 6.17: All IR camera readings of the winding temperature where one or more sensors were used simultaneously, along with first degree curve fit lines for the Optris and Thermocouple 1 temperatures.

The average of the two curvefit lines in Figure 6.17 is shown in equation 6.1. For all cases where the temperature was measured with both the IR camera and at least one other more accurate temperature sensor, the reading from the IR camera will be closer to the more accurate value if it is converted through the formula of equation 6.1.

$$T_{\mathrm{mean}} = 0.8346\,T_{\mathrm{IRcam}} + 4.3244 \tag{6.1}$$

For the rest of the measurements of the prototype this is assumed to hold true also for the cases where no other sensor but the IR camera was used. So, if the IR camera for one case shows $100\,°\mathrm{C}$ on the windings, the best estimation of the actual hotspot temperature of the windings is $100\,°\mathrm{C} \cdot 0.8346 + 4.3244 = 87.8\,°\mathrm{C}$. Had the temperatures of all sensors been as spread as those from thermocouple sensor 2 compared to the readings from the

IR camera, the assessment would be that the temperature data from the IR camera is too inconsistent to be used, but since two different sensors are so consistent and with almost equal curvefit functions, the assessment is that the readings from the IR camera are usable with the mentioned conversion formula (equation 6.1).

Winding temperatures of non-visible poles

In order to test whether there is any substantial difference in cooling for the poles at the bottom of the rotor (the coolant air inlets and outlets are positioned at upper half of the end plates), the winding temperature was measured with the IR camera as the rotor was rotated one pole at time (45 degrees) and then 180 degrees, while applying 75 A of rotor current. The results are shown in Figure 6.18 where the shaft rotations are marked with vertical dotted lines. The first four rotations are 45 degrees and the last is 180 degrees. Four of the five rotations result in an increased measured windings temperature (of 0.3 to 1.4 degrees), and one in a decreased temperature (of -1 degree). This change is then followed by a similar decrease in temperature in the following reading before the next rotation. This suggests that there are some individual differences between the poles (in resistance and/or coolant flow) and that the uppermost poles are cooled slightly better, as they offer a straighter and shorter path for the coolant air, but the difference is not significant and certainly not something that invalidates the concept.

Winding temperatures with shaft rotations
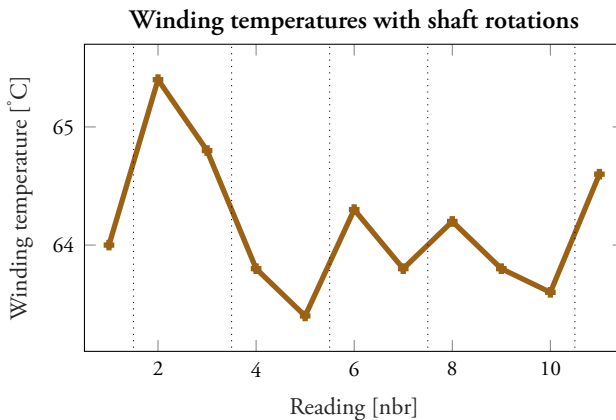
Figure 6.18: Winding temperatures as shaft is rotated to pot different poles in the outlet window. Shaft rotations are marked with vertical dotted lines.

Temperature measurements of rotating shaft

In order to verify whether the IR camera can measure the winding temperature while the shaft is rotating, an experiment was made where the winding hotspot temperatures were

measured as the machines were set alternatively rotate and be still, while readings were taken from the IR camera. The results can be seen in Figure 6.19 (with converted data from the IR camera). The results show that the IR camera tends to give deviating readings when the shaft is rotating, while the readings from the stationary shaft are more consistent. This shows that the IR camera should not be used to read the winding temperature while the shaft is rotating. Instead, for tests with rotating shafts, the machine should be stopped and the reading taken immediately after the shaft is stopped before the temperature has time to adapt to the standstill case.
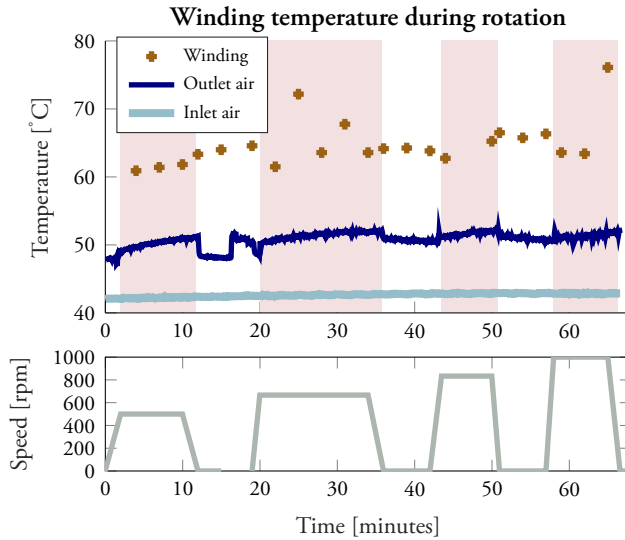
Figure 6.19: Winding temperatures as shaft is rotating or being still in turns. The shaft is rotating during the areas marked in pink. Rotational speed is shown in bottom figure. (After 12 minutes the outlet thermocouple sensor is accidentally put out of place before being reinstated at around 20 minutes.)

6.4.2 Steady-state behavior

The goal of the direct cooled rotor is to achieve a high continuous power, not just a high temporary peak. Therefore, the focus is to evaluate steady-state events and not just transients. In order to understand the nature of the steady-state events of the prototype, two initial long-time measurements were made.

The first one, shown in Figure 6.20, shows the temperature development of the prototype with no rotor current applied during 28 minutes. The results show that the inlet air from the air supply is about 40 °C after around 10 minutes. At the beginning of the measurements, the vacuum cleaner is cool (room temperature) as well as the aluminium pipe around the airflow sensor, and before they are both heated up to around 40 °C, the air at the inlet of the prototype machine is lower than that. After 28 minutes the inlet temperature has reached

41.4 °C (marked with the dotted line in the figure) and the outlet temperature has reached 38.1 °C. One conclusion to draw from this is that the inlet airflow needs to be activated for at least 15 minutes before other measurements are started, since the inlet temperature has not fully developed until then.

The whole prototype is still not at full equilibrium after 28 minutes, but the development is slow and other factors start to come into effect. This is better explained in the next measurement. The winding temperature was measured with a thermocouple sensor and had not fully cooled down to ambient temperatures, from earlier experiments, which is the reason why it is not cooler than the outlet air.
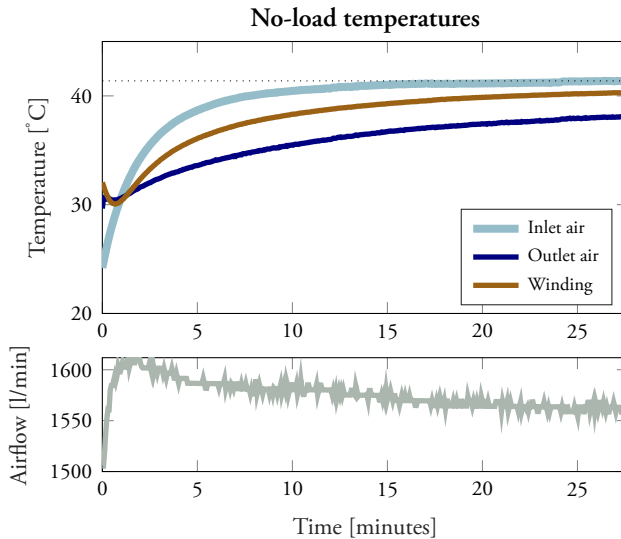
Figure 6.20: Temperature of inlet and outlet air, as well as no-load windings, over a timespan of 28 minutes after the air supply is turned on.

The rotor current was set to 75 A and the temperatures were measured for 32 minutes, shown in Figure 6.21. In the previous measurement the coolant air was heating up the machine rather than cooling it. In this measurement it is clear the opposite is happening. The winding temperature (measured with thermocouple sensor) is significantly higher than the air temperatures, so now the coolant air is really cooling the machine.

There are several things to be noted. One is the spike in the outlet air temperature at about 4 minutes, which is caused by a manual measurement of the winding temperatures with the optical temperature sensor. The hindered coolant air causes a small spike in the thermocouple measured winding temperature too, but this effect is quickly eliminated.

Another thing to note is that the winding temperature is still not 100 % stable after 32 minutes, when it has reached 64.1 °C. After 20 minutes, the measured temperature is 63.2 °C,

one degree lower. After 10 minutes, the measured temperature is 61.9 °C, 2.3 degrees lower and after 5 minutes the measured temperature is 60.8 °C, 3.4 degrees lower. But during this time, other factors change as well. During the 32 minutes of this measurement, the ambient temperature increases by 1.2 degrees and the airflow decrease by about 2 % (Note the narrow band of the y-axis of the airflow plot). On top of this, there is a slow heating up of the rest of the machine bench, slightly affecting the heat flows in the machine.
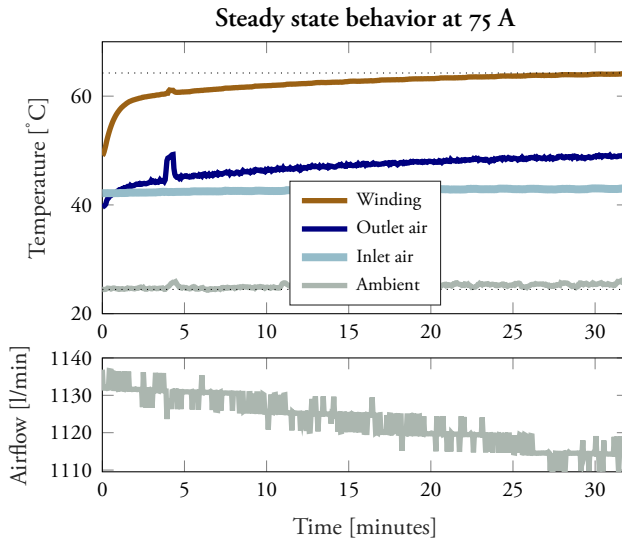
Figure 6.21: Temperature of windings with an excitation current of 75 A over a timespan of 32 minutes.

The conclusion is that there is no true equilibrium to expect in practice. Important variables as the ambient temperature, machine bench temperature and airflow will vary over time it takes for the whole prototype to reach completely stable temperature, thus constantly changing the circumstances of an equilibrium before it is reached. However, only 10 minutes after the transient begins the winding temperature change rate is about 1 degree per 10 minutes (a rate which would be declining further if given time). So, as further experiments are concluded when considering limits of the cooling with respect to the maximum allowed copper temperature, a safety margin of some degrees should be used rather than an assumption about steady-state being reached for the prototype.

6.4.3 Temperatures for various rotor currents

The hotspot winding temperature was measured for various rotor currents, with results shown in Figure 6.22. The temperatures shown are converted readings from the IR camera. The current was set to 50 A, 70 A, 85 A, 100 A, 120 A and 135 A (An attempt was made to reach 140 A, but the voltage supply to the current controller could not maintain this

current at the time, so it was set down to 135 A. The voltage supply was changed after this experiment.)
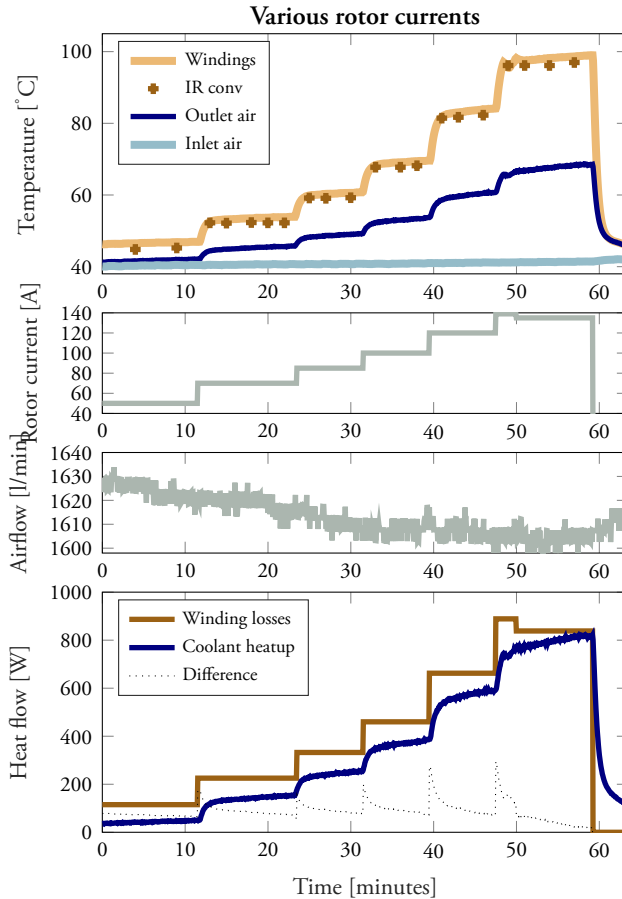
Figure 6.22: Winding temperatures for various rotor currents.

The figure also shows the estimated winding losses and the heat carried away by the coolant. The winding losses are derived from the current and the measured resistance at 101.4 °C (see section 6.3.2). This means the temperature dependence of the winding resistance is not taken into account (This is integrated in the thermal model presented in section 6.5.4). The coolant heatup is derived by multiplying the volume flow and the temperature rise from air inlet to outlet (along with the density and specific heat capacity of air at 40 °C). The difference between the energy flows is shown in the dotted black line, which consistently decreases to around 80 W after each transient. This suggests that around 80 W of heat is escaping the prototype via convection to the rest of the motor bench and via passive convection to ambient air. Note that this involves comparing the continuous data from the data logger (of the air temperatures and airflow) with the manually logged data (of

the IR camera and the applied current). The lines between the +-signs are thus not to be interpreted as continuous development.
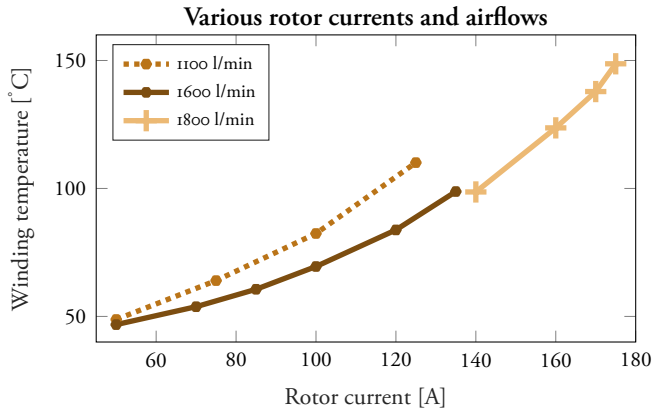
Figure 6.23: Winding temperatures at various rotor currents and airflow rates.

The final temperature reading from each current level is shown in Figure 6.23.

6.4.4 Temperatures for various airflow rates

The hotspot winding temperatures for various airflow rates were measured, presented in Figure 6.24. The temperatures shown are converted IR camera readings. The air supply was varied from 695 l/min to 827 l/min, 1029 l/min and finally to 1067 l/min.

The results show that the air inlet temperature is lower for lower flow rates. Less airflow generates a smaller pressure drop along the supply tube and through the prototype, which means the air is heated up less through compression in the pump of the vacuum cleaner. The outlet air temperature is more or less constant through all flow rates, making the temperature difference between inlet and outlet larger for lower flow rates, but lower flow rates also means less mass flow of the coolant. In the end, the heat carried away by the coolant is more or less constant through the different airflows. In other words, while the airflow rate is increased by over 50 % (from 695 to 1067 l/min), the outlet temperature increases only 2 degrees (from 57 to 59 °C), the heat dissipated by the coolant decreases some 1 % (from 299 to 296 W) and the winding temperature decreases 4.1 degrees (from 88.6 to 84.5 °C). The changes in inlet temperature are accounted for in the mass flow used to derive the dissipated heat.

In Figure 6.25 the average winding temperatures for each flow rate are shown (except 1030 which is omitted due to too close to 1067) .
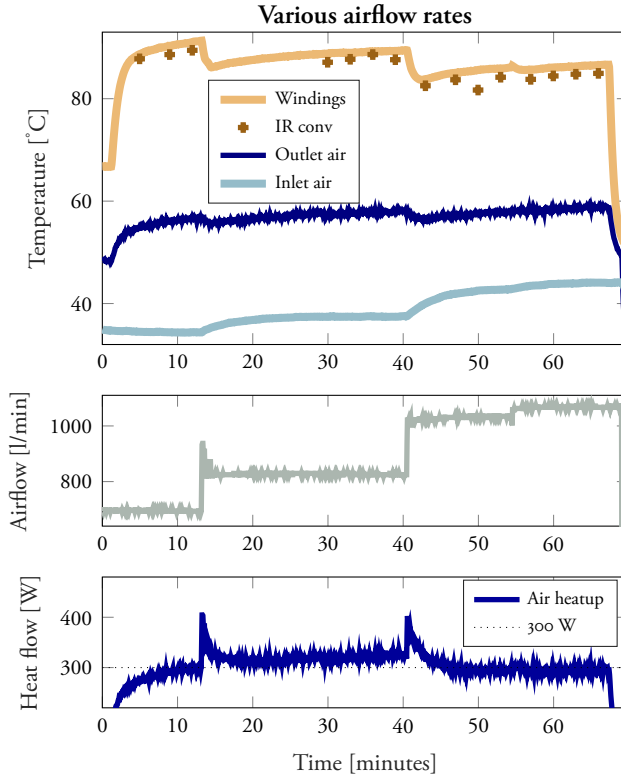
Figure 6.24: Winding temperatures for various airflow rates.
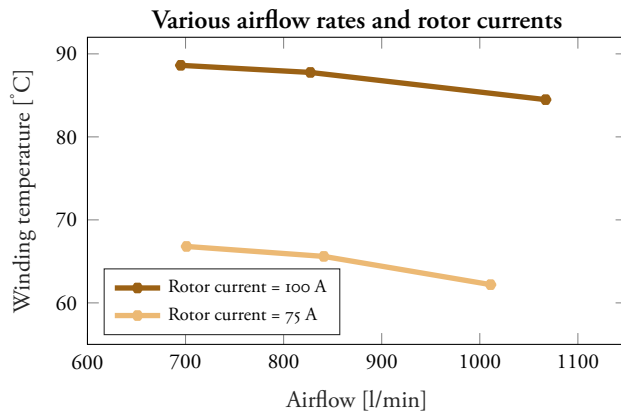
Figure 6.25: Winding temperatures for various airflow rates.

6.4.5   Temperatures for various rotational speeds

The hotspot winding temperatures were measured just after stopping the machine from various speeds (0, 330, 670 and 1000 rpm) and the results are shown in Figure 6.26. Since the machine was rotating, the only sources for the winding temperatures are the IR camera and the optical sensor. The optical sensor was only used for the first measurements after shaft was stopped (and it can be seen in the outlet air temperature how it disturbs the airflow, heating up the outlet air). The IR camera readings from when the shaft is spinning can not be trusted to be accurate (even if they are converted).
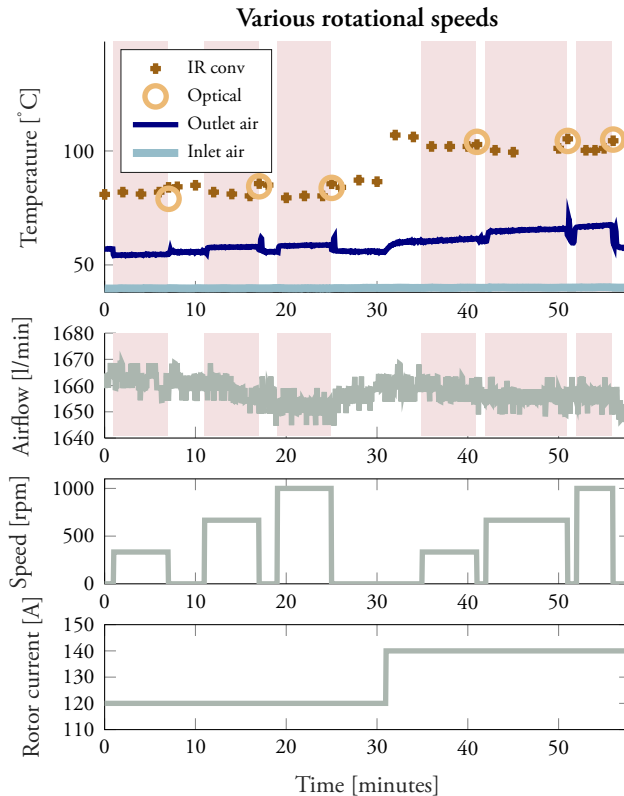
Figure 6.26: Winding temperatures for various rotational speeds. Pink areas mark that the machine is spinning.

The outlet air temperature increases slightly when the machine is accelerated to 670 rpm or higher. It is unknown whether this is because of the core losses increasing the machine temperature which results in more heat being transferred to the coolant air, or if the cooling is more effective when the rotor is spinning (either by increasing the flow around the end turns or making the airflow faster and more turbulent through the cooling channels) thus decreasing the temperature of the whole machine while it is spinning. To know this,

temperature sensors in the machine core are necessary.

The winding temperature readings (summed up in Figure 6.27 together with more series than the one shown in Figure 6.26) did not clarify this. The readings are too inconsistent to draw any conclusions from. Is the winding temperature generally increasing or decreasing as the speed increases? It can not be told from the data available. If anything, it can be concluded that the rotation of the rotor does not have a significant effect on the winding temperature. But this only holds for low speeds. The behavior at higher speeds, such as 5000 rpm or 10 000 rpm, can not be predicted from these experiments. It may be that the cooling is more effective the higher the speed, or there may be a certain rotational speed where the turbulence around the inlets to the cooling channels becomes so vicious that it completely seals of the inlets, disabling the cooling altogether.
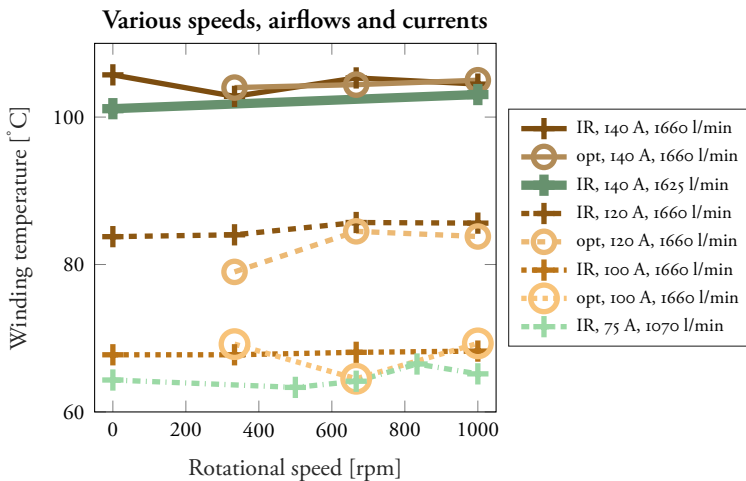
Figure 6.27: Winding temperatures for various airflow rates.

6.4.6 Case consistency

Throughout the different experiments, some case (of current, speed, airflow) occurred several times. The winding temperatures of the most common case, 140 A at 0 rpm and ~1600 l/min, is shown in Figure 6.28. The different cases shown are taken on different days, or on the same day with with several other cases in between. The temperatures have some variation (5 degrees from smallest to largest), some of which can be attributed to small measurable differences in airflow rate or inlet air temperature. Other factors that will affect the outcome are differences in how well developed the steady state case was allowed to get before the measurement was taken (as well as what the previous case was and if the transient came from a higher or lower temperature), the ambient air temperature, the temperature

of the motor bench and the accuracy of the IR camera (which was the only sensor used for most of the cases).
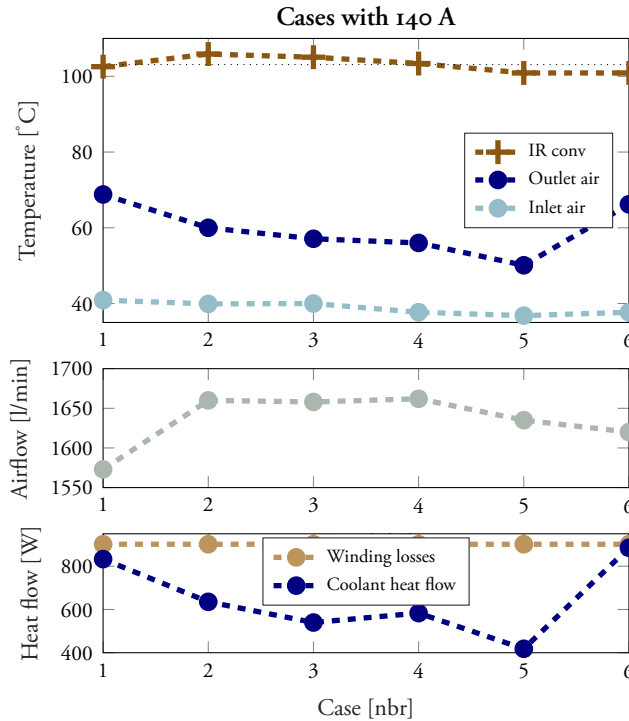
Figure 6.28: Winding temperatures for similar cases of 140 A, 0 rpm and approximately 1600 l/min. Average winding temperature for the cases is shown with a black dotted line (103.1˚C).

The conclusion is that one case can not be exactly replicated due to too many unknown and/or uncontrollable factors. The outcome of a case may vary by ±2.5 degrees without any measurable difference in the input variables of the case.

It can also be noted from these results that the measurements of the outlet air temperature too are inconsistent between the different experiments. The difference between the estimated winding losses and heat dissipated by the airflow is so large in some cases that it most likely depends on the outlet thermocouple sensor being moved slightly out of the outlet air, so that it is too mixed with ambient air.

6.4.7   All tested cases

All cases run during the experimental tests of the prototype are shown in Figure 6.29 along with the results from the steady-state thermal model. They are shown in the order they were

performed in, no matter if the cases are made the same day or different days. Compared to the test plan presented in section 6.2.5 the cases focus on the following:
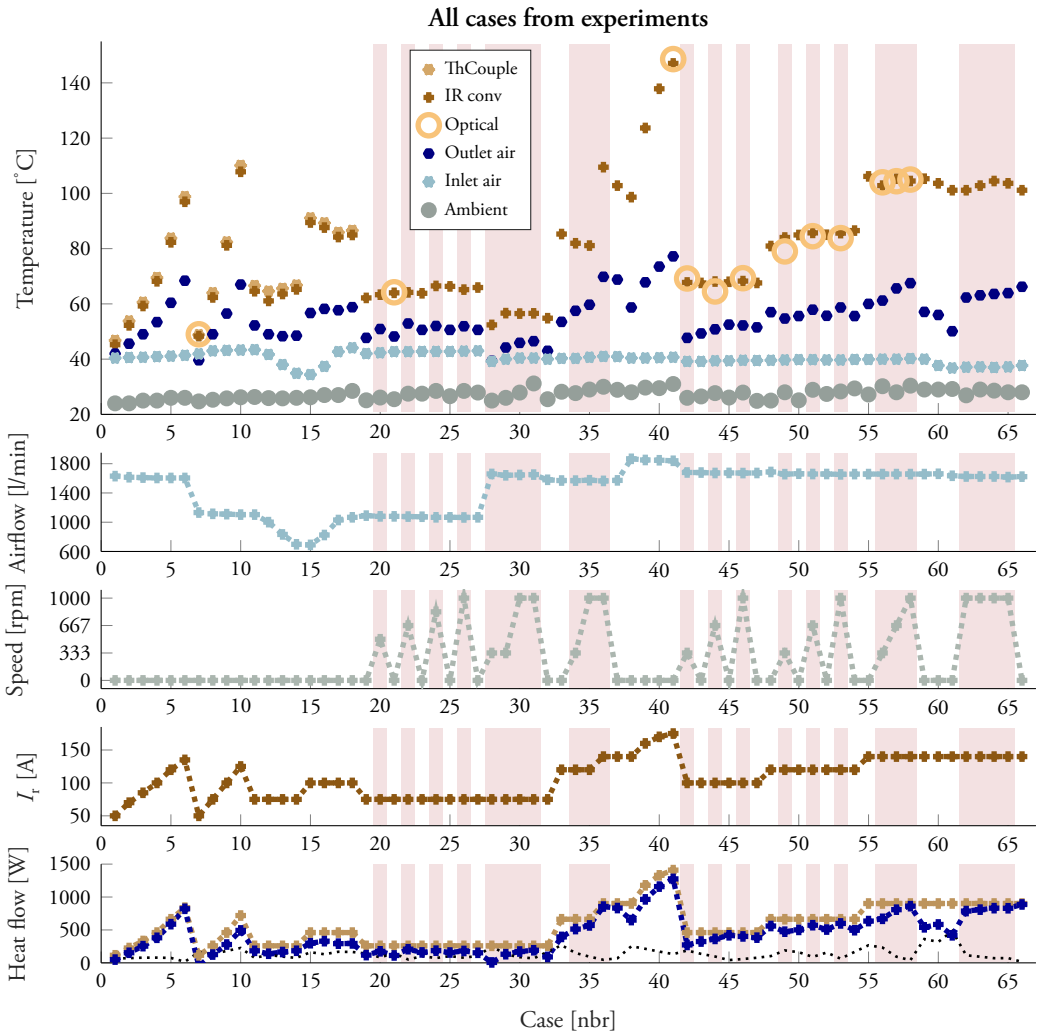
Figure 6.29: All cases tested during prototype experiments. The bottommost plot shows the estimated resistive losses (bronze), the heat taken up by the air (blue) and the difference between them (dotted).

1. The results of the **electromagnetic characterization** are not of use to compare with the steady-state thermal model and are thus not shown here. These results are instead compared with simulation results as they are presented above.

2. The **initial thermal characterization** in transient cases are not of use to compare with the steady-state thermal model. The post-transient results are shown as cases 1-19,

with various rotor currents (two sweeps up to 140 A in cases 1-10) and various airflow rates (cases 11-19).

3. The **continued thermal characterization** is shown in cases 20-37 focusing on various rotational speeds and airflow rates.

4. The **extended thermal characterization** at higher rotor currents is made in cases 38-41. This is done without rotational speed in order to better monitor the temperatures to avoid failures.

5. The **deepened thermal characterization** in cases 42-66 are focusing on various speeds again to try to make sense of the behavior at different speeds.

6.5   Steady-state thermal model of prototype

A thermal model of the prototype is developed in Matlab. The foundation of the model is the combined empirical-lumped parameter model from the Licentiate thesis [9, 16], which is a steady-state model that iterates between empirical models for the forced convective heat transfer to the coolant, and a lumped parameter thermal network model of the temperature distribution of the hollow windings (summed up in chapter 1.2.3). The prototype studied in this thesis need some adjustments to that model. The main differences are:

1. The thermal network must be extended to model the whole machine and not just look at the windings as an isolated system. The updated thermal network is shown in Figure 6.31 and 6.32 and involves heat transfer to the surrounding ambient air, the machine bench, to the coolant air from the end turns of the windings, and from the core losses of the machine.

2. The coolant air is allowed to flow in three different types of channels, namely inside the hollow windings, in the prism shaped gaps between the coils, and in the airgap. This means that the forced convection model is run as three separate systems with different dimensions where different heat transfer coefficients are derived.

3. Because of the three different types of channels for the coolant air, the distribution of the total airflow must be determined.

The steps of the model are described below with all the involved variables explained.
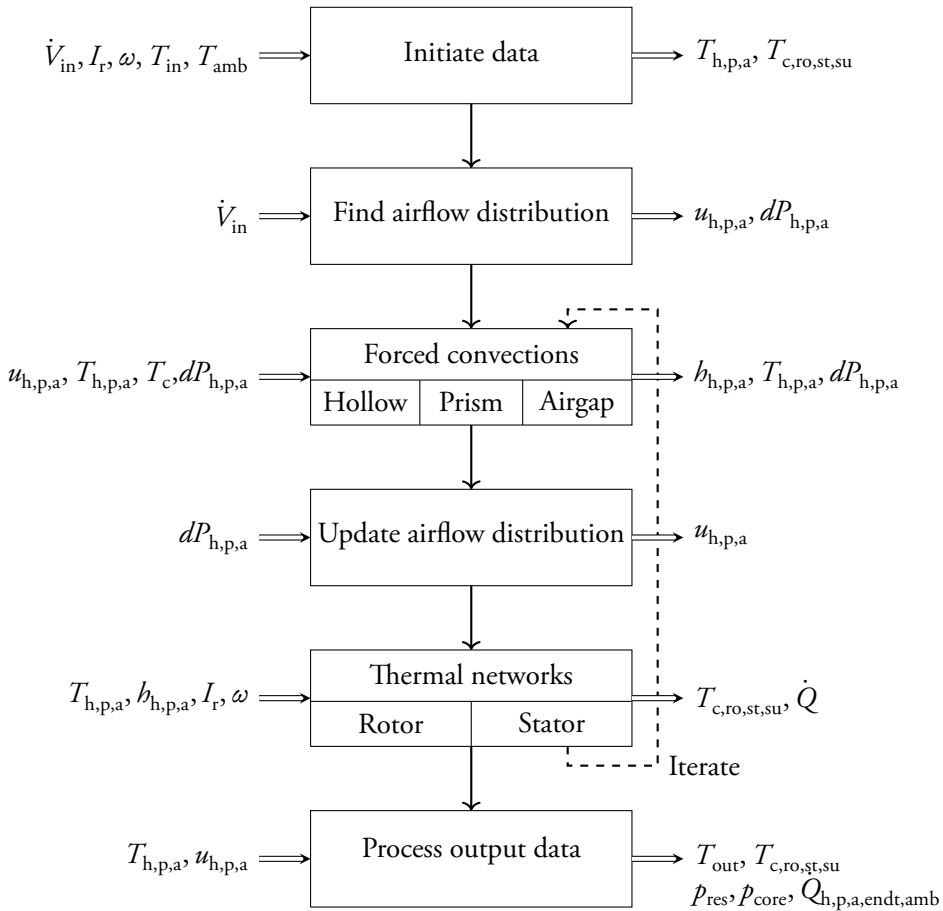
Figure 6.30: Flow chart of the steady-state thermal model of the prototype, with the in and out data of each step.

6.5.1 Initiation

The model variables given input data are shown in Table 6.1 and the variables given initial values are shown in Table 6.2.

6.5.2 Finding airflow distribution

The channels where the air is allowed to flow (hollow channels, prism channels and airgap) all have different dimensions and shapes, which results in different pressure drops along each of the channels if the air velocity is set to the same value. Therefore the channels will have different air velocities, that will result in an equal air pressure drop along the different channels, $dP_{h,p,a}$. These three different air velocities (u_h, u_p and u_a) are found in an iterative

Table 6.1: Model input data

Symbol	Description
$\dot{V}_{\text{in}}$	Inlet airflow (in liters/minute)
I_{r}	Current provided to the windings
ω	Rotational speed of the machine
T_{in}	Coolant air temperature at the inlet
T_{amb}	Temperature of the ambient air surrounding the machine

Table 6.2: Parameters given initial values

Symbol	Description
T_{h}	Air bulk temperatures along the hollow channels
T_{p}	Air bulk temperatures along the prism channels
T_{a}	Air bulk temperatures along the airgap
T_{c}	Temperature along the copper windings
T_{ro}	Temperatures along the rotor teeth
T_{st}	Temperature along the stator teeth
T_{su}	Temperatures along the stator surface

process, so that all channels have equal pressure drops and the total airflow of all channels (summing up all 12x8 hollow windings, 2x8 prisms and 360 degrees of airgap) equal to $\dot{V}_{\text{in}}$.

When the air in the channels are heated up due to cooling of the machine, the pressure drop will change. Thus, the distribution of the airflow is updated for each iteration of the model.

6.5.3   Forced convection empirical model

As mentioned above, the three different types of air channels are solved as separate problems. The hollow channels are solved as circular pipes, the prisms as triangular pipes and the airgap as an infinitely long parallel gap (when it comes to the parameters that depend on the shape of the channel, but for the surface area, the airgap has a finite width). The hydraulic diameter (and the friction factor in laminar flow) is different for these different shapes, as well as the dimensions themselves. Apart from that, the models of all three channel types are the same as in the Licentiate thesis [9], solved as constant wall temperature forced convection problems for each section of the pipes. The pipes are divided into 120 sections, one for each millimeter of the machine.

The models use the air velocity and the bulk air temperature to determine the heat trans-

fer coefficient for each section (h), the pressure drop over each section (dP) and other air properties (such as density, viscosity and thermal conductivity). They then use the heat transfer coefficient and the wall and fluid temperatures (T_c, T_h, T_p, T_a) to determine the heat flow and coolant temperature rise so that the next section is solved with the correct air bulk temperature and air velocities.

The summed pressure drops of each channel types are then used to find an updated air-flow distribution, so that the updated air velocities (u_h, u_p, u_a) are used for the next forced convection iteration.

6.5.4 Thermal network

The thermal network of the prototype is divided into two separate networks, one for the rotor and one for the stator. The radial plane of the networks are shown in Figure 6.31.

The rotor network has one temperature node for all of the copper in the cross section of a half pole. This means the model does not distinguish temperature differences between individual winding turns, but assumes the heat transfer is so effective within the cross section of each conductor, as well as between them, so that any temperature difference along this plane can be neglected.

The other temperature node in the rotor network is in the rotor teeth, and represents the temperature of the rotor tooth surface towards the airgap. The rotor core losses are generated along this edge. The thermal resistance between the copper and rotor tooth is calculated using the physical dimensions of the tooth size and the thermal conductivity estimated to 22 W/mK. The heat transfer coefficient to each channel type (hollow, prism, airgap) is provided by the forced convection model. Note that even though the figure shows two whole prisms, only two halfs prisms (equal to one whole) are used for the half pole symmetry of the model.

The stator network has one node at the teeth surface towards the airgap, much like the corresponding rotor node, where the forced convection to/from the airgap coolant as well as the stator core losses are situated. The other stator node corresponds to the surface of the housing bordering to the surrounding ambient air. The natural convection heat transfer coefficient to the ambient air is set to 20 W/m²K. The thermal conductivity from the tooth node to the surface node is 22 W/mK

There are 120 such transverse planes along the machine, one for each section of the forced convection model. The radial plane of the network is shown in Figure 6.32.

The thermal conductivity between the stator teeth nodes in this plane is 60 % of 22 W/mK because of the laminated steel. Between the rotor teeth nodes it is 50 % of 22 W/mK.
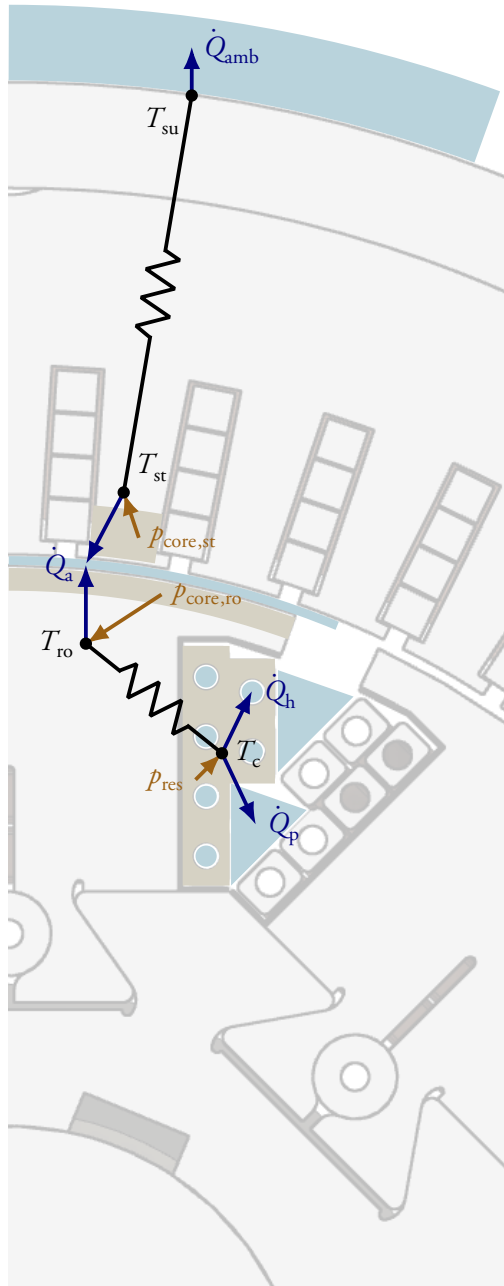
Figure 6.31: The thermal network(s) in the transverse cross-section. Cooling areas marked with bright blue, heating areas marked with beige. Stator and rotor core temperatures are modeled as the surface temperature towards the airgap.
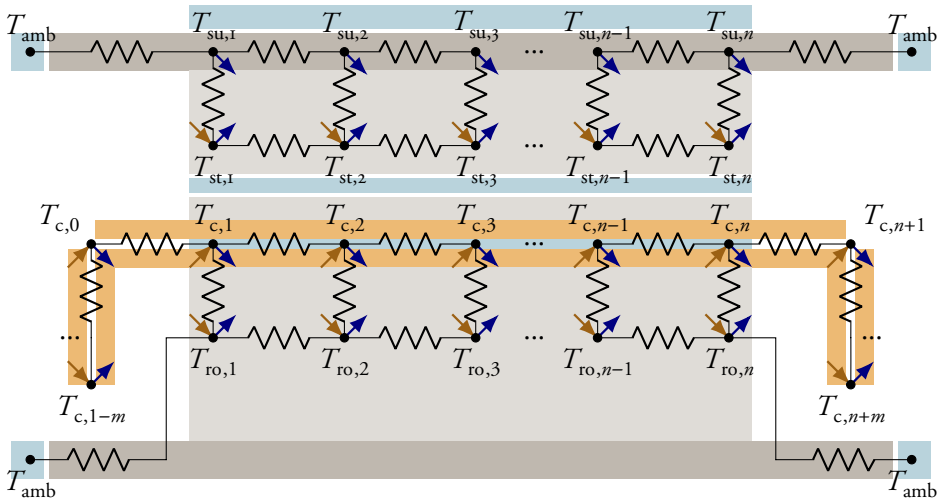
Figure 6.32: The thermal network(s) in the radial cross-section. Blue arrows mark heat dissipated through heat transfer coefficient (to air cooling or ambient air) and bronze arrows mark heat generated (by resistance or iron losses), but the direction of arrows does not indicate physical direction within the machine.

The housing nodes are connected with a thermal conductivity of 200 W/mK (aluminium). The first and last nodes of the stator surface and the rotor teeth have a thermal resistance coupled to a node with ambient temperature, to take into account the heat flowing via the housing and the shaft to the motor bench.

The thermal conductivity along the copper is 390 W/mK. The copper nodes continue on the the end turns with the same properties as along the active length of the machine, but with a natural convection heat transfer coefficient of 15 W/m²K, instead of the heat transfer coefficients from the forced convection model. The extra resistance of the bend of the end turn is added. The model for the resistance is explained in section below.

Copper resistance

The resistance of the copper in the model is based on the measured resistance of the whole rotor windings. However, knowing only that the total resistance is 46 mΩ at 117.5 °C is not enough to make an accurate resistance model. The measured temperature is the hotspot temperature and the temperature distribution along the windings is not known from the measurements. The resistance known is thus an average resistance over the windings at an unknown temperature distribution, but a known hotspot temperature. The problem with this is that a winding with a different temperature distribution, but with an equal hotspot temperature, will in practice have a different total resistance and the model needs to be able to take this into account.

The model produces a temperature distribution along the copper, as a result of the thermal network. This temperature distribution is mainly dependent on the cooling (coolant temperature, heat transfer coefficient and other heat fluxes in the machine), but also slightly dependent on the copper temperature itself. Since the resistance of each copper node is temperature dependent, the losses will increase wherever the temperature is higher, which will further increase the temperature. For this reason the model needs to estimate the local temperature dependent resistance along the windings.

To model this a temperature coefficient (k_T) must be assumed and used for the only known measured temperature ($R_{ref} = 46 \, m\Omega$ at $T_{ref} = 117.5 \, °C$) as shown in equation 6.2.

$$R_{total} = R_{ref}(1 + k_T(T_{hotspot} - T_{ref})) \tag{6.2}$$

This means that at a measured winding hotspot temperature $T_{hotspot}$, the total winding resistance is assumed to be R_{total}, given that k_T is accurate. This also means that the average local resistance along the windings can be calculated with equation 6.3.

$$R_{average \, per \, section} = R_{ref}(1 + k_T(T_{hotspot} - T_{ref}))/L_{tot} \tag{6.3}$$

where $R_{average \, per \, section}$ is the average resistance of each 1 mm section of the windings if L_{tot} is given in millimeters. With this equation the whole resistance of the windings are hotspot temperature dependent, but independent of the local temperature distribution along the windings.

In order to be able to input the local winding temperature rather than only the hotspot temperature, the whole equation can be shifted by $T_{hotspot} - T_{average}$ degrees, if $T_{hotspot}$ is the hotspot temperature from the model and $T_{average}$ is the average winding temperature from the model, as shown in equation 6.4.

$$R_{local} = R_{total}(1 + k_T(T_{local} - T_{ref} + (T_{hotspot} - T_{average})))/L_{tot} \tag{6.4}$$

With this equation a local temperature resistance can be calculated that will result in a total winding resistance of 46 mΩ at $T_{hotspot} = 117.5 \, °C$. However, how big this temperature dependency is depends on the value of k_T. This means that both the real local resistances in the reference case are still unknown, as well as the local and total resistances in any other case until k_T is calibrated.

An assumption could be to set k_T to the data table value of copper temperature dependence (of about 0.0039). However, this would not make the model accurate. A lower excitation current will not only result in a lower hotspot temperature but also in a flatter temperature

distribution and a smaller $T_{hotspot} - T_{average}$. In other words, the higher the average winding temperature gets, the larger will also $T_{hotspot} - T_{average}$ be. If the average winding resistance per millimeter was known for a specific average temperature, instead of hotspot temperature, a data sheet value for the temperature coefficient of copper resistance could be used accurately. But when the resistance is only known for a certain hotspot temperature, the function will have a built-in exaggeration of the temperature dependency.

This means that the temperature coefficient for the resistance function should have a lower value than the data sheet value (of about 0.0039). How much lower it should be can be derived from the model based on a few assumptions, but the uncertainty of those assumptions makes it more reasonable to use k_T as calibration parameter for the model. This way, k_T will be set to a value that makes the model as good as possible at predicting the thermal behavior of the prototype.

Core losses

The core losses are measured as described above, but those measurements provide no information of where in the core those losses are originating from. For this reason a series of simulations were run. These simulations show that 97-100 % of the core losses are originated in the stator. The simulations could also be used to establish how the core losses depend on the excitation current and the rotational speed. The function shown in equation 6.5 was found for the core losses p_{core}, using Matlab curve fit on the simulated core loss results, where a was then used to curve fit on the measured core losses to determine a value of $a = 0.0283$, ω is the rotational speed of the prototype (in rad/s) and I_r is the rotor current per turn.

$$p_{core} = a(0.004341\omega^2 + \omega)\sqrt{I_r} \qquad (6.5)$$

Since the purpose of the model is to describe the rotor winding temperature, a conservative value was chosen for the distribution of core losses. The model thus always assumes 2.5 % of the core losses are originated in the rotor teeth and 97.5 % in the stator teeth. The amount of total core losses are a function of rotational speed and excitation current as described above, but the fraction of core loss origin is always 2.5/97.5 %.

6.5.5 Iterations and output data of the model

The three main steps of the model (forced convection, updated airflow distribution and thermal networks) are repeated 40 times. This results in converging results.

A small extra step is made before the re-iteration, which is adding the heat from the end turn convection to the air temperatures. On the inlet side, the heat from the end turn convection is added to the temperature of the inlet air to the cooling channels (T_h for the hollow channels, T_p for the prism channels and T_a for the airgap). On the outlet side the heat from the end turn convection is added to the summed up total outlet air temperature.

The output air temperature is calculated by taking the average of the bulk temperatures of the channel types, weighted by the total mass flow in each channel type (and then adding the heat from the hot side end turn convection). The model then outputs the temperatures ($T_{h,p,a}$, T_{out}, $T_{c,ro,st,su}$) and heat flow rates (added heat from winding resistance and core losses, heat taken away by each channel type respectively, end turns and to ambient via surface, housing and shaft), scaled to the total of the prototype rather than the symmetric half pole used in the rest of the model.

6.5.6   Limitations and simplifications of the model

The model has some weaknesses.

One of them is that the model is a steady-state model and not a transient model. The experimental cases, however, has not been able to run until full steady-state is reached. The longer it runs, the vacuum cleaner providing the coolant air will keep on heating up, so if the inlet air temperature starts at 37 °C, it will be 40 °C 30 minutes later. On top of that, the ambient temperature is increasing slowly as the room heats up, and the motor bench itself will also keep heating up for a long time before it (if ever) reaches its equilibrium. So the cases of the steady-state model are based on the final stages (inlet air temperature, ambient temperature, etc) of the transient experimental case.

Another weakness is that the thermal model is not a fully accurate thermal model of the whole machine, but only for the parts that have an effect on the cooling of the windings. For instance, the model has no built-in support for adding stator winding currents, as they were never present in the experiments. Stator winding currents with losses could be added quite easily to the model, but the thermal paths for the losses would not be accurately modeled without adding complexity to the thermal networks.

The model is also not very accurate at estimating the housing, bearing or slip ring temperatures, since these are present in the model only to the extent that they improve the accuracy of modeling the winding temperatures. This applies also to how the thermal path from the rotor to (and away from) the shaft is modeled. It would be more correct to add at one row of thermal nodes along the shaft, thermally connected to both the rotor and copper nodes. However, this would cause higher complexity and more uncertain parameters, which would not improve the model accuracy.

The model only models one temperature for each cross section of the windings. It does not take into account the potential temperature differences between the inner or outer layers of the windings, or the individual variations in air inlet dimensions, etc. With the measurements available, modeling the copper temperature with so much extra complexity would be hard to verify and thus add little value of the model.

Some parameters in the model are hard to estimate. They are listed and described below.

- *The exact dimensions of the cooling channels*: The hollow channels and the airgap are assumed to be quite well represented as circular with radius 1 mm, and as a straight gap pf 0.75 mm, but the real dimensions of the prisms are harder to know. In the real prototype, they are not as perfectly triangular shaped as in the 3D CAD model, and they are also different from one another. The triangular shape is however kept throughout the model, but the dimensions of the triangles are allowed to be used as a calibration parameter.

- *The temperature coefficient* of the copper resistance, described in Equation 6.4.

- *Other material properties* of the machine which are not accurately known, such as the thermal resistance between the nodes in the thermal networks and the heat transfer coefficients to ambient air and from the end windings.

These parameters are kept within reasonable intervals, and are varied in the model until a good correlation throughout the experimental cases were met.

6.5.7 Model validation

The model was used to simulate all cases from the experimental tests on the prototype. The only inputs to the model are the rotor current, the airflow rate, the inlet air temperature, the ambient temperature and the rotational speed.

The results are shown in Figure 6.33. In the topmost graph, the temperatures are compared. The simulated temperatures are shown as rings and if the measured temperatures are centered in these rings, the model outputs the same results. In the bottom-most graph the difference between the modeled hotspot winding temperature and the measured temperatures (average of thermocouple, optical and/or converted IR camera data, depending on which are available for each case). A positive difference means the model overestimates the winding temperature.

In the middle graph the modeled winding losses plus core losses are shown and compared to the heat dissipated by the coolants. The difference (in the fourth graph) is equivalent to the
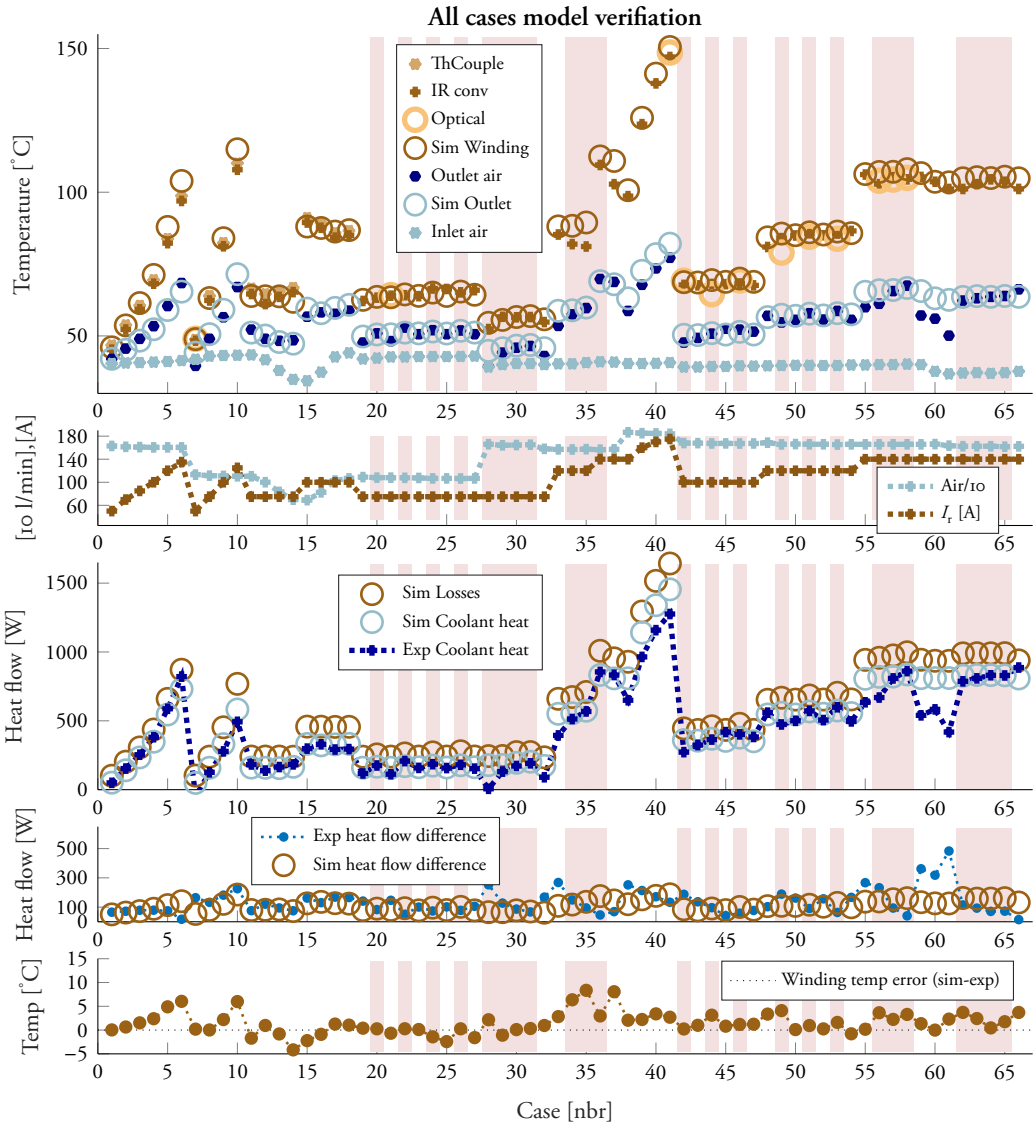
Figure 6.33: All cases from the experimental measurements and the model results from the same cases.

heat dissipated to ambient air or the motor bench. These data are based on the airflow and the difference between inlet and outlet temperature and as already concluded, the measured outlet temperatures are not completely consistent, which explains some of the variations between the model and experimental results. These results are not the main purpose of the model (namely correctly modelling the winding temperature), but they do say a lot about the accuracy of the thermal model. If the model correctly estimates the heat flow to

ambient air and the motor bench, the model can be better trusted to be correct also for cases that have not been validated against measurements. Considering the inconsistencies in the measured results overall and in the outlet air temperature specifically, full correlation can not be expected. 85 % of the cases have a difference between -10 % and +20 % difference, see histogram in Figure 6.34 (percentage of the total simulated losses), so for most of the cases the results are at least in the correct magnitude.
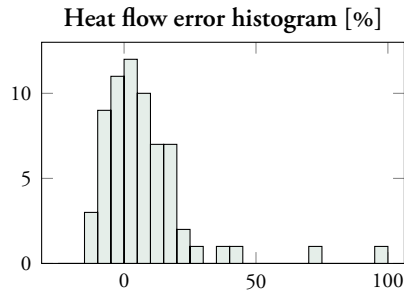
Figure 6.34: Histogram over the differences between model and measured data for heat dissipated into ambient air and motor bench. Positive value means model overestimates this heat flux.

In order to get an overview of the errors in the winding temperature, the results are plotted against different parameters in Fig 6.35. The first 7 graphs show the winding temperature error (modeled hotspot temperature minus the average of the measured temperatures for all cases). To clarify, in the first graph the winding temperature error is plotted against all measured air inlet temperatures. Most dots are places above 0 degrees, which means the model mostly overestimated the winding temperature. The largest negative error coincide with the lowest measured inlet temperatures. The blue line show a first degree curve fit of the data, showing that the errors in average tend to overestimate the winding temperature slightly more for higher air inlet temperatures.

This analysis has been made for all parameters. All parameters have a positive derivative of the first degree curve fit line. This means that the model overestimates the winding temperature more for higher rotor currents, higher airflow rates, higher rotational speeds, etc. This makes the model conservative and less likely to give an overoptimistic result of the cooling concept's ability to handle higher loads. Except this conservative tendency, there are no clear correlations to mention. Except for the cases with the lowest air inlet temperatures and airflow rates, no extreme values are gathered alone in one end or the other, but rather coincide with parameter values that have a majority of the cases in smaller error rates (For instance the highest extremes of the Winding temperature graph are situated at around 80 and 105 °C, for which temperatures there are many dots between 0 and 5 degrees). This is, of course, due to the fact that these results are the outcome of a calibration process, where the more unknown parameters of the prototype in the model have been adjusted to remove clear correlations and ensure that the model results are not too optimistic at higher loads.
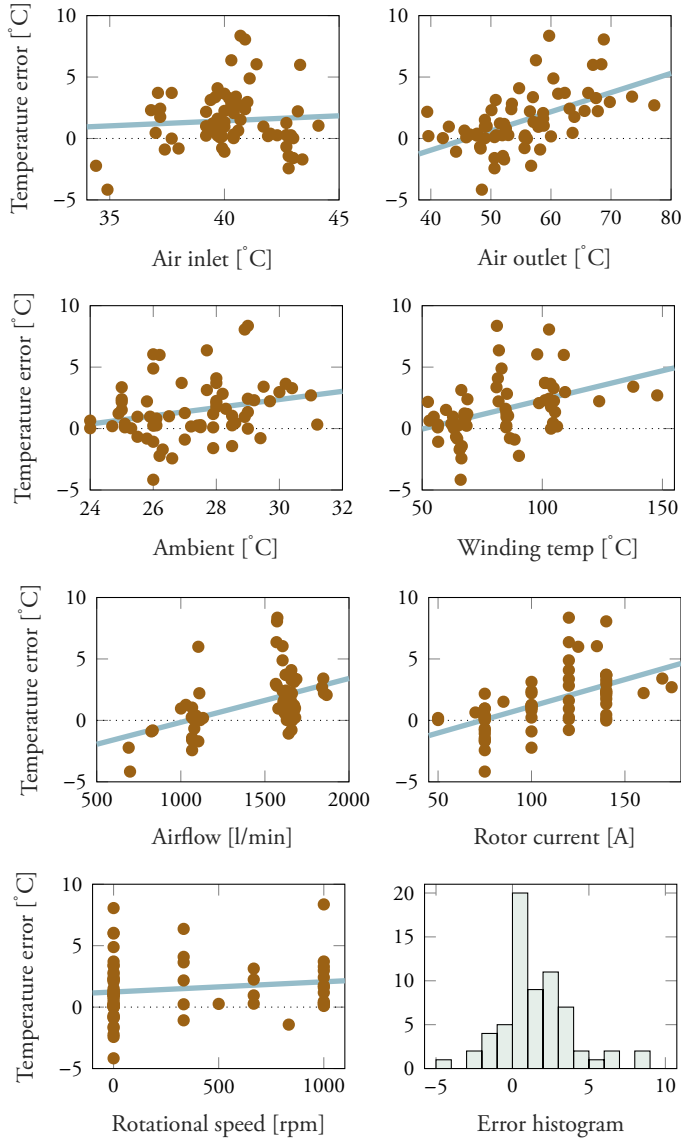
Figure 6.35: Correlation between temperature error (modeled temperature minus measured) and model parameters (as a function of the x-axis unit of each graph).

This process is explained more in the coming section.

The histogram in the final graph show that 71 % of all cases (47 out of 66 cases) result in the model estimating a winding temperature between 0 and 4 degrees higher than the measured result. The average error is 1.46 degrees and the standard deviation of the errors is 2.34 degrees.

The conclusion is that the model is able to successfully predict the thermal behavior of the machine to the extent that the model can be used to contribute to a better understanding of the cooling process. The model may also provide a reasonable prediction of cases that could not be tested experimentally, particularly cases not too far from the prototype experimental conditions.

Model calibration

Most parameters in the model were known and could be set to fixed values. Examples of this are the spacial dimensions (like the length of the machine) or well tested material properties (like air density in different pressure and temperatures, the thermal conductivity of copper, etc.). Other parameters were harder to estimate accurately. These parameters had to be calibrated from the experimental results.

One of these parameters is the shape and size of the prisms. They are all more or less individual in shape and size and change along the length of the machine as the outer turns switch level. An estimate of the dimensions from the CAD version was used as a base value and this value was then adjusted. In general, a higher value or larger prisms (like 100 % of the CAD dimensions of the prisms) makes the cooling less effective, increasing the winding temperatures, especially for higher rotor currents, but also spreads the temperature errors, resulting in a more spread out histogram. A lower value (like 90 %) will result in the model underestimating the winding temperatures. The value used in the final version is 95.5 % of the prism size according to the CAD model.

The temperature coefficient k_{T} of the resistance function (shown in equation 6.4) will tilt the temperatures for higher or lower rotor currents. The parameter value used in the final version of the model is 0.0026, which is lower than the table value of around 0.0039, but considering the resistance model has to take the winding temperature distribution into account, a lower value than the data table value is expected. A higher value than 0.0026 will make the derivatives of the winding temperature and rotor currents curve fits become higher than wanted.

Other unknown parameters are the heat transfer coefficients to the ambient air from the surface of the machine and from the end turn windings to the air inside the housing, as well as the thermal resistances from the housing and the shaft to the motor bench. These values have a huge effect on the thermal behavior of the model. The heat transfer coefficient from the housing surface to the ambient air was set to 20 $\mathrm{W/m^2K}$, which is a reasonable value considering the air is not completely still around the machine. The heat transfer coefficient from the end turns to the air inside the housing was set to 50 $\mathrm{W/m^2K}$, which is reasonable since the air is flowing abundantly there, but not as fast as inside the cooling channels. The motivation of these values are mostly that they are reasonable, and not mainly because of

the effect they have on the model results. These coefficients are used as constants and not as functions of the inlet airflow. Especially for the end turn heat transfer coefficient, this could have an effect on the accuracy of the model at higher and lower airflow rates than the ones used most in the prototype tests.

The thermal resistances to the motor bench from the stator housing and rotor shaft were set to an equivalent of 1 m length of housing walls and rotor shaft to large masses of metal of ambient temperature. This value was hard to predict, and has the effect of moving all temperatures up or down.

The main procedure of calibration was setting the heat transfer coefficients to reasonable values, then setting the resistance temperature coefficient to a value that gives an acceptable winding temperature curve fit derivative (not too high), and then setting the prism factor to a value as low as possible, gathering the histogram shape, but not so low so the housing and shaft thermal resistances had to be set unreasonably high without getting an average temperature error that was too low (or even negative).

6.6   Cooling concept assessment

The steady-state thermal model of the prototype can be used to better understand what happens in the prototype where no measurements have been made or in cases that could not be tested in the test bench used.

If nothing else is mentioned, the standard input parameters of the cases are shown in Table 6.3 (from which at least one parameter is varied in each of the following sections).

Table 6.3: Standard case parameters

Parameter		Value	
Rotor current	I_r	140	A
Airflow rate	$\dot{V}$	1600	l/min
Rotational speed	ω	0	rpm
Inlet air temperature	T_{in}	40	°C
Ambient temperature	T_{amb}	27	°C

6.6.1   Higher rotational speeds

One such case is higher speeds than the 1000 rpm the test bench allowed. This assumes that the model correctly estimates the core losses in higher speeds. The core loss model

of the steady-state thermal model is based on to things. First, it has the magnitude of the losses measured experimentally, within the region of speeds that was tested. Second, for rotational speeds higher than those tested it has its speed dependency derived from ANSYS simulations. The temperatures of the machine over some different rotational speeds are shown in Figure 6.36.
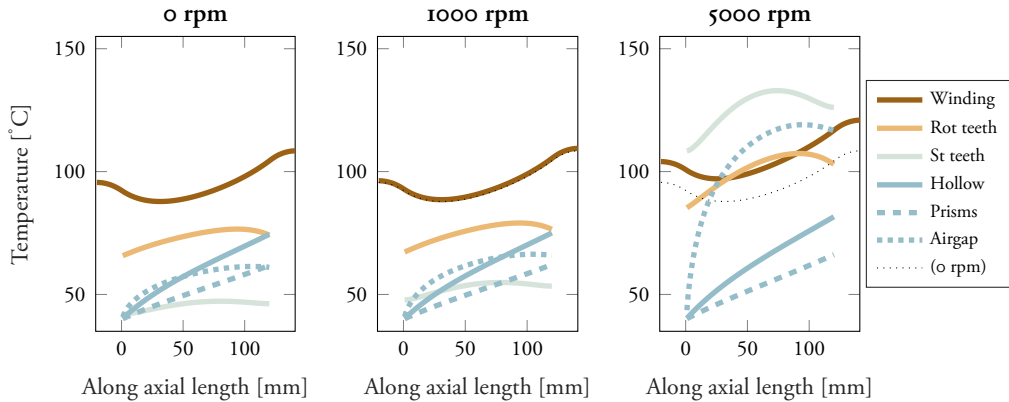
Figure 6.36: Model results of machine temperatures at different speeds. The winding temperatures of the 0 rpm case are shown as a black dotted line in the 1000 and 5000 rpm cases for comparison. The rotor current for all cases is 140 A.

From 0 to 1000 rpm the difference is very small. The winding temperature is only 1 degree higher. The most noticeable (although barely) difference is the temperatures of the stator teeth and the airgap coolant temperatures. In these first two cases it is apparent that the temperature of the rotor teeth is higher than that of the stator teeth, and the coolant in the airgap cools down the rotor teeth about as much as it heats up the stator teeth (the curve of the airgap air is between the rotor and stator teeth temperatures).

At 5000 rpm the difference is much bigger. The winding temperatures are 12 degrees hotter than in the 0 rpm case, but the stator teeth are even hotter. This is because in the prototype there is no cooling for the stator (other than the airgap). The original machine has stator cooling through water jackets in the housing, which would be needed to drive the machine with stator currents as well, and not just with rotor currents as in the testing of this prototype. This cooling is also designed to handle the core losses of the stator (and not just the stator winding losses), which is where almost all of the core losses are generated. In other words, simulating cases with 5000 rpm (or higher) is not relevant for a rotor cooling system other than to prove this point. If this rotor with direct air cooling is to be used in a real application, the stator cooling has to work too, either with original housing water jackets or (preferably) with direct cooled stator windings too. With this in place, the present cooling method for the rotor is enough to handle higher speeds, since the amount of core losses in the rotor are so low even at higher speeds (according to the model, 33 W in the whole rotor at 10 000 rpm with 100 A of rotor current) compared to the winding losses (over 1000 W

at high rotor currents), so their effect on the winding temperature is small.

6.6.2 Higher airflow rates

During the experimental tests, the highest airflow rat that could be achieved with the available air supply was around 1800 l/min. The practical limit for the airflow is higher than that. When a fluid reaches around Ma = 0.3 (30 % of the speed of sound in the medium) the behavior of the of the fluid starts to change so that the speed of the fluid will increase very little with further increments of the pressure. Increasing the heat transfer coefficient thus becomes very expensive energy-wise for air velocities above around 100 m/s. This corresponds to around 3200 l/min for the cooling channels in the prototype.
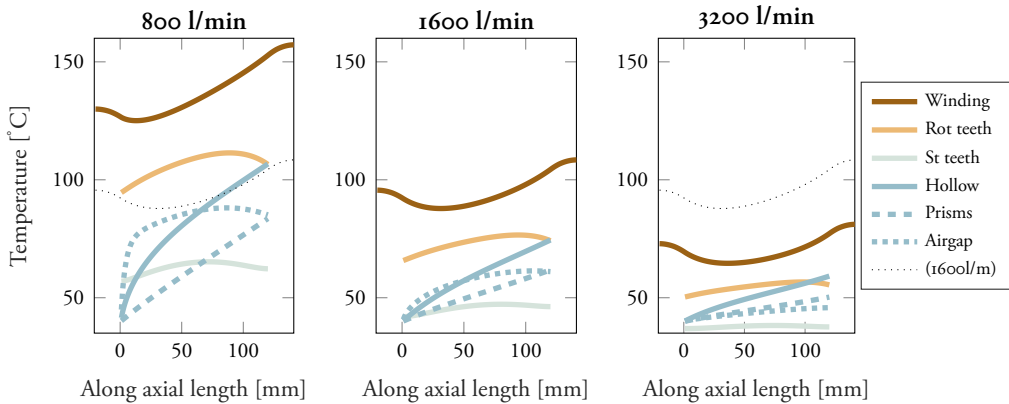
Figure 6.37: Model results of temperatures in the machine at different airflow rates. The rotor current is 140 A in all three cases. The winding temperatures of the 1600 l/min case are shown as a dotted black line in the other cases for comparison.

The model results of three different airflow rates are shown in Figure 6.37. Obviously, more airflow cools the machine to lower temperatures. A higher airflow helps the cooling in three ways. The first and more intuitive way is that higher air velocities means higher heat transfer coefficients. Higher air velocities also means there is less time for the coolant to heat up. In the 800 l/min case this is evident in the airgap airflow, which very quickly rises to a temperature centered between the rotor and stator teeth temperatures. The third way is that higher air velocities means higher inlet air pressure is needed, which means that the air will cool down more as the pressure drops along the cooling channel, keeping the coolant cooler.

It should be mentioned that all of these cases are simulated with the same (40 °C) inlet temperature. In a real application, the higher pressure needed for the higher airflow would also mean the air is heated up more in the compression, which would result in a higher inlet air temperature. It can be seen in the experimental results that the cases with lower

airflow rate also have a lower inlet air temperature. This is not taken into consideration in the input data of these results. However, it is clear that if a high airflow rate can be achieved, the current density of the windings could be increased further.

6.6.3   Higher rotor currents

If the goal of the cooling concept can be narrowed down to one single parameter it is the maximum continuous current enabled in the rotor slots. In Figure 6.38, the model results with various rotor currents are shown. The airflow rate is the most common airflow from the prototype tests, 1600 l/min. With this airflow rate, a rotor current of 172 A can be allowed without exceeding 160 °C in the windings. In the tests 175 A of rotor current was tested once, but with around 1800 l/min. At this point the winding hotspot temperature was measured to 148.7 °C (with the IR camera).
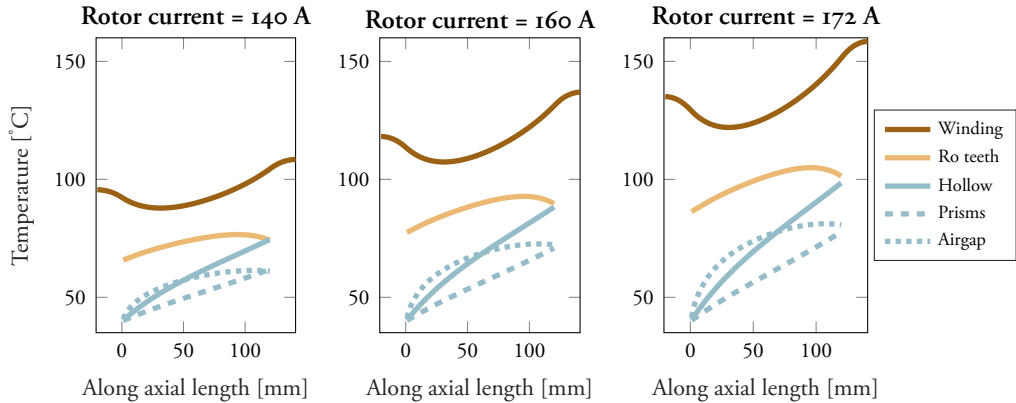
Figure 6.38: Model results of winding temperatures at various rotor currents with an airflow rate of 1600 l/min.

If a higher airflow rate could be used, higher rotor currents would be possible. An example with 2400 l/min is shown in Figure 6.39, enabling a maximum current of 197 A per turn. This corresponds to a copper current density 23.7 A/mm^2 and is higher than the rotor current predicted in Chapter 2.2.1, using a thermal model of the winding as thermally isolated from the rest of the machine and blowing up to 110 m/s through the cooling channels. With 2400 l/min, the air velocity is less than that (around 45 m/s in the hollow channels and 75 m/s in the prisms), but the heat is also allowed to dissipate in other ways from the windings than solely to the coolant. However, assuming the model is accurate for higher airflow rates (which previous analysis predicts that it is in a conservative way), and that an airflow rate of 2400 l/min is practically achievable, it is hereby confirmed that the cooling concept can indeed reach the performance target set in chapter 2.
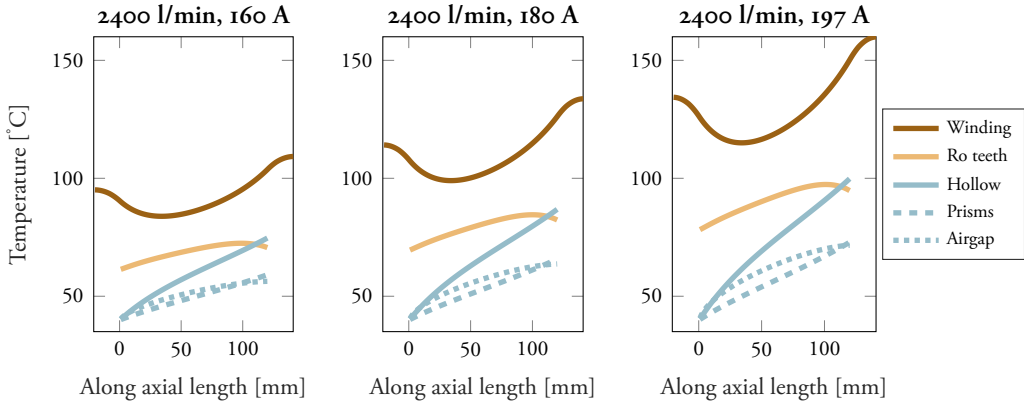
Figure 6.39: Model results of winding temperatures at various rotor currents with an airflow rate of 2400 l/min.

6.6.4   Sweep of rotor currents and airflow rates

For a wider range of rotor currents and airflow rates, Figure 6.40 is presented. It shows the hotspot winding temperature across 200-3200 l/min and 50-215 A. The maximum line shown is 155 °C, in order to take into account the imperfections of the measurements and the model (leaving a 5 degree margin to 160 °C for unpredicted temperature variations, a more developed steady-state and inaccuracies in the model). These results also assume 0 rpm (but with stator cooling in a real application, higher speeds are assumed to have little to no effect on the rotor winding temperatures), 40 °C inlet air temperature (which does not vary with the airflow rate in these simulations) and 27 °C ambient temperature. The maximum rotor current that can be reached in these circumstances is 213 A.

The bumps seen in the results at around 250 l/min and 800 l/min are due to the air in different cooling channels crossing an air velocity which makes the flow transition from laminar to turbulent (see next section).

6.6.5   Airflow distribution in cooling channels

In the assembly of the prototype it was decided that the prisms and the airgap would be left open for the air to flow through. This was done partly to save time by not having to solve the details about how the airgap and/or the prisms would be closed off, but also in order to be able to draw some conclusions about the cooling effect of different cooling channels.

The model can estimate how much of the total airflow that flows in each type of cooling channel. The distribution of the total airflow for different total airflow rates is shown in Figure 6.41. The shares show some peculiar behavior which needs the Reynolds number to
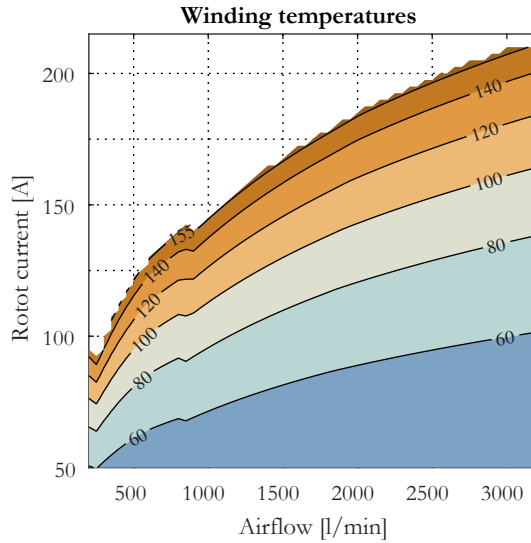
Figure 6.40: The hotspot winding temperatures at various rotor currents and airflow rates.

be explained. The model assumes the flow is laminar up to a Reynolds number of 2300. Between 2300 and 4000 the flow starts to develop into turbulent flow, which is fully developed in the model at a Reynolds number of 4000. The narrow channels of the hollow windings transition from laminar already at a low velocity (around 250 l/min) while the prism channels starts to transition at 850 l/min. This is the reason for why the airflow distribution change so peculiarly at these airflow rates.
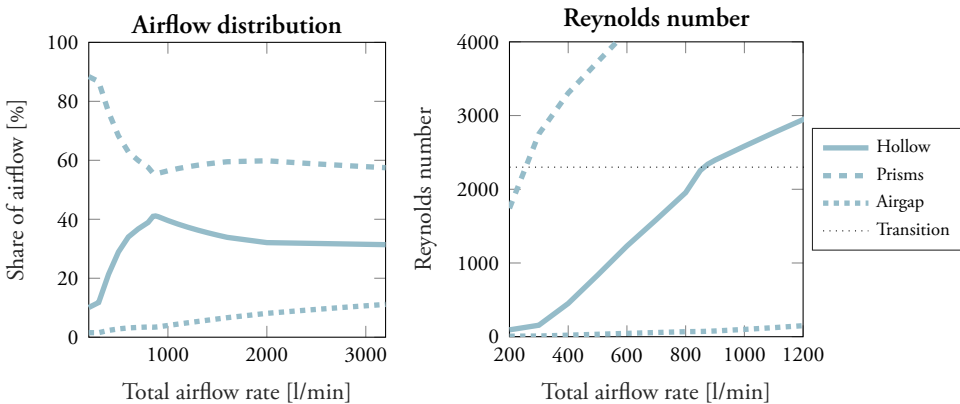
Figure 6.41: Airflow distribution for the air channels at different total airflow rates, along with the Reynolds numbers up to 1200 l/min.

Table 6.4 shows that the different cooling channels have similar cross section areas, but the size of each channel varies more. The hydraulic diameter is a way of representing a

channels shape and size, as a fraction between the cross section area and the perimeter of the channel, $D_h = 4A/p$. This number corresponds better than any other parameter for the channel dimensions how the airflow divides between the channels. It explains why so much of the total airflow ends up in the prism channels (57 % of the total airflow), despite all channel types having roughly the same cross section area. One way to put this is that large channels offer less friction for the air and can thus swallow more of the airflow than the more narrow channels. This lower friction also allow the air to flow faster through the channel, which in turn results in a higher heat transfer coefficient for the channel.

One crucial parameter determining the channel cooling remains to be examined. The wall area, or the total surface that creates an interface between the solid and the fluid in the cooling process. The larger prism channels can absorb a lot of air, but offer little contact between the air and the windings they are supposed to be cooling. The smaller hollow channels offer more wall area, so in the end, both the hollow channels and the prism channels dissipate an equal amount of heat from the windings. Compared to the total cross section area, the hollow channels are the most effective.

Table 6.4: Cooling channel comparison at 1600 l/min and 140 A

	Hollow	Prisms	Airgap	
Total cross section area	302	341	289	mm^2
Cross section area/channel	3.14	21.3	289	mm^2
Number of channels	96	16	1	-
Hydraulic diameter	2	3.8	0.75	mm
Airflow rate	542	952	106	l/min
Airflow percentage	34	60	7	%
Air velocity	32	49	6.5	m/s
Avg heat transfer coefficient	155	237	152	W/m^{2}K
Total wall area	603	325	770	mm^2
Heat dissipated	375	382	44	W
Dissipation/cross section area	1.24	1.12	0.152	W/mm^2
Dissipation/airflow	41.5	24.1	24.9	J/l

The case shown in Table 6.4 is the same base case as in previous sections, with 1600 l/min of airflow and 140 A of rotor current. This current generates resistive losses of 960 W, out of which 800 W are dissipated in the cooling channels. 22 W are dissipated into the ambient air, 53 W spread to the motor bench from the stator housing, 44 W spread to the motor bench via the rotor shaft and 41 W are dissipated from the end turns into the air of either side of the machine in the housing. The heat dissipation for the airgap channel is not representative with respect to its cooling effectiveness, since it is cooling the rotor teeth and heating up the stator teeth at the same time, moving more heat across the airgap air

(75 W) than out with it (44 W).

6.6.6   Closing off cooling channels

Closing off the cooling channels is easier to do in the model than in practice. Any cooling channel can be closed off in the model, and the total airflow will b distributed correctly over the remaining channels. Figure 6.42 shows two examples of this where only the prisms or the hollow channels are open, while the rest are sealed.
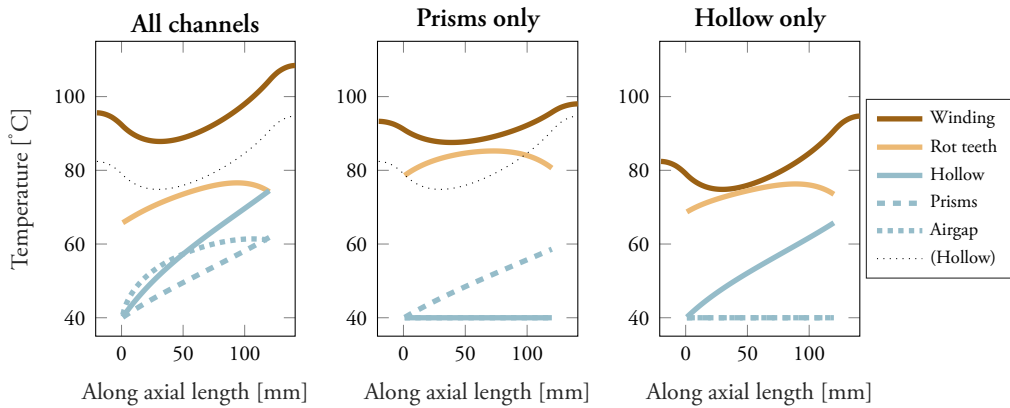
Figure 6.42: Model results with different cooling channels open. The first case has hollow, prism and airgap channels open for coolant, just like the prototype machine. The second case allows air only trough the prisms, filling in the hollow channels with copper to reduce resistance. The final case allows coolant only through the hollow windings.

Obviously, the cooling is more effective when concentrated to fewer channels. This is partly because the same airflow rate will be forced through fewer channels, resulting in higher air velocities and thus higher heat transfer coefficients (and getting rid of the thermal resistance to the airgap, if the airgap is closed off). The hotspot temperature of the windings is 108 °C for the case with all channels open (like the real prototype), 98 °C with only the prisms open and 95 °C with only the hollow channels open. This is despite the fact that the case with only prism cooling adds extra copper in the hollow channels and thus reducing the resistance and heat generated.

These results in particular prove that the point that the performance of an electrically excited rotor can be improved by removing copper in order to give space for direct cooling with channels that increase the heat transfer coefficient (hollow only gives 459 W/m^{2}K and prism only gives 372 W/m^{2}K for their respective cases) and increase the wall area.

6.7   Cooling concept possibilities

The prototype tested in this chapter achieved a rotor current of 175 A per turn, corresponding to a copper current density of 21.1 A/mm^2. With a higher available air supply (3200 l/min, which takes the airflow to the limit where compressibility effects can be neglected), the current could be increased to 213 A, corresponding to a copper current density of 25.6 A/mm^2. Closing off all other channels than the hollow windings makes the cooling more effective with the same coolant flow rate, but at around 2000 l/min the air velocity in the hollow channels reaches the limit. At this flow rate, the maximum allowable rotor current is 204 A (to keep windings below 155 °C according to model), which means that the additional cooling channels in the prototype will allow a higher maximum currents than with only the hollow windings open. The extra airflow (up to 3200 l/min) is not as effective at cooling the windings, but they push the limit where the effect of the compressibility of the air can no longer be neglected. All of this assumes the model predicts the thermal behavior of the machine at higher currents and airflow rates than tested experimentally, but the model was designed to give conservative results in these conditions, so there is reason to be optimistic about the outcomes of the prototype.
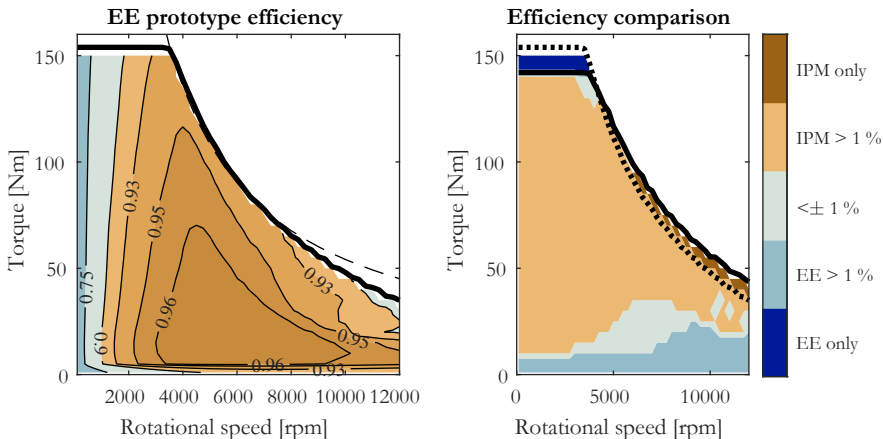
Figure 6.43: Left: Simulated efficiency for the final prototype design with maximum 213 A rotor current. Right: Efficiency and performance comparison between prototype and IPM rotor. Maximum torque for the IPM rotor is marked with a solid black line, while the maximum torque of the prototype is marked with a dashed line.

Electromagnetic ANSYS FEM simulations show that the performance of the prototype with a maximum rotor current of 213 A (Figure 6.43) is close to that of the original design (Figures 2.5 and 2.6). The current of 213 A is higher than the 195 A predicted in chapter 2, since the extra cooling channels and the heat dissipated to the ambient air and motor bench were not included in the initial thermal model. But the dovetail joint reduces the magnetic flux slightly, which basically balances out the advantage of the higher current. The simulated maximum torque of the prototype is 154 Nm, compared to 156 for the initial

design. The power output at higher speeds is comparable, but the efficiency is slightly poorer for the prototype.

Compared to the IPM rotor from the initial study (right in Figure 6.43), the prototype reaches a higher torque (142 Nm for the IPM rotor) but a lower power output at higher speeds. The IPM is still more efficient for most operating points, but the prototype is more efficient at lower power and high speed, and comparable to the IPM in a small region at the high torque and base speed (This region is smaller than for the initial EE design).

6.8   Chapter summary

This chapter presents the experimental setup and the results of the testing of the prototype. The delimitations from the experimental setup are considered in the test plan which is then carried out successfully. Some measured electromagnetic aspects are measured and compared with simulated results.

The thermal measurements of the prototype conclude that the rotor windings can maintain a rotor current of 175 A continuously with a hotspot winding temperature below 150 °C (corresponding to a copper current density of 21.1 A/mm^2).

A steady-state thermal model is designed and calibrated to the prototype in order to better understand the thermal behavior of the prototype. The outcomes from the model of the prototype are used to estimate the maximum rotor current to 213 A with a higher airflow rate than the rate used in the experimental tests. The model also predicts that with a limited airflow, allowing the air to pass only through the hollow channels would provide better cooling than allowing it in the airgap and the prism shaped areas between the coils as well. However, allowing a higher airflow rate will result in the cooling effects to saturate in the hollow channels, while a more effective cooling is achieved with airflow paths as in the prototype (resulting in a maximum current per turn of 213 A as mentioned above).

With these results the prototype, with a higher airflow rate than used in the experimental tests, simulations indicate that the final prototype almost matches the performance predicted with the initial EE rotor design in Chapter 2 with just slightly lower maximum continuous torque and efficiency. Compared to the equivalent IPM rotor the maximum continuous torque as well as efficiency in some high speed operating points are higher for the prototype.

Rotational speeds above 1000 rpm were not possible in the test setup, so the behavior of the air flow at higher speeds are not assessed in this work. Except for this aspect, the cooling concept with hollow direct air cooled windings for an EE rotor is verified to work in practice.

Chapter 7

Conclusions and topics for future research

The aim of this thesis is to validate whether the concept of hollow direct air cooled windings work in practice. This has been done through background studies, and a design process aimed at manufacturability, before manufacturing and testing the prototype. The progress of the work is described in this thesis.

7.1 Discussion

The work has been a process with many steps of struggle. The prototype is built, the cooling concept is validated and the outcome is a success, confirming the predictions. The cooling concept does work, it can be realized in a real machine and it makes the machine reach performance levels which proves that the concept very promising for further development.

Air cooling is a slightly neglected topic for thermal management of electric machines for EV traction. It is understandable since the cooling capabilities of other coolants like water and oil are so much better. However, as shown in chapter 2, direct air cooling can be as effective as water jacket cooling or end turn oil-spray cooling. In particular, direct air cooling may be the most effective concept for the cooling of rotor windings in EE machines. However, prior to this work there was no proof of concept in practice to be found in literature. This work has aimed at, and succeeded to verify that this cooling concept works in practice.

EE machines are important because of the world's need of the transportation electrification transition being successful. Right now over 80 % of the EV market relies on REE PMs for that transition while the market is growing rapidly. The EE machine is one of the few

PM-less topologies that can match the performance of PM machines for EV traction while reducing the use of critical raw materials such as REEs. Direct air cooling can strengthen the performance of EE machines, making it a stronger competitor on the EV market and can thus potentially relieve the EV market from risky raw material dependencies.

There are several ways direct air cooled rotor windings can be implemented, such as laminated windings, cooling tubes in the winding slots, additive manufactured coils or hollow windings. Hollow windings have been showed to be promising in theory (in the Licentiate thesis, [9]) and also seem to solve some of the practical problems of laminated windings (as discussed in Chapter 1.2.2). Still, building a prototype with hollow windings faced several challenges along the work.

The manufacturability aspect the rotor design is one such challenge. It was evident early in the work that winding the hollow conductors directly onto a conventional EE rotor (such as rotor design option 1 in Chapter 3) is not likely to render in a satisfying result. Instead a way of allowing the installing of prewound coils is needed, and for that a dovetail joint was deemed to be the way forward for this project. To design a dovetail joint with satisfying mechanical and electromagnetic properties is a challenge in itself. Furthermore, machining and assembly of the joints presented another challenge. The dovetail joint that this work produced meets the mechanical needs of the prototype testing, but for a product in industry more development work is likely needed.

The testing of the prototype involves the biggest weakness in the verification of the cooling concept, due to the limited range of rotational speeds that was possible to achieve in the motor bench. This means there are not theoretical or practical attempts to show how the fluid dynamics of higher rotational speeds may affect to airflow.

Apart from that, the cooling concept is shown to work as predicted. The model and simulations behind the initial performance predictions (of a maximum current of 195 A per turn with a winding hotspot temperature below 160 °C and the efficiency map in Figure 2.5 has two major differences from the prototype. In the initial model and simulations: 1) The windings are modeled as thermally isolated from the rest of the machine which makes the thermal predictions conservative; 2) No dovetail joints are present which makes the performance simulations optimistic. Experimentally 175 A per turn was reached with a winding hotspot temperature below 150 °C and a lower airflow velocity than in the initial prediction. The thermal model estimates a maximum current of 213 A per turn is possible with the same airflow velocity as in the initial prediction, which is higher than the conservative figure of 195 A. The resulting efficiency map (Figure 6.43) has a higher maximum torque which is expected due to the higher maximum rotor current, and a lower efficiency which is expected due to the added dovetail joints. With these differences considered the cooling concept can thus be said to perform as expected.

The maximum rotor current of 213 A per turn means that the simulated results of the

EESM can produce more torque at low speeds than an IPM designed according to the original specifications of the stator used. The IPM machine has a higher efficiency in most operating points and a higher maximum power at high speeds, but there is an operating area at high speed and low power where the EESM is more efficient. This means vehicles and drive cycles with more of its operating points in this area can be overall more efficient with an EESM than an IPMSM. This result is true for the two specific machines compared, but it should be noted that this comparison is made with a stator originally designed for an IPM rotor with this specific airgap flux density. The EESM produces more linked magnetic flux which means that it would likely perform better in a stator with a different slot/core ratio than in the stator used. If this comparison were instead made with that hypothetical stator instead, it is likely the EE would outperform the IPMSM in even more operating conditions than in this comparison.

The experimentally validated thermal model allows a deeper insight into the thermal behavior of the cooling concept as it is implemented in this work, and can also open doors to other ways of implementing it. It does confirm that pure direct cooling is better if the goal is to aspire for higher levels of performance (when it comes to torque and power).

7.2 Conclusions

Summed below are the most important scientific conclusions from the work:

- End turn cooling with oil spray, housing jacket with water or direct cooling in hollow windings with air have similar potentials to cool a generic winding effectively, although the specific application will weigh in on which cooling concept that is more suitable.

- Hollow windings can be used with different coolants, such as water, oil and air, with their different pros and cons. When it is done, the optimal number of parallel and serial flows are different for each coolant, where a minimum number of parallel flows is more efficient for water, a maximum number of parallel flows is more efficient for air.

- Direct cooling can potentially make up for the increased losses due to copper being sacrificed for cooling channels by reducing the winding temperature so much that the total resistance becomes lower compared to conventional cooling methods.

- Dovetail joints can be used to allow the installation of prewound coils on an EE rotor with limited electromagnetic and mechanical effects. This conclusion includes the finding that wire discharge machining can be used to cut the dovetail joints in laminated steel. The measured induced voltage indicates that the dovetail joint adds

an equivalent airgap of 0.01 mm. A tensile test is used to validate the mechanical simulations which show the contact region is best represented with a friction factor of 0.35.

- Measuring the temperatures on the rotating rotor is done with an IR camera. Despite the reflectivity of the windings, the IR camera readings are shown to be consistent with measurements with thermocouple sensors and an optical thermal sensor on a static rotor.

- As expected, the experimental measurements show that a higher current results in higher winding hotspot temperatures, while a higher airflow rate decreases the temperatures. Within the span of rotational speeds that were tested, no clear correlation between rotational speed and temperatures or cooling capabilities were detected. The maximum current per turn tested and maintained continuously was 175 A, corresponding to a copper current density of 21.1 A/mm^2, with an airflow rate of around 1800 l/min. With this current the winding hotspot temperature was measured to be below 150 °C.

- A thermal model is calibrated to the experimental results and with this validated model, the thermal behavior of the prototype is better understood. This model is used to draw a number of conclusions.

  - The prototype windings can maintain an excitation current of 213 A per turn, which corresponds to a copper current density of 25.6 A/mm^2, given that a coolant airflow rate of 3200 l/min can be provided.

  - If the airgap and prism channels were to be blocked from coolant airflow, only allowing airflow inside the hollow windings, the cooling would be more efficient, allowing a higher maximum current for a total airflow rate of 1600 l/min. However, at higher airflow rates it is more effective to have all cooling channels open, maximizing the current at 213 A, compared to a maximum current of 204 A with air only in the hollow conductors.

  - About equal amounts of heat is dissipated from the air going through the hollow conductors compared with the air in the prism channels. The channels of the prisms allow almost twice as much airflow compared to the hollow channels in total, which means the hollow channels dissipate 70 % more heat per liter air. This is consistent with initial model results (Figure 2.18) showing that larger cooling channels are less efficient than the ones used in the prototype efficient than larger channels. However, the larger channels may be more effective, allowing a higher maximum current density, which is also consistent with the conclusion above. Since the prototype was tested without operation in the stator windings (due to lack of stator cooling), the airflow in the airgap

mostly moved air from the rotor to the stator rather than dissipating air from the rotor to the air outlet.

- The results above are used for simulations of the total operation of the prototype and compared with a PM rotor. The machine simulated with a PM rotor is overall more efficient than the EE rotor and can produce higher powers at high rotational speeds, but the EE rotor produces a higher maximum torque at lower speeds and a higher efficiency at high speed-low power operating points. This means that for vehicles working mainly in this operating area, such as heavy duty trucks spending most of the time at high speeds, the direct air cooled EE rotor machine can be more efficient than PMSMs.

7.3 Topics for future research

The following topics have been identified during the work as important for future research:

- **Further evaluation of the prototype**: For the continued evaluation of the prototype itself, more insight can be gained by testing the machine at higher rotational speeds, higher airflow rates and with some stator currents (but the lack of stator cooling in the prototype limits the latter). More advanced fluid dynamics models can make better predictions about the airflow at higher speeds, in supplement with the experimental tests. For this work other temperatures measurement methods that enables continuous temperature readings on a rotating rotor, such as onboard thermocouple sensors or laser based temperature measurements, would be suitable.

- **Cooling concept**: The cooling concept itself should be further studied and developed. One step is to design a whole machine with direct cooling, combining stator and rotor cooling. Another step is to investigate the feasibility of additive manufacturing, since a more free design of the coils will open up huge potentials when it comes to maximizing the cooling effectiveness. Yet another step could be to study improving or boosting the air cooling with other gases such as vaporized water.

- **Dovetail joints**: The dovetail joints can also be studied further. Their structural integrity can certainly be improved further design-wise, The long-term durability can also be investigated further along with the settling of a safety factor. The choice of machining method can also be in need of a look in order to reduce manufacturing time and/or costs.

- **Towards transport electrification**: But most importantly, the strive should be continued towards the transport electrification transition. In this strive several of the conclusions of this work can be used:

– Dovetail joints can be used to enable EE rotors instead of PM rotors by allowing the installation of prewound direct cooled EE windings, but also to enable modular electric machines (EMs) so that individual (or all) windings can be exchanged simply at the end of their lifetime without having to scrap the whole machine.

– Direct air cooling with hollow conductors can be used to improve the thermal management of traction machines (stator and/or rotor), which can in turn reduce the use of materials in total due to increased power density overall.

– Direct air cooling can also be used to strengthen the use of more EESMs in the EV market instead of REE PMs which will reduce the market's dependency on a singly critical raw material along with reducing the environmental footprint that is connected to the extraction of REE extraction or PM manufacturing. If EESMs are used in the right type of vehicles, or in combination with other topologies so that the EESM is used in the operating points where it has the best efficiency, then the total energy consumption over the vehicle's lifetime can also be reduced, leading to less environmental impact from the energy production as well as lower operation costs.

– Aluminium windings can potentially also be used with lower costs and higher efficiency than copper windings. Direct cooling can potentially help enable the use of aluminium windings and should definitely be studied further. This would decrease the environmental impact of the machines and further reduce the EV market's dependency on critical raw materials.

Nomenclature

List of Symbols

p_{fric} Pipe friction losses [W]   p. 33

p_{res} Resistive loss [W]   p. 5, 141, 144

Pr Prandtl number [-]   p. 28, 29

$\dot{Q}$ Heat flux [W]   p. 4, 19, 20, 24, 141, 144

R Resistance [Ω]   p. 6, 24, 37, 63, 146

r Radius [m]   p. 27, 71

Re Reynolds number [-]   p. 28, 30

T Temperature [K, ˚C]   p. 4, 19, 20, 24, 37, 141, 145–148

T_{a} Air bulk temperatures in airgap [K, ˚C]   p. 142, 143, 148

T_{amb} Ambient temperature [K, ˚C]   p. 141, 142, 145, 154

T_{b} Air (bulk) temperature [K, ˚C]   p. 4–6

T_{c} Copper (winding) temperature [K, ˚C]   p. 4–6, 24, 141–145

T_{h} Air bulk temperatures in hollow channels [K, ˚C]   p. 142, 143, 148

T_{in} Coolant inlet temperature [K, ˚C]   p. 141, 142, 154

T_{p} Air bulk temperatures in prism channels [K, ˚C]   p. 142, 143, 148

T_{ro} Rotor tooth temperature [K, ˚C]   p. 142, 144, 145

T_{st} Stator tooth temperature [K, ˚C]   p. 142, 144

T_{su} Stator surface temperature [K, ˚C]   p. 142, 144

u Coolant flow velocity [m/s]   p. 4, 6, 28–30, 141

$\dot{V}$ Coolant volume flow [l/min,m^3/s]   p. 33, 141, 142, 154

λ Thermal conductivity [W/mK]   p. 19, 20, 24, 27, 29, 30

λ_{c} Thermal conductivity of copper/conductor [W/mK]   p. 5

ν Dynamic viscosity [mPa·s]   p. 28–30

ω Angular velocity [rad/s]   p. 71, 141, 142, 147, 154

ρ Density [kg/m^3]   p. 28–30

Notes on symbols: Symbols may be preceded with d, δ or Δ to mark a part or change of the variable. Indices may contain numbers or variables ($_{1,2,n,m}$) to mark the number of an element where n or m is the number of elements in the set. Indices may also contain entire words describing them, in which case they are not listed here. Symbols containing multiple letters are written as upright text (such as Pr, Re, etc) in order to not confuse them with single letter variables (such as $Pr = P \cdot r$).

The following indices always have the same meaning when they appear on symbols without being listed explicitly above: S_a airgap, S_c copper/conductor, S_h hollow, S_p prism, S_{ro} rotor tooth, S_{st} stator tooth, S_{su} stator surface.

Acronyms

BEV battery EV   p. 1, 2

EE electrically excited   p. 9, 12, 13, 15–18, 26, 40, 55, 67, 70, 84, 107, 110, 112, 122, 163, 165–167, 169, 170

EESM electrically excited synchronous machine   p. 3–8, 16, 44, 167, 170

EM electric machine   p. 170

EMSM electrically magnetized synchronous machine   p. 3

EV electric vehicle   p. 1–3, 8, 18–20, 70, 84, 165, 166, 170

FEM finite element model   p. 6, 162

IEA Industrial electrical engineering and automation   p. 3

IM induction machine   p. 3

IPM interior permanent magnet   p. 5, 6, 8, 12, 13, 15–18, 40, 122, 162, 163, 167

IPMSM interior permanent magnet synchronous machine   p. 3, 16, 167

PHEV plug-in hybrid EV   p. 1, 2

PM permanent magnet   p. 1, 2, 7, 9, 13, 70, 115, 165, 166, 169, 170

PMSM permanent magnetized synchronous machine   p. 115, 169

REE rare earth element   p. 1, 2, 165, 166, 170

SMC soft magnetic composite   p. 49, 52, 55

Glossary

current density *Current density* means how the average current through a certain cross section area.

For *copper current density* this cross section area is the copper area of a conductor (8.31 mm^2 for the conductors used in this work).

For *slot current density* the cross section area is the whole inner area of a slot, including slot liner and any unfilled area (123.3 mm^2 for a half rotor slots of the prototype). For *conductor current density* the cross section area includes whole rectangular area that the conductor takes up, including the inner channel and area outside the rounded corners (12 mm^2 for the conductors used in this work).   p. 4, 7, 14, 16, 21, 32–37, 40, 157, 162, 163, 168

fill factor *Fill factor* generally means the ratio of copper that fills a certain cross-section area. However, in this work fill factor refers specifically to the area that is not taken up by the inner cooling channel compared to the total area. This means the term does not take varnish thickness, rounded corners or how tightly packed the conductors are in a slot, but is used exclusively when comparing the effect of different dimensions of the inner cooling channel.   p. 4, 32–35, 41

mortise *Mortise* refers to one of the parts in a dovetail joint, namely the hole, socket or groove into which the tenon is inserted. For woodwork applications the Swedish translation is tapphål. Sometimes the mortise is called the female part of the joint. p. 76–79

pipe friction losses *Pipe friction losses* in this work refers to the product of multiplying the pressure drop across a channel and the volume flow rate passing through it. This power is sometimes referred to as pump power (and is used that way in the licentiate thesis [9]), but since it may be interpreted as the power the pump producing the flow needs the term is changed to pipe friction losses instead. It does correspond to the output power of a pump, but the input power to the pump is higher due to the efficiency of the pump.   p. 31, 33–37, 40

tenon *Tenon* refers to one of the parts in a dovetail joint, namely the protruding part which is inserted into the mortise. For woodwork applications the Swedish translation is tapp. Sometimes the tenon is called the male part of the joint.   p. 76–80, 83, 91

References

[1] International Energy Agency. Global EV outlook 2024 / Data explorer, April 2024 (accessed on 2024-04-25). https://www.iea.org/data-and-statistics/data-tools/global-ev-data-explorer.

[2] J. Edmondson and S. Siddiqi. Electric motors for electric vehicles 2024-2034. Technical report, IDTechEx, June 2023 (accessed on 2023-09-15). https://www.idtechex.com/en/research-report/electric-motors-for-electric-vehicles-2024-2034/941.

[3] J. Marx, A. Schreiber, P. Zapp, and F. Walachowicz. Comparative life cycle assessment of ndfeb permanent magnet production from different rare earth deposits. *ACS Sustainable Chemistry & Engineering*, 6(5):5858–5867, 2018.

[4] U.S. Geological Survey. Mineral commodity summaries, January 2024 (accessed on 2024-05-02). https://pubs.usgs.gov/periodicals/mcs2024/mcs2024.pdf.

[5] Adamas Intelligence. Rare earth magnet market outlook to 2040, April 2023 (accessed on 2023-09-26). https://www.adamasintel.com/report/rare-earth-magnet-market-outlook-to-2040/.

[6] European Commission. Critical raw materials: ensuring secure and sustainable supply chains for eu's green and digital future, March 2023 (accessed on 2023-09-19). https://ec.europa.eu/commission/presscorner/detail/en/ip_23_1661/.

[7] European Commission, Entrepreneurship Directorate-General for Internal Market, Industry, SMEs, M Grohol, and C Veeh. *Study on the critical raw materials for the EU 2023 – Final report*. Publications Office of the European Union, 2023 (accessed on 2023-09-27). https://data.europa.eu/doi/10.2873/725585.

[8] Carrara S, Bobba S, Blagoeva D, Alves Dias P, Cavalli A, Georgitzikis K, Grohol M, Itul A, Kuzov T, Latunussa C, Lyons L, Malano G, Maury T, Prior Arce A, Somers J, Telsnig T, Veeh C, Wittmer D, Black C, Pennington D, and Christou M. Supply chain analysis and material demand forecast in strategic technologies and sectors in the EU? A foresight study. Policy assessment, Anticipation and foresight, Risk assessment KJ-NA-31-437-EN-N (online),KJ-NA-31-437-EN-C (print), Luxembourg (Luxembourg), 2023.

[9] S. Estenlund. *Air cooling of an EMSM field winding*. Licentiate thesis, Lund University, 2018.

[10] C. Högmark, R. Andersson, A. Reinap, and M. Alaküla. Electrical machines with laminated winding for hybrid vehicle applications. In *2012 2nd International Electric Drives Production Conference (EDPC)*, pages 1–6, Oct 2012.

[11] R. Andersson, C. Högmark, A. Reinap, and M. Alaküla. Modular three-phase machines with laminated winding for hybrid vehicle applications. In *2012 2nd International Electric Drives Production Conference (EDPC)*, pages 1–7, Oct 2012.

[12] A. Reinap, F. J. Márquez-Fernández, R. Andersson, C. Högmark, M. Alaküla, and A. Göransson. Heat transfer analysis of a traction machine with directly cooled laminated windings. In *2014 4th International Electric Drives Production Conference (EDPC)*, pages 1–7, Sept 2014.

[13] R. Andersson, A. Reinap, and M. Alaküla. Design and evaluation of electrical machine for parallel hybrid drive for heavy vehicles. In *2012 XXth International Conference on Electrical Machines*, pages 2622–2628, Sept 2012.

[14] R. Andersson and A. Reinap. Loss mapping of an insert permanent magnets synchronous machine for parallel hybrid electric heavy vehicles. In *2016 XXII International Conference on Electrical Machines (ICEM)*, pages 1847–1853, Sept 2016.

[15] Luvata. Hollow conductors, 2023 (accessed on 2024-05-02). https://www.luvata.com/en/Products/Hollow-Conductors/.

[16] S. Estenlund. Model of air cooled windings for traction machine. In *2018 XIII International Conference on Electrical Machines (ICEM)*, pages 1278–1284, Sep. 2018.

[17] International Electrotechnical Commission. IEC 60172, 2020 (accessed on 2024-03-13). https://webstore.iec.ch/publication/68059.

[18] Dahrén Group. Product data sheets, 2023 (accessed on 2024-03-13). https://www.dahrengroup.com/documents/product-data-sheets/.

[19] Zhe Huang, Francisco J. Márquez-Fernández, Yury Loayza, Avo Reinap, and Mats Alaküla. Dynamic thermal modeling and application of electrical machine in hybrid drives. In *2014 International Conference on Electrical Machines (ICEM)*, pages 2158–2164, 2014.

[20] E. L. Brancato. Estimation of lifetime expectancies of motors. *IEEE Electrical Insulation Magazine*, 8(3):5–13, May 1992.

[21] W.Q. Chu, Z.Q. Zhu, J. Zhang, X. Ge, X. Liu, D. Stone, and M. Foster. Comparison of electrically excited and interior permanent magnet machines for hybrid electric vehicle application. In *Electrical Machines and Systems (ICEMS), 2014 17th International Conference on*, pages 401–407, Oct 2014.

[22] L. Cuiping, G. Zhengwei, L. Junhui, Z. Bing, and D. Xiucui. Optimal design of cooling system for water cooling motor used for mini electric vehicle. In *2017 20th International Conference on Electrical Machines and Systems (ICEMS)*, pages 1–4, 2017.

[23] V. Josefsson, A. Carlsson, S. Nategh, and D. Ekholm. An experimental investigation of hybrid cooling solution for high performance traction motor. In *IECON 2022 – 48th Annual Conference of the IEEE Industrial Electronics Society*, pages 1–6, 2022.

[24] P. Liang, F. Chai, K. Shen, and W. Liu. Water jacket and slot optimization of a water-cooling permanent magnet synchronous in-wheel motor. *IEEE Transactions on Industry Applications*, 57(3):2431–2439, 2021.

[25] R. Pechánek and L. Bouzek. Analyzing of two types water cooling electric motors using computational fluid dynamics. In *2012 15th International Power Electronics and Motion Control Conference (EPE/PEMC)*, pages LS2e.4–1–LS2e.4–5, 2012.

[26] S. Wu, D. Hao, and W. Tong. Cooling system design and thermal analysis of modular stator hybrid excitation synchronous motor. *CES Transactions on Electrical Machines and Systems*, 6(3):241–251, 2022.

[27] A. Acquaviva, O. Wallmark, E. A. Grunditz, S. T. Lundmark, and T. Thiringer. Computationally efficient modeling of electrical machines with cooling jacket. *IEEE Transactions on Transportation Electrification*, 5(3):618–629, Sep. 2019.

[28] Z. Huang, S. Nategh, V. Lassila, M. Alaküla, and J. Yuan. Direct oil cooling of traction motors in hybrid drives. In *2012 IEEE International Electric Vehicle Conference*, pages 1–8, 2012.

[29] H. Li, K. W. Klontz, V. E. Ferrell, and D. Barber. Thermal models and electrical machine performance improvement using encapsulation material. *IEEE Transactions on Industry Applications*, 53(2):1063–1069, 2017.

[30] P. S. Ghahfarokhi, A. Podgornovs, A. Kallaste, T. Vaimann, A. Belahcen, and A. J. Marques Cardoso. Oil spray cooling with hairpin windings in high-performance electric vehicle motors. In *2021 28th International Workshop on Electric Drives: Improving Reliability of Electric Drives (IWED)*, pages 1–5, 2021.

[31] C. Liu, Z. Xu, D. Gerada, F. Zhang, Y. C. Chong, M. Michon, J. Goss, C. Gerada, and H. Zhang. Experimental investigation of oil jet cooling in electrical machines with hairpin windings. *IEEE Transactions on Transportation Electrification*, 9(1):598–608, 2023.

[32] A. La Rocca, A. Fregni, S. La Rocca, and C. Gerada. Numerical thermal modelling of multiphase spray cooling of hairpin windings. In *2020 International Conference on Electrotechnical Complexes and Systems (ICOECS)*, pages 1–5, 2020.

[33] V. Madonna, A. Walker, P. Giangrande, G. Serra, C. Gerada, and M. Galea. Improved thermal management and analysis for stator end-windings of electrical machines. *IEEE Transactions on Industrial Electronics*, 66(7):5057–5069, July 2019.

[34] E. Fairall, C. Rheberhegen, E. Rowan, J. Lo, B. Bilgin, and A. Emadi. Maximizing thermal effectiveness and minimizing parasitic loss in a liquid cooled switched reluctance machine. In *2016 IEEE Transportation Electrification Conference and Expo (ITEC)*, pages 1–7, 2016.

[35] Payam Shams G., A. Podgornovs, A. Kallaste, A. J. Marques Cardoso, A. Belahcen, and T. Vaimann. The oil spray cooling system of automotive traction motors: The state of the art. *IEEE Transactions on Transportation Electrification*, 9(1):428–451, 2023.

[36] Z. Xu, Y Xu, W Liu, and Y Wang. Stator winding hotspot temperature rise characteristic study of an axially forced air-cooled motor with air deflector. *Applied Thermal Engineering*, 224:120108, 2023.

[37] J. Pecotich, D. Klink, G. Heins, and B. Bahrani. Additively manufactured electric machine conductors with integrated end turn heat exchangers. In *2022 International Conference on Electrical Machines (ICEM)*, pages 1498–1504, 2022.

[38] N. Simpson, C. Tighe, and P. Mellor. Design of high performance shaped profile windings for additive manufacture. In *2019 IEEE Energy Conversion Congress and Exposition (ECCE)*, pages 761–768, 2019.

[39] S. Vahid, T. Chowdhury, S. Koushan, and A. EL-Refaie. Electrical characteristics of additively manufactured hollow conductor coils with integrated heat pipes for electric aircraft applications. In *2022 IEEE Energy Conversion Congress and Exposition (ECCE)*, pages 1–8, 2022.

[40] Mingda Liu, Yingjie Li, Hao Ding, and B. Sarlioglu. Thermal management and cooling of windings in electrical machines for electric vehicle and traction application. In *2017 IEEE Transportation Electrification Conference and Expo (ITEC)*, pages 668–673, June 2017.

[41] F. Chai, Y. Bi, and L. Chen. Thermal investigation and cooling enhancement of axial flux permanent magnet motors for vehicle applications. In *2019 22nd International Conference on Electrical Machines and Systems (ICEMS)*, pages 1–5, 2019.

[42] A. La Rocca, S. La Rocca, T. Zou, C. Liu, M. Moslemin, C. Gerada, and A. Cairns. Performance assessment of standard cooling strategies for hairpin windings. In *2022 International Conference on Electrical Machines (ICEM)*, pages 1163–1169, 2022.

[43] Z. Li, B. Xiong, and K. Huang. Research on stator core axial pipe to improve cooling performance of permanent magnet synchronous motor. In *2022 25th International Conference on Electrical Machines and Systems (ICEMS)*, pages 1–5, 2022.

[44] W. Yu, D. Fangmian, W. Jiankangy, C. Ziran, and R. Lin. Effect of evaporative cooling of stator core on electromagnetic field of large horizontal generator. In *2022*

25th International Conference on Electrical Machines and Systems (ICEMS), pages 1–4, 2022.

[45] R. Zhou, G. J. Li, Z. Q. Zhu, M. P. Foster, D. A. Stone, C. J. Jia, and P. McKeever. Novel liquid cooling technology for modular consequent-pole PM machines. In *2021 IEEE International Electric Machines & Drives Conference (IEMDC)*, pages 1–7, 2021.

[46] W. Liu, W. Li, S. Luo, X. Huang, D. Li, Z. Li, and G. Xu. Influence of a novel stator teeth internal ventilation structure on air-cooled turbo-generator parameters and stator temperature. *IEEE Access*, 8:122422–122433, 2020.

[47] Gokhan Cakal, Ahmed Hembel, and Bulent Sarlioglu. Additively manufactured heat exchanger for improved cooling of electric machines. In *2022 IEEE Energy Conversion Congress and Exposition (ECCE)*, pages 1–6, 2022.

[48] US Patent US7489057B2 Tesla, 2009.

[49] N. A. Rahman, E. Bostanci, and B. Fahimi. Thermal analysis of switched reluctance motor with direct in-winding cooling system. In *2016 IEEE Conference on Electromagnetic Field Computation (CEFC)*, pages 1–1, Nov 2016.

[50] M. Schiefer and M. Doppelbauer. Indirect slot cooling for high-power-density machines with concentrated winding. In *2015 IEEE International Electric Machines Drives Conference (IEMDC)*, pages 1820–1825, May 2015.

[51] M. Al-Ani, A. Al-Timimy, P. Giangrande, M. Degano, P. Lindh, H. Zhang, M. Galea, and C. Gerada. Influence of manufacturing and drive effects in high-speed, high-power-density PM machine for flooded pump application. In *2021 IEEE Workshop on Electrical Machines Design, Control and Diagnosis (WEMDCD)*, pages 88–94, 2021.

[52] G. Venturini, G. Volpe, and M. Popescu. Slot water jacket cooling system for traction electrical machines with hairpin windings: Analysis and comparison. In *2021 IEEE International Electric Machines & Drives Conference (IEMDC)*, pages 1–6, 2021.

[53] F. R. Ismagilov, V. Y. Vavilov, V. I. Bekuzin, V. V. Ayguzina, and D. Y. Permin. Analysis of stator cooling methods of fault-tolerant electric machines. In *2020 Fifteenth International Conference on Ecological Vehicles and Renewable Energies (EVER)*, pages 1–6, 2020.

[54] S. Yamamoto, A. Nishikawa, T. Harada, and W. Kosaka. Development of direct cooling stator structure using high thermal conductive epoxy molding compounds. In *2022 International Power Electronics Conference (IPEC-Himeji 2022- ECCE Asia)*, pages 407–414, 2022.

[55] Z. Yao, R. K. Mandel, and F. P. McCluskey. In-slot cooling for high power density electric motor with encapsulation channels. In *2022 21st IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (iTherm)*, pages 1–6, 2022.

[56] C. Di, I. Petrov, and J. Pyrhönen. Estimation of continuous power of a permanent magnet synchronous machine equipped with direct-liquid-cooling winding for propulsion applications by efa. In *2020 International Conference on Electrical Machines (ICEM)*, volume 1, pages 854–859, 2020.

[57] P. Lindh, I. Petrov, A. Jaatinen-Värri, A. Grönman, M. Martinez-Iturralde, M. Satrústegui, and J. Pyrhönen. Direct liquid cooling method verified with an axial-flux permanent-magnet traction machine prototype. *IEEE Transactions on Industrial Electronics*, 64(8):6086–6095, Aug 2017.

[58] C. Wohlers, P. Juris, S. Kabelac, and B. Ponick. Design and direct liquid cooling of tooth-coil windings. *Electrical Engineering*, 100:1–10, 12 2018.

[59] C. Wohlers and B. Ponick. First estimations of stator dimensions for permanent magnet synchronous machines with tooth-coil windings and direct liquid cooling. In *2020 International Conference on Electrical Machines (ICEM)*, volume 1, pages 462–468, 2020.

[60] A. Reinap, M. Andersson, F. J. Márquez-Fernández, P. Abrahamsson, and M. Alaküla. Performance estimation of a traction machine with direct cooled hairpin winding. In *2019 IEEE Transportation Electrification Conference and Expo (ITEC)*, pages 1–6, 2019.

[61] A. Reinap, M. Gabassi, M. Alaküla, and M. Andersson. Assessment of cooling integration with direct cooled windings. In *2018 IEEE International Conference on Electrical Systems for Aircraft, Railway, Ship Propulsion and Road Vehicles International Transportation Electrification Conference (ESARS-ITEC)*, pages 1–6, Nov 2018.

[62] S. Estenlund. Hollow direct air cooled windings: Rotor design process. In *2020 International Symposium on Power Electronics, Electrical Drives, Automation and Motion (SPEEDAM)*, pages 489–496, 2020.

[63] S. Estenlund, A. Tokat, J. Engqvist, and M. Alaküla. Dovetail design for direct cooled rotor: Design and manufacturing. In *2022 International Conference on Electrical Machines (ICEM)*, pages 2121–2127, 2022.

[64] S. Ayat, C. Serghine, T. Klonowski, S. Yon, A. Mutabazi, and S. McDaniel. The use of phase change material for the cooling of electric machine windings formed with hollow conductors. In *2019 IEEE International Electric Machines & Drives Conference (IEMDC)*, pages 1195–1201, 2019.

[65] D. Bobba, Z. Yao, J. Swanke, R. Mandel, P. McCluskey, T. Jahns, and B. Sarlioglu. Multi-physics based analysis and design of stator coil in high power density PMSM for aircraft propulsion applications. In *2021 AIAA/IEEE Electric Aircraft Technologies Symposium (EATS)*, pages 1–9, 2021.

[66] S. Haller, J. Persson, P. Cheng, and K. Bertilsson. Multi-phase winding with in-conductor direct cooling capability for a 48v traction drive design. In *2020 International Conference on Electrical Machines (ICEM)*, volume 1, pages 2118–2124, 2020.

[67] J. Zhang, X. Zhu, Z. Zhang, and Y. Xia. AC loss calculation and analysis of hollow conductor for doubly salient brushless DC generator. *IEEE Transactions on Magnetics*, 58(8):1–5, 2022.

[68] X. Chen, J. Wang, A. Griffo, and A. Spagnolo. Thermal modeling of hollow conductors for direct cooling of electrical machines. *IEEE Transactions on Industrial Electronics*, 67(2):895–905, Feb 2020.

[69] Dynamic E Flow GMBH, 2024 (accessed on 2024-04-29). https://www.dynamiceflow.com/.

[70] C. Liu, Z. Xu, D. Gerada, J. Li, C. Gerada, Y. C. Chong, M. Popescu, J. Goss, D. Staton, and H. Zhang. Experimental investigation on oil spray cooling with hairpin windings. *IEEE Transactions on Industrial Electronics*, 67(9):7343–7353, 2020.

[71] Material properties from engineering toolbox. https://www.engineeringtoolbox.com/. (accessed on 2023-05-05).

[72] Viscosity of engine oil from Anton Paar Wiki. https://wiki.anton-paar.com/en/engine-oil/. (accessed on 2023-05-05).

[73] Z. Huang, A. Reinap, and M. Alaküla. Dielectric properties modeling and measurement of single tooth coil insulation system under accelerated degradation test. In *2016 XXII International Conference on Electrical Machines (ICEM)*, pages 2698–2703, 2016.

[74] X. Wang, S. Zhou, L. Wu, M. Zhao, and C. Hu. Iron loss and thermal analysis of high speed PM motor using soft magnetic composite material. In *2019 22nd International Conference on Electrical Machines and Systems (ICEMS)*, pages 1–4, 2019.

[75] Y. Ling Lim, W. L. Soong, N. Ertugrul, and S. Kahourzade. Embedded stator end-windings in soft magnetic composite and laminated surface PM machines. In *2018 IEEE Energy Conversion Congress and Exposition (ECCE)*, pages 5387–5394, 2018.

[76] Höganäs. Somaloy 3P, June 2018 (accessed on 2023-09-29). https://www.hoganas.com/en/powder-technologies/products/somaloy/somaloy-3p/.

[77] Tata steel. Electrical steel - M235-35A, 2008 (accessed on 2023-09-29). https://www. tatasteeleurope.com/engineering/products/electrical-steel.

[78] W. W. Salsig and H. W. Vogel. Extended flattop operation of bevatron results from studies of generator dovetail failure. *IEEE Transactions on Nuclear Science*, 16(3):699–703, 1969.

[79] Y. Hidaka. High-torque density magnet-assisted wound field motor using a field-unit-type rotor. In *2022 International Conference on Electrical Machines (ICEM)*, pages 185–191, 2022.

[80] M. K. Kheto, N. C. Mahendra Babu, and J. Madan. Fretting fatigue analysis in dovetail joint of compressor through numerical simulation. *SASTech Journal*, 8(1):71–76, April 2009.

[81] L. Witek. Failure analysis of turbine disc of an aero engine. *Engineering Failure Analysis*, 13(1):9–17, 2006.

[82] M.M.I. Hammouda, R.A. Pasha, and A.S. Fayed. Modelling of cracking sites/development in axial dovetail joints of aero-engine compressor discs. *International Journal of Fatigue*, 29(1):30–48, 2007.

[83] M. Allara, S. Zucca, and M. Gola. Effect of crowning of dovetail joints on turbine blade root damping. *KEY ENGINEERING MATERIALS*, 347:317–322, 09 2007.

[84] Femi Robert, Amalin Prince A, and Jac Fredo A R. Influence of wire electrical discharge machine cutting parameters on the magnetization characteristics of electrical steel laminations. *Materials Today: Proceedings*, 52:746–750, 2022. International Conference on Smart and Sustainable Developments in Materials, Manufacturing and Energy Engineering.

Division of Industrial Electrical
Engineering and Automation
Department of Biomedical Engineering

Faculty of Engineering, Lund University