POLITECNICO DI TORINO

Master’s degree in Energy and Nuclear Engineering



Master’s Degree Thesis

Investigation of direct rotor cooling in an Electrically Excited Synchronous Machine (EESM) using single-phase and two-phase dielectric fluids

Supervisors:

Prof. Andrea Tonoli

Prof. Vittorio Ravello

Prof. Gianmario Pellegrino

Candidate:

Andrea Concini

A.Y. 2025/2026

Abstract

The increasing adoption of high power density electric machines requires increasingly effective cooling solutions to ensure high performance, reliability, and long-term durability. In this work, the direct rotor cooling of an Electrically Excited Synchronous Machine (EESM) is investigated using dielectric fluids flowing through cooling channels integrated within the rotor.

The three-dimensional geometry of the cooling system was developed using SolidWorks, including the fluid ducts, rotor windings, laminated core, and insulating components. Subsequently, a Computational Thermal Fluid Dynamics (CtFD) analysis was performed using Ansys Fluent to investigate heat transfer phenomena, temperature distribution, pressure losses, and refrigerant phase-change behavior.

In the first stage of the study, the cooling performance of the system was analyzed using the single-phase dielectric oil AmpCool AC-110, with particular attention given to the influence of the coolant inlet temperature. Four operating conditions were considered, corresponding to inlet temperatures of 20 °C, 40 °C, 60 °C and 80 °C, in order to quantify the effect of the supply conditions on the cooling of the rotor windings and magnetic laminations.

Subsequently, a comparative study was carried out using four two-phase refrigerants characterized by different thermophysical properties: Novec 7000, Opteon SF33, R134a, and FC-72. For each fluid, evaporation phenomena and latent heat transport were investigated through multiphase simulations, evaluating temperature distribution, vapor fraction, cooling effectiveness, and pressure drop along the rotor cooling channels.

The results demonstrate that the use of two-phase refrigerants significantly enhances the heat removal capability compared to the single-phase solution, owing to the exploitation of phase change and the latent heat of vaporization. Furthermore, the study made it possible to identify the most effective operating configurations in terms of winding temperature reduction and to determine the optimal compromise between thermal performance and hydraulic requirements of the cooling system.

Contents

List of Figures……6

List of Tables……12

1. Introduction.....14

- **1.1** Background and Motivation.....14
- **1.2** Objectives of the Thesis.....16

2. Electrically Excited Synchronous Motors (EESMs)……18

- **2.1** Operating Principle.....18
  - 2.1.1 The Additional Degree of Freedom: Excitation Current……...19
  - 2.1.2 Systematic Comparison with Competing Technologies……20
  - 2.1.3 Why EESM Technology Is Gaining Ground……23
- **2.2** Rotor Structure……24
  - 2.2.1 The Excitation Winding……25
  - 2.2.2 Rotor Excitation Supply Systems……..26
- **2.3** Advantages ……..27
  - 2.3.1 Critical Issues ……...27
- **2.4** Heat Sources in the Rotor: Joule Losses in the Windings, Iron Losses, and Mechanical Losses……28
  - 2.4.1 Joule Losses in the Excitation Winding……29
  - 2.4.2 Rotor Iron Losses……..29
  - 2.4.3 Mechanical Losses……30
  - 2.4.4 Industrial Context: Adoption of EESMs……30
- **2.5** Why Rotor Cooling Represents the Main Thermal Challenge of this Technology……..31
  - 2.5.1 Heat Dissipation Path……32
  - 2.5.2 Centrifugal Force Constraints……32
  - 2.5.3 Thermal Feedback of Joule Losses……32
  - 2.5.4 Direct Relationship Between Thermal and Electromechanical Performance...33

3. Cooling Techniques for Electric Motors……34

- **3.1** Indirect Cooling (Housing and Water Jacket)……..34
  - 3.1.1 Advantages and Limitations……..35
- **3.2** Direct Cooling: Hairpin/Oil Spray and Rotor Channel Cooling……..36

- 3.2.1 Direct Stator Cooling: Hairpin Cooling and Oil Spray……36
- 3.2.2 Direct Rotor Cooling: Internal Rotor Channels……37
- 3.2.3 Electrical Compatibility of the Coolant……38

- **3.3** Single-Phase Cooling……...39
  - 3.3.1 Two-Phase Cooling and Latent Heat……40
  - 3.3.2 Critical Issues in Two-Phase Cooling: Critical Heat Flux and Instabilities…..41
- **3.4** State of the Art: Review of Scientific Literature on Two-Phase Cooling of Rotors and Stators……42
  - 3.4.1 Two-Phase Stator Cooling……42
  - 3.4.2 Two-Phase Rotor Cooling……43
  - 3.4.3 Positioning of the Present Work with Respect to the State of the Art……...43
- **3.5** Working Fluid Selection and Selection Criteria……...44
  - 3.5.1 Candidate Fluids Analyzed in the Present Work……...45
  - 3.5.2 Methodology for the Verification of Thermophysical Properties……46

4. Numerical Modelling and CFD Setup……47

- **4.1** Geometry and Computational Domain.....47
  - 4.1.1 Domain Reduction Criteria……47
  - 4.1.2 Description of the Modelled Geometry…..……48
  - 4.1.3 The Two Cooling Channels……...50
  - 4.1.4 Component Materials……50
  - 4.1.5 Heat Source in the Copper Windings……52
  - 4.1.6 Geometry Preparation in Ansys SpaceClaim……53
- **4.2** Mesh and Computational Grid Generation Strategy……54
  - 4.2.1 Grid Configurations Analyzed……..……..55
  - 4.2.2 Methodology and Results of the Grid-Independence Assessment……57
- **4.3** Physical Model for the Single-Phase Case……..……..58
  - 4.3.1 Physical Model for the Two-Phase Case……..59
- **4.4** Boundary Conditions for the Single-Phase Case……61
  - 4.4.1 Boundary Conditions for the Multiphase Case……62
- **4.5** Under-Relaxation Factors and Solver Setup……63
  - 4.5.1 Solver Setup for the Single-Phase Case……66
  - 4.5.2 Solver Setup for the Two-Phase Case……67

5. Case Study: Single-Phase Cooling……..69

- **5.1** Coolant: AmpCool AC-110 Dielectric Oil……...69
  - 5.1.1 Operating Conditions of the Case Study……...71
- **5.2** Simulation Results……71
  - 5.2.1 Simulation with $T_{in} = 20°C$……..71
  - 5.2.2 Simulation with $T_{in} = 40°C$……..72
  - 5.2.3 Simulation with $T_{in} = 60°C$……..73

- 5.2.4 Simulation with $T_{in} = 80°C$……..74
- **5.3** Presentation of the Results……...76
  - 5.3.1 Analysis of Temperature Fields and Hot Spots……77
  - 5.3.2 Pressure Drop Analysis……80
- **5.4** Critical Analysis of the Limitations of Single-Phase Cooling……81

6. Case Study: Two-Phase Cooling — Comparison of Four Fluids……82

- **6.1** Methodology for the Retrieval of Thermophysical Properties……82
- **6.2** Thermophysical Properties of the Liquid and Vapor Phases……83
- **6.3** Operating Conditions and Latent Heat of Vaporization……...84
- **6.4** Rationale for the Selection of the Four Fluids……..85
- **6.5** Comparison Methodology.....86
  - 6.5.1 Parametric Study of the Mass Flow Rate……..87
  - 6.5.2 Criteria for the Selection of the Optimal Operating Point……88
  - 6.5.3 Performance Metrics Adopted for the Final Comparison……..89
- **6.6** Parametric Study of the Mass Flow Rate: : Identification of the Optimal Mass Flow Rate……...……..89
  - 6.6.1 First-Estimate Mass Flow Rate……90
  - 6.6.2 Results and Limitations of the First-Estimate Mass Flow Rate……91
  - 6.6.3 Progressive Increase in Mass Flow Rate: Results for the Four Fluids……...……...92
  - 6.6.4 Identification of the Optimal Operating Point……102
- **6.7** Comparison of the Four fluids at the Optimal Mass Flow Rate…….....104
  - 6.7.1 Novec 7000……..105
  - 6.7.2 FC-72.....111
  - 6.7.3 Opteon SF-33……...117
  - 6.7.4 R134a……...122
- **6.8** Comparative Analysis of the Four Fluids and Comparison with the Single-Phase Case…….....127

7. Discussion of the Results.....130

- **7.1** Overall Performance Comparison.....130
- **7.2** Design Implications.....133
- **7.3** Limitations of the Numerical Model and Potential Sources of Uncertainty……..134

8. Conclusions and Future Developments.....137

- **8.1** Summary of the Main Results……137
- **8.2** Possible Extensions.....142

Bibliography……147

Technical data sheets…….....149

List of Figures

Figure 1 – EESM Motor Operating Principle

Figure 2 – PMSM Motor

Figure 3 – IM Motor

Figure 4 – SynRM Motor

Figure 5 – EESM Rotor

Figure 6 – Losses in an Electric Motor

Figure 7 – Water Jacket

Figure 8 – Oil Spray Technique

Figure 9 – Direct Oil Cooling

Figure 10 – 3D SolidWorks Model of the Rotor, Front View

Figure 11 – 3D SolidWorks Model of the Rotor, Isometric View

Figure 12 – Developed Rotor Computational Domain, Isometric View

Figure 13 – Analyzed Rotor Computational Domain, Front View

Figure 14 – Plastic Partition, Front View

Figure 15 – Plastic Partition, Isometric View

Figure 16 – Surface Mesh Generation, Isometric View

Figure 17 – Volume Mesh Generation Using Poly-Hexcore, Isometric View

Figure 18 – Detail of the Boundary-Layer Mesh in the Fluid Domains, Front View

Figure 19 – Pressure Drop Report, Inlet Temperature 20 °C

Figure 20 – Average Cooling-Channel Temperature Report, Inlet Temperature 20 °C

Figure 21 – Average Winding Temperature Report, Inlet Temperature 20 °C

Figure 22 – Maximum Cooling-Channel Temperature Report, Inlet Temperature 20 °C

Figure 23 – Maximum Winding Temperature Report, Inlet Temperature 20 °C

Figure 24 – Cooling-Channel Outlet Temperature Report, Inlet Temperature 20 °C

Figure 25 – Pressure Drop Report, Inlet Temperature 40 °C

Figure 26 – Average Cooling-Channel Temperature Report, Inlet Temperature 40 °C

Figure 27 – Average Winding Temperature Report, Inlet Temperature 40 °C

Figure 28 – Maximum Cooling-Channel Temperature Report, Inlet Temperature 40 °C

Figure 29 – Maximum Winding Temperature Report, Inlet Temperature 40 °C

Figure 30 – Cooling-Channel Outlet Temperature Report, Inlet Temperature 40 °C

Figure 31 – Pressure Drop Report, Inlet Temperature 60 °C

Figure 32 – Average Cooling-Channel Temperature Report, Inlet Temperature 60 °C

Figure 33 – Average Winding Temperature Report, Inlet Temperature 60 °C

Figure 34 – Maximum Cooling-Channel Temperature Report, Inlet Temperature 60 °C

Figure 35 – Maximum Winding Temperature Report, Inlet Temperature 60 °C

Figure 36 – Cooling-Channel Outlet Temperature Report, Inlet Temperature 60 °C

Figure 37 – Pressure Drop Report, Inlet Temperature 80 °C

Figure 38 – Average Cooling-Channel Temperature Report, Inlet Temperature 80 °C

Figure 39 – Average Winding Temperature Report, Inlet Temperature 80 °C

Figure 40 – Maximum Cooling-Channel Temperature Report, Inlet Temperature 80 °C

Figure 41 – Maximum Winding Temperature Report, Inlet Temperature 80 °C

Figure 42 – Cooling-Channel Outlet Temperature Report, Inlet Temperature 80 °C

Figure 43 – Average Cooling-Channel Temperature as a Function of Refrigerant Inlet Temperature

Figure 44 – Maximum Cooling-Channel Temperature as a Function of Refrigerant Inlet Temperature

Figure 45 – Average Winding Temperature as a Function of Refrigerant Inlet Temperature

Figure 46 – Maximum Winding Temperature as a Function of Refrigerant Inlet Temperature

Figure 47 – Average Plastic-Partition Temperature as a Function of Refrigerant Inlet Temperature

Figure 48 – Maximum Plastic-Partition Temperature as a Function of Refrigerant Inlet Temperature

Figure 49 – Average Rotor Ferromagnetic-Structure Temperature as a Function of Refrigerant Inlet Temperature

Figure 50 – Maximum Rotor Ferromagnetic-Structure Temperature as a Function of Refrigerant Inlet Temperature

Figure 51 – Refrigerant Outlet Temperature as a Function of Refrigerant Inlet Temperature

Figure 52 – Pressure Drop as a Function of Refrigerant Inlet Temperature

Figure 53 - Report temperature outlet all four fluids

Figure 54 - Report temperature average duct all four fluids

Figure 55 – Report temperature max duct all four fluids

Figure 56 - Report temperature average winding all four fluids

Figure 57 - Report temperature max winding all four fluids

Figure 58 - Report temperature average lamination all four fluids

Figure 59 - Report temperature max lamination all four fluids

Figure 60 - Report temperature average plastic all four fluids

Figure 61 - Report temperature max plastic all four fluids

Figure 62 - Report pressure drop all four fluids

Figure 63 - Report liquid fraction all four fluids

Figure 64 - Report vapor fraction all four fluids

Figure 65 - Report vapor fraction outlet all four fluids

Figure 66 – Average Cooling-Channel Temperature Report, Novec 7000

Figure 67 – Maximum Cooling-Channel Temperature Report, Novec 7000

Figure 68 – Average Winding Temperature Report, Novec 7000

Figure 69 – Maximum Winding Temperature Report, Novec 7000

Figure 70 – Novec 7000 Outlet Temperature Report

Figure 71 – Pressure Drop Report, Novec 7000

Figure 72 – Liquid Fraction Report, Novec 7000

Figure 73 – Vapor Fraction Report, Novec 7000

Figure 74 – Outlet Vapor Fraction Report, Novec 7000

Figure 75 – Electric Rotor Computational-Domain Contour, Outlet Side, Isometric View

Figure 76 – Electric Rotor Computational-Domain Contour, Inlet Side, Isometric View

Figure 77 – Electric Rotor Computational-Domain Contour, Outlet Side, Front View

Figure 78 – Electric Rotor Computational-Domain Contour, Inlet Side, Front View

Figure 79 – Copper-Winding Contour, Isometric View

Figure 80 – Copper-Winding Contour, Top View

Figure 81 – Plastic-Partition Contour, Outlet Side, Isometric View

Figure 82 – Plastic-Partition Contour, Inlet Side, Isometric View

Figure 83 – Plastic-Partition Contour, Side View

Figure 84 – Rotor Ferromagnetic-Structure Contour, Outlet Side

Figure 85 – Rotor Ferromagnetic-Structure Contour, Inlet Side

Figure 86 – Vapor-Phase Contour in the Cooling Channels, Inlet Side

Figure 87 – Vapor-Phase Contour in the Cooling Channels, Outlet Side

Figure 88 – Vapor-Phase Contour in the Cooling Channels, Side View

Figure 89 – Pressure-Drop Contour in the Cooling Channels, Inlet Side

Figure 90 – Average Cooling-Channel Temperature Report, FC-72

Figure 91 – Maximum Cooling-Channel Temperature Report, FC-72

Figure 92 – Average Winding Temperature Report, FC-72

Figure 93 – Maximum Winding Temperature Report, FC-72

Figure 94 – FC-72 Outlet Temperature Report

Figure 95 – Pressure Drop Report, FC-72

Figure 96 – Liquid Fraction Report, FC-72

Figure 97 – Vapor Fraction Report, FC-72

Figure 98 – Outlet Vapor Fraction Report, FC-72

Figure 99 – Electric Rotor Computational-Domain Contour, Inlet Side, Isometric View

Figure 100 – Electric Rotor Computational-Domain Contour, Outlet Side, Isometric View

Figure 101 – Electric Rotor Computational-Domain Contour, Inlet Side, Front View

Figure 102 – Electric Rotor Computational-Domain Contour, Outlet Side, Front View

Figure 103 – Copper-Winding Contour, Isometric View

Figure 104 – Copper-Winding Contour, Top View

Figure 105 – Plastic-Partition Contour, Outlet Side, Isometric View

Figure 106 – Plastic-Partition Contour, Inlet Side, Isometric View

Figure 107 – Plastic-Partition Contour, Side View

Figure 108 – Rotor Ferromagnetic-Structure Contour, Outlet Side

Figure 109 – Rotor Ferromagnetic-Structure Contour, Inlet Side

Figure 110 – Vapor-Phase Contour in the Cooling Channels, Inlet Side

Figure 111 – Vapor-Phase Contour in the Cooling Channels, Outlet Side

Figure 112 – Vapor-Phase Contour in the Cooling Channels, Side View

Figure 113 – Pressure-Drop Contour in the Cooling Channels, Inlet Side

Figure 114 – Average Cooling-Channel Temperature Report, Opteon SF-33

Figure 115 – Maximum Cooling-Channel Temperature Report, Opteon SF-33

Figure 116 – Average Winding Temperature Report, Opteon SF-33

Figure 117 – Maximum Winding Temperature Report, Opteon SF-33

Figure 118 – Opteon SF-33 Outlet Temperature Report

Figure 119 – Pressure Drop Report, Opteon SF-33

Figure 120 – Liquid Fraction Report, Opteon SF-33

Figure 121 – Vapor Fraction Report, Opteon SF-33

Figure 122 – Outlet Vapor Fraction Report, Opteon SF-33

Figure 123 – Electric Rotor Computational-Domain Contour, Inlet Side, Isometric View

Figure 124 – Electric Rotor Computational-Domain Contour, Outlet Side, Isometric View

Figure 125 – Electric Rotor Computational-Domain Contour, Inlet Side, Front View

Figure 126 – Electric Rotor Computational-Domain Contour, Outlet Side, Front View

Figure 127 – Copper-Winding Contour, Isometric View

Figure 128 – Copper-Winding Contour, Top View

Figure 129 – Plastic-Partition Contour, Outlet Side, Isometric View

Figure 130 – Plastic-Partition Contour, Inlet Side, Isometric View

Figure 131 – Plastic-Partition Contour, Side View

Figure 132 – Rotor Ferromagnetic-Structure Contour, Outlet Side

Figure 133 – Rotor Ferromagnetic-Structure Contour, Inlet Side

Figure 134 – Vapor-Phase Contour in the Cooling Channels, Inlet Side

Figure 135 – Vapor-Phase Contour in the Cooling Channels, Outlet Side

Figure 136 – Vapor-Phase Contour in the Cooling Channels, Side View

Figure 137 – Pressure-Drop Contour in the Cooling Channels, Inlet Side

Figure 138 – Average Cooling-Channel Temperature Report, R134a

Figure 139 – Maximum Cooling-Channel Temperature Report, R134a

Figure 140 – Average Winding Temperature Report, R134a

Figure 141 – Maximum Winding Temperature Report, R134a

Figure 142 – R134a Outlet Temperature Report

Figure 143 – Pressure Drop Report, R134a

Figure 144 – Liquid Fraction Report, R134a

Figure 145 – Vapor Fraction Report, R134a

Figure 146 – Outlet Vapor Fraction Report, R134a

Figure 147 – Electric Rotor Computational-Domain Contour, Inlet Side, Isometric View

Figure 148 – Electric Rotor Computational-Domain Contour, Outlet Side, Isometric View

Figure 149 – Electric Rotor Computational-Domain Contour, Inlet Side, Front View

Figure 150 – Electric Rotor Computational-Domain Contour, Outlet Side, Front View

Figure 151 – Copper-Winding Contour, Isometric View

Figure 152 – Copper-Winding Contour, Top View

Figure 153 – Plastic-Partition Contour, Outlet Side, Isometric View

Figure 154 – Plastic-Partition Contour, Inlet Side, Isometric View

Figure 155 – Plastic-Partition Contour, Side View

Figure 156 – Vapor-Phase Contour in the Cooling Channels, Outlet Side

Figure 157 – Vapor-Phase Contour in the Cooling Channels, Inlet Side

Figure 158 – Vapor-Phase Contour in the Cooling Channels, Side View

Figure 159 – Pressure-Drop Contour in the Cooling Channels, Inlet Side

List of Tables

Table 1 – Characteristics of Different Electric Motor Types

Table 2 – Thermophysical Properties of the Materials Used in the Model

Table 3 – Geometric and Thermal Data of the Winding (for Each Modelled Half-Winding)

Table 4 – Mesh Generation Parameters for the Three Mesh Configurations

Table 5 – Number of Nodes, Cells, and Faces for the Three Mesh Configurations

Table 6 – Results of the Grid-Independence Assessment (Oil Inlet Temperature of 40 °C)

Table 7 – Boundary Conditions for the Single-Phase Case

Table 8 – Boundary Conditions for the Multiphase Case

Table 9 – Under-Relaxation Factors for the Single-Phase Case

Table 10 – Under-Relaxation Factors for the Multiphase Case

Table 11 – Main Thermophysical and Dielectric Properties of AmpCool AC-110 Oil

Table 12 – Results of the Single-Phase Simulations at Different Inlet Temperatures

Table 13 – Overall Energy Balance of the Computational Domain at Different Inlet Temperatures

Table 14 – Thermophysical Properties of the Four Two-Phase Fluids (Liquid and Vapor Phases under Saturation Conditions)

Table 15 – Saturation Conditions and Latent Heat of Vaporization of the Four Fluids (T_{sat} = 60 °C)

Table 16 – Simulation Results at a Mass Flow Rate of 0.000568 kg/s per Channel

Table 17 – Temperature outlet all four fluids

Table 18 - T average duct all four fluids

Table 19 - Temperature max duct all four fluids

Table 20 - Temperature average winding all four fluids

Table 21 - Temperature max winding all four fluids

Table 22 - Temperature average lamination all four fluids

Table 23 - Temperature max lamination all four fluids

Table 24 - Temperature average plastic

Table 25 - Temperature max plastic all four fluids

Table 26 – Pressure drop all four fluids

Table 27 – Liquid fraction all four fluids

Table 28 – Vapor fraction all four fluids

Table 29 – Vapor fraction outlet all four fluids

Table 30 - Average values and thermal benefit all four fluids

Table 31 – Novec 7000 Two-Phase Simulation Results (0.0035 kg/s per Channel)

Table 32 – Energy Balance of the Novec 7000 Two-Phase Simulation (0.0035 kg/s per Channel)

Table 33 – FC-72 Two-Phase Simulation Results (0.0035 kg/s per Channel)

Table 34 – Energy Balance of the FC-72 Two-Phase Simulation (0.0035 kg/s per Channel)

Table 35 – Opteon SF-33 Two-Phase Simulation Results (0.0035 kg/s per Channel)

Table 36 – Energy Balance of the Opteon SF-33 Two-Phase Simulation (0.0035 kg/s per Channel)

Table 37 – R134a Two-Phase Simulation Results (0.0035 kg/s per Channel)

Table 38 – Energy Balance of the R134a Two-Phase Simulation (0.0035 kg/s per Channel)

Table 39 – Comparison of the Properties and Two-Phase Simulation Results of the Four Fluids

Table 40 – Single-Phase Simulation Results (Refrigerant Inlet Temperature of 60 °C)

Table 41 – Comparison between Two-Phase Refrigerants and Single-Phase Oil

Table 42 – Summary of the Key Findings of the Thesis Study

Table 43 – Summary of the Main Quantities from the Two-Phase and Single-Phase Simulations

Chapter 1

Introduction

1.1 Background and Motivation

The electrification process of propulsion systems, which has been underway for more than a decade in the automotive sector and is progressively extending to the aerospace, railway, and industrial sectors, has imposed increasingly stringent performance requirements on electrical machine designers in terms of power density, torque density, and energy efficiency. The drive towards reducing the mass and volume of electric drives, dictated by packaging constraints, range requirements, and the overall system cost, has resulted in a systematic increase in electromagnetic and thermal loads for a given machine size, pushing electrical machines towards operating conditions increasingly close to the thermal limits of the insulating and magnetic materials employed.

Within this context, Electrically Excited Synchronous Motors (EESMs) have emerged as a viable alternative to the more widely adopted Permanent Magnet Synchronous Motors (PMSMs), owing to the possibility of actively controlling the rotor excitation flux, the absence of rare-earth permanent magnets—with consequent advantages in terms of cost, raw-material availability, and sustainability of the manufacturing supply chain—and greater operating flexibility in high-speed field-weakening regions. This functional advantage is, however, counterbalanced by a critical issue inherent to the EESM topology: the presence of a DC-fed excitation winding on the rotor introduces a concentrated source of Joule losses in a region of the machine that has historically been more difficult to cool than the stator, due to rotation, limited available space, and the need to ensure electrical insulation and mechanical integrity under high centrifugal forces.

If not adequately dissipated, the losses generated in the rotor lead to an increase in the temperature of the excitation windings, which may result in the thermal rating of the insulation being exceeded, accelerated degradation of polymeric materials and, in the most critical cases, performance degradation due to thermal derating imposed by the control system or, under extreme conditions, failure caused by insulation breakdown. Effective rotor cooling therefore represents an enabling factor for increasing the power density of EESMs, since it allows operation at higher current densities in the excitation windings while maintaining the same thermal margin. This directly translates into a reduction in machine mass and volume or, for a given machine size, an increase in the torque and power that can be delivered continuously.

Cooling solutions traditionally adopted for the rotor, based on single-phase forced convection using air or liquid coolants (typically water-glycol mixtures) circulating through channels integrated into

the rotor core or shaft, exhibit inherent limitations associated with the sensible heat capacity of the fluid. Heat removal occurs solely through an increase in the fluid temperature along the channel, resulting in the development of significant axial thermal gradients and the occurrence of hot spots in the inlet regions or in areas corresponding to the most concentrated heat sources. To overcome these limitations, recent scientific literature has devoted increasing attention to two-phase cooling techniques, in which the refrigerant undergoes a liquid-to-vapor phase transition within the cooling channel, exploiting the latent heat of vaporization—typically one or two orders of magnitude greater than the sensible heat that can be exchanged over the same temperature range—to achieve significantly higher convective heat transfer coefficients and improved surface temperature uniformity for the same temperature rise above ambient conditions.

The use of engineered dielectric fluids, such as fluoroketones, hydrofluoroethers (HFEs) belonging to the Novec family, fluorinated electronic liquids, together with other dielectric or low-environmental-impact refrigerants, also enables the possibility of direct cooling of the windings without the interposition of additional electrical insulation barriers, thereby further reducing the overall thermal resistance of the system. This thesis is positioned within this research framework and aims to investigate, through Computational Fluid Dynamics (CFD) simulations, the thermal performance of a rotor cooling channel in an EESM. The analysis first considers a reference single-phase cooling configuration and subsequently a set of two-phase cooling configurations employing four different working fluids, with the aim of characterizing their thermofluid-dynamic behavior and identifying their optimal operating conditions.

Despite the growing scientific interest in two-phase cooling techniques outlined above, a systematic review of the state of the art (presented in detail in Section 3.4) highlights that several aspects specifically relevant to the direct two-phase cooling of EESM rotors remain only partially addressed by the current literature. First, the majority of the available numerical and experimental studies on rotor cooling are limited to a single working fluid or a single cooling technology (e.g., oil-spray impingement, single-phase channel cooling, or one specific two-phase refrigerant), which prevents a direct and methodologically consistent comparison of the relative thermal and fluid-dynamic performance achievable with different candidate fluids. Second, studies specifically addressing the rotor of electrically excited machines, characterized by a continuous, DC-fed excitation winding acting as a concentrated heat source, are considerably scarcer than those concerning permanent-magnet rotors or, more generally, stator windings, for which two-phase immersion cooling has already been investigated through a broader set of experimental and numerical works. Third, the combined effect of key operating parameters, such as inlet subcooling, mass flow rate, and operating pressure, on the maximum winding temperature and pressure drop has rarely been examined for multiple dielectric refrigerants within a unified computational framework, sharing identical geometry, boundary conditions, and numerical methodology. This last aspect is further compounded by the increasingly restrictive requirements introduced by the European F-gas Regulation (EU 2024/573) on the use of high-Global-Warming-Potential refrigerants, which calls for an extension of the comparative analysis to a wider range of low-GWP dielectric fluids. The present thesis is positioned to address this research gap by providing a rigorous, internally consistent numerical comparison between a single-phase reference configuration and four candidate two-phase dielectric refrigerants (Novec 7000, R134a, Opteon SF-33, and FC-72) for the direct cooling of an EESM rotor channel, thereby contributing quantitative, design-oriented insight to a research area in which systematic multi-fluid benchmarking data are, to date, still limited.

1.2 Objectives of the Thesis

In light of the growing demand for increased power density and energy efficiency in electrical machines intended for automotive applications, thermal management of the rotor in Electrically Excited Synchronous Machines (EESMs) plays a key role in ensuring system reliability, performance, and durability. Within this context, the present research aims to thoroughly investigate the potential of direct cooling using dielectric fluids, with particular emphasis on two-phase applications, through the development of an advanced numerical model capable of describing the main thermofluid-dynamic phenomena governing the heat removal process within the rotor. The first objective of this thesis is to characterize the thermal problem associated with the EESM by analyzing the loss sources that contribute to the increase in rotor temperature and, in particular, the Joule losses generated within the excitation windings. This preliminary phase is aimed at understanding the heat generation mechanisms and identifying the regions most susceptible to local overheating, which may compromise the integrity of the insulating materials and limit machine performance. In parallel, a computational domain representative of the actual physical configuration is defined and appropriately simplified in order to reduce the computational burden without compromising the significance of the results obtained. In particular, the domain under investigation consists of a double cooling channel located between two adjacent pole expansions, which is considered sufficiently representative of the thermal and fluid-dynamic conditions occurring within the rotor.

Subsequently, the research focuses on the development and validation of a three-dimensional CFD model implemented in Ansys Fluent, aimed at simulating conjugate heat transfer between the heat sources distributed within the windings and the coolant flowing through the channel. The model is designed to accurately represent both heat conduction within the solid components and convective heat transfer within the fluid, thereby allowing the overall thermal behavior of the system to be assessed under realistic operating conditions. For configurations involving phase change, the modelling activity also includes the description of liquid-to-vapor evaporation processes, which are particularly relevant due to the high heat removal capability associated with the latent heat of vaporization. For this purpose, a multiphase approach available within the software is employed and analyzed, namely the Mixture model coupled with the Lee mass-transfer mechanism, assessing its accuracy, numerical stability, and ability to correctly represent the expected boiling regime within the cooling channels.

Once the numerical methodology has been defined, the analysis of the reference single-phase case is carried out, representing the starting point for evaluating the performance of the cooling system. Through this simulation, the three-dimensional temperature distribution within the windings and surrounding components, the location of hot spots, the effectiveness of heat removal along the fluid flow path, and the pressure drop associated with the flow through the channel are determined. The analysis of the single-phase case also provides a quantitative benchmark against which the results obtained from the subsequent two-phase configurations can be rigorously compared, highlighting the inherent limitations of sensible-heat-based cooling compared with cooling based on phase change.

The subsequent phase of the research extends the investigation to four different candidate dielectric refrigerants, namely Novec 7000, Opteon SF-33, FC-72, and R134a, selected on the basis of their thermophysical properties, electrical safety characteristics, and different saturation temperatures. For each fluid, the main boundary conditions, such as the thermal power generated by the windings, inlet temperature, and saturation temperature, are kept constant, while the operating pressure is varied to promote evaporation at the same temperature, thereby enabling a direct comparison of their respective performance. The objective is to assess the influence of refrigerant properties on heat transfer

phenomena, temperature distribution throughout the domain, and the degree of evaporation achieved as the fluid flows through the channels. At the same time, a parametric study of the mass flow rate is conducted to identify the optimal operating conditions. This analysis makes it possible to investigate the trade-off between improved thermal performance and increased pressure losses, identifying the mass flow rate capable of minimizing the maximum winding temperature without introducing excessive fluid-dynamic and energy-related penalties.

Finally, the results obtained are subjected to a systematic comparison using consistent and directly comparable performance indicators. In particular, the maximum temperature reached by the windings, the temperature distribution within the adjacent components, the vapor fraction at the channel outlets, the overall pressure drop, and the energy balance of the computational domain are analyzed. The combined interpretation of these parameters enables a quantitative assessment of the effectiveness of the different fluids and operating conditions investigated, allowing the most promising configurations for the considered application to be identified. Furthermore, comparison of the results makes it possible to highlight the main critical aspects of the system, the limitations arising from the modelling assumptions adopted, and potential future developments of the numerical model. Overall, while acknowledging the necessary simplifications introduced in the modelling approach, the achievement of the proposed objectives is intended to contribute to a better understanding of direct two-phase cooling phenomena applied to the rotor of EESMs and to the definition of useful design criteria for the development of high-power-density solutions. Although limited to a specific case study, this research addresses a scientific field that remains relatively underexplored compared with analogous stator cooling problems and aims to contribute to the advancement of the available knowledge regarding the use of two-phase dielectric fluids for the thermal management of next-generation electrical machines.

Chapter 2

Electrically Excited Synchronous Motors (EESMs)

2.1 Operating Principle

Electrically Excited Synchronous Motors (EESMs), also referred to in the literature as Wound Rotor Synchronous Motors (WRSMs), Wound Field Synchronous Motors (WFSMs), or, in more recent terminology associated with automotive traction applications, Externally Excited Synchronous Motors, belong to the family of three-phase synchronous machines and share their fundamental operating principle. The stator, which is structurally identical to that of a PMSM or an induction motor of comparable size, accommodates a distributed three-phase winding (or a concentrated winding in fractional-slot configurations) which, when supplied by a set of sinusoidal currents displaced by 120 electrical degrees through a three-phase PWM inverter, generates a rotating magnetic field at the synchronous speed

$$n_s = \frac{60 \cdot f}{p} \ [rpm]$$

where f is the fundamental frequency of the stator currents and p is the number of pole pairs of the machine. The rotor, magnetized by direct current through the excitation winding wound around the pole expansions, behaves as a rotating electromagnet whose magnetic field, under synchronous steady-state operating conditions, remains stationary with respect to the rotating stator field: the rotor magnetically “locks” onto the stator field and rotates at the same synchronous speed, without slip, unlike what occurs in induction machines.

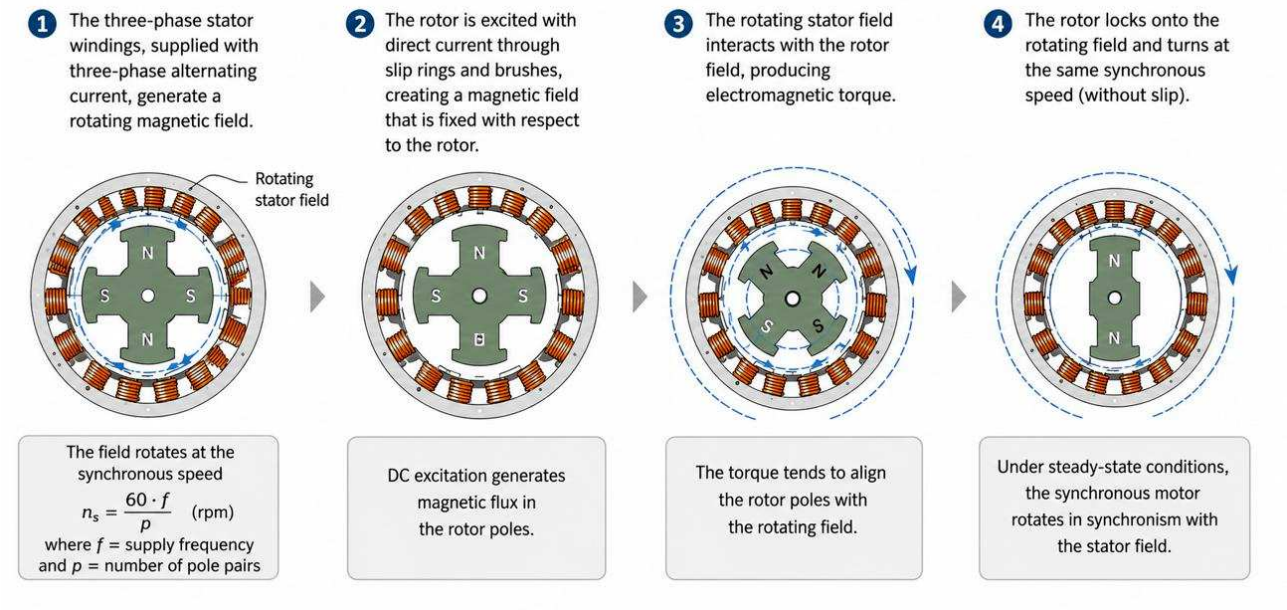


Figure 1 - EESM Motor Operating Principle

The electromagnetic torque produced by an EESM can be expressed, in the rotating reference frame fixed to the rotor (the d-q reference frame according to the Park transformation), as the sum of two distinct contributions:

$$T_{em} = \frac{3}{2} \cdot p \cdot \left[\Psi_f \cdot i_q + \left(L_d - L_q\right) \cdot i_d \cdot i_q\right]$$

The first term, proportional to the excitation flux Ψ_f (directly related to the excitation current I_f through the stator-rotor mutual inductance) and to the quadrature-axis component of the stator current i_q, represents the excitation torque, analogous to the torque produced by a PMSM. The second term, proportional to the difference between the direct-axis synchronous inductance L_d and the quadrature-axis synchronous inductance L_q, represents the reluctance torque, generated by the possible magnetic anisotropy of the rotor (saliency) resulting from the physical presence of the pole expansions and the interpolar spaces. Unlike surface-mounted PMSMs, in which $L_d \approx L_q$ and the reluctance torque is negligible, the salient-pole rotor typically employed in EESMs generally exhibits $L_q > L_d$, making the reluctance contribution non-negligible and providing an additional degree of freedom in the control strategy through the simultaneous optimization of i_d, i_q and the excitation current I_f in order to maximize torque for a given level of losses (a strategy known as extended Maximum Torque per Ampere, or Maximum Efficiency Control when the optimization of the rotor current is also included).

2.1.1 The Additional Degree of Freedom: Excitation Current

The truly distinctive feature of the EESM topology compared with other synchronous machines lies in the availability of an additional control degree of freedom—the excitation current I_f, which can be adjusted almost instantaneously through the power electronics controlling the rotor supply system. This degree of freedom is absent in PMSMs (where the magnetic flux is fixed by the permanent

magnets) and, although in a different manner, in induction motors and SynRMs. Within the physical limits of the machine, this additional degree of freedom makes it possible to decouple the magnetic flux level from the rotational speed, thereby enabling control strategies that would otherwise not be feasible:

- **Flux optimization under varying load conditions:** at low speeds and high loads, operation with a high excitation flux is advantageous in order to maximize torque for a given stator current (thereby minimizing stator Joule losses); conversely, under partial-load conditions, reducing the excitation flux makes it possible to minimize iron losses and overall Joule losses, resulting in efficiency improvements at low-load operating points, which account for the predominant share of operating time under real-world driving cycles.
- **Active field weakening at high speeds:** above the base speed, directly reducing the excitation current makes it possible to keep the induced electromotive force within the voltage limits imposed by the DC bus of the power conversion system, thereby extending the constant-power operating range without the need, except to a limited extent, to inject demagnetizing direct-axis stator current. In PMSMs, this technique represents the only available means of field weakening and results in additional Joule losses throughout the entire duration of field-weakening operation.
- **Power factor regulation:** by adjusting the excitation current, it is possible to control the phase angle between the stator voltage and current, allowing the EESM, when required by the application, to operate under overexcited conditions with capacitive reactive power exchanged with the grid or the supply inverter, a function that has no practical significance in a fixed-flux PMSM.

2.1.2 Systematic Comparison with Competing Technologies

PMSMs (Permanent Magnet Synchronous Motors): PMSMs generally exhibit higher torque and power densities for a given machine volume, owing to the high energy density of modern permanent magnets—typically based on Neodymium-Iron-Boron (NdFeB), with the addition of Dysprosium or Terbium to improve resistance to thermal demagnetization—as well as higher efficiency at low and medium speeds, since they are not affected by the rotor Joule losses typical of EESMs. On the other hand, the magnetic flux in PMSMs cannot be actively controlled: its magnitude is determined by the choice of permanent magnet and by the machine geometry, making field weakening through the injection of demagnetizing current necessary, with the previously discussed penalties, while also exposing the magnets to the risk of irreversible demagnetization under fault conditions (for example, in the event of a three-phase inverter short circuit, which may generate transient currents of sufficient magnitude to partially demagnetize the magnets if the operating temperature is already close to the allowable limit).

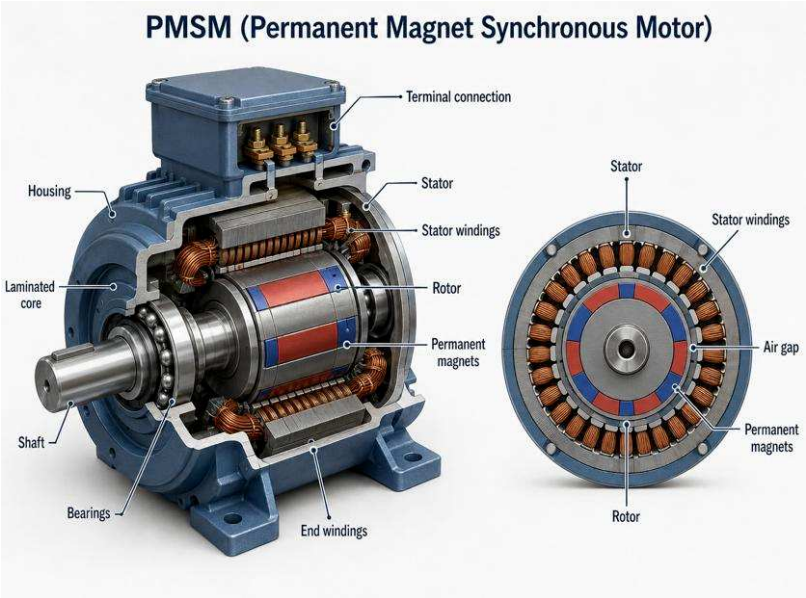


Figure 2 – PMSM Motor

Induction Motors (IMs): similarly to EESMs, induction motors do not require permanent magnets; however, the rotor magnetic field is generated through electromagnetic induction by exploiting the slip between the rotating stator field and the mechanical speed of the rotor, according to the principle of a rotating transformer. This results in Joule losses in the rotor bars (typically made of die-cast aluminum or, in high-performance designs, copper), which are intrinsically associated with the operating principle and cannot be completely eliminated even under no-load conditions, as well as a generally lower power factor compared with synchronous machines, due to the magnetizing current that must necessarily be supplied through the stator. Although EESMs also exhibit rotor Joule losses, these losses are concentrated within a DC-fed winding whose excitation can be controlled independently according to the operating point; furthermore, EESMs do not require any magnetizing current component from the stator and are therefore generally more efficient than induction motors under partial-load operating conditions.

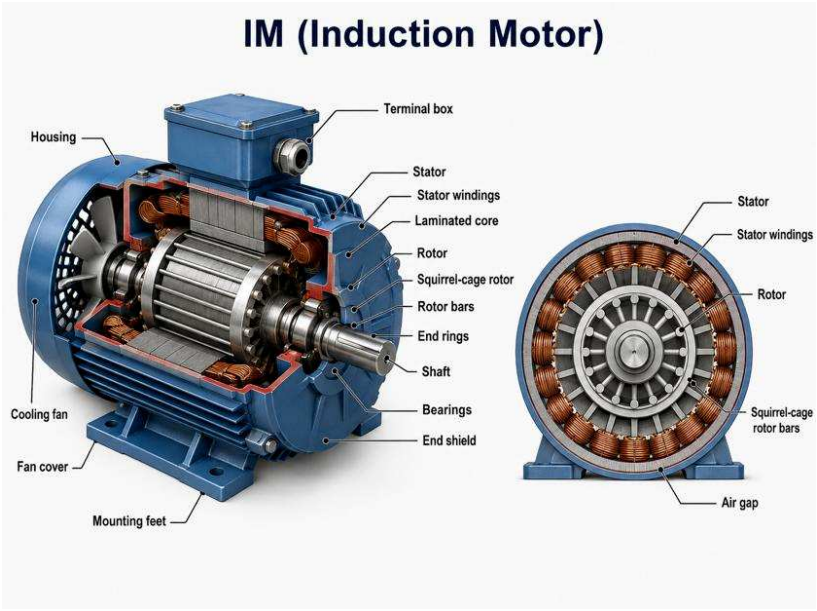


Figure 3 – IM Motor

Synchronous Reluctance Motors (SynRMs): SynRMs generate torque exclusively by exploiting the magnetic anisotropy of the rotor, without requiring either permanent magnets or rotor windings. Consequently, among the topologies considered, they are the most cost-effective, robust, and straightforward to cool, since they are free from rotor loss sources associated with electrical excitation. However, the reluctance-based torque generation mechanism is less effective in terms of torque density than that of excited synchronous machines, making SynRMs less competitive in applications where system volume or mass represents the dominant constraint. Permanent-Magnet-Assisted Synchronous Reluctance Motors (PMASynRMs) partially bridge this performance gap, although at the cost of reintroducing dependence on rare-earth materials.

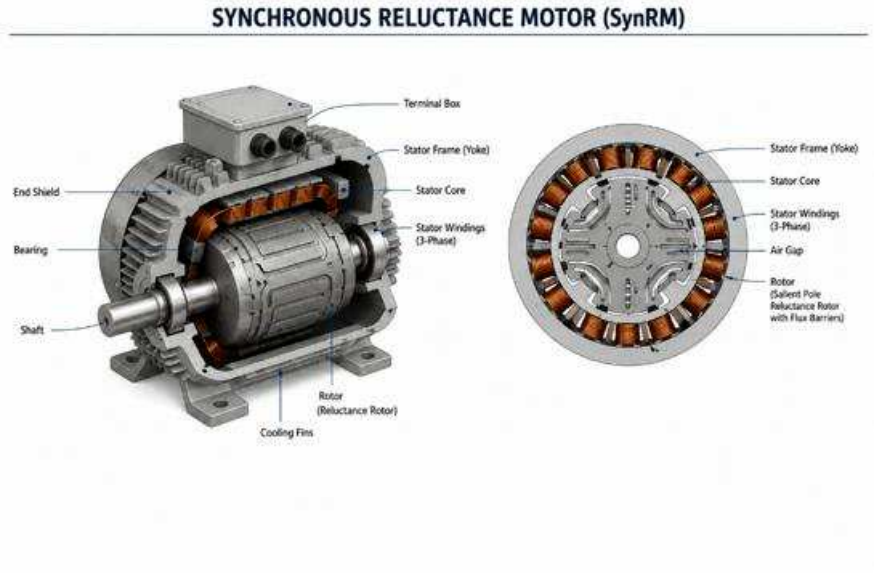


Figure 2 - Synchronous Reluctance Motors (SynRMs)

Table 1 summarizes the main characteristics of the electric motors discussed above.

Table 1 – Characteristics of Different Electric Motor Types

Criteria	PMSM	EESM	IM	SynRM
Torque density	Very high	High	Medium	Medium-low
Part Load efficency	High	High (with I_f optimization)	Medium-low	High
Flux controllability	None	Full	Indirect (slip)	Absent
Dependence on rare-earth materials	Yes	No	No	No (if not assisted)
Intrinsic rotor losses	None	Excitation Joule losses	Induced Joule losses	None
Rotor power supply system complexity	None	High (brushes or exciter)	None (or present only in wound rotor configurations)	None
Demagnetization risk	Present	Absent	Absent	Absent

In summary, EESMs occupy an intermediate position within the range of available technological solutions: they sacrifice the maximum torque density achievable with the highest-performing PMSMs and introduce rotor Joule losses that are absent in the latter; however, in return, they offer controllability of the excitation flux, independence from rare-earth materials, and greater inherent safety under fault conditions.

2.1.3 Why EESM Technology Is Gaining Ground

The growing interest in EESMs within the electric traction industry, following years of substantial dominance of PMSMs in the automotive sector, can be attributed to the convergence of several technical, economic, and geopolitical factors, which should be examined individually in order to fully understand the relevance of the topic addressed in this thesis.

- **Economic factors and rare-earth supply:** the high-performance permanent magnets employed in PMSMs require rare-earth elements (particularly Neodymium, Praseodymium, Dysprosium, and Terbium), whose extraction and refining are geographically concentrated in a limited number of regions worldwide, consequently exposing the automotive supply chain to risks associated with price volatility, export restrictions, and geopolitical constraints. The elimination of permanent magnets from the rotor, which is inherent to EESMs, removes this dependency at its source, providing manufacturers with a degree of strategic independence and cost stability that would otherwise not be available.
- **Sustainability and end-of-life considerations:** the extraction and processing of rare-earth elements have a significant environmental impact, both in terms of energy consumption and the generation of associated radioactive waste (rare-earth elements are often found in mineral deposits together with thorium and uranium). Furthermore, recycling permanent magnets at the end of the motor's service life is technically complex, whereas the copper used in the excitation windings of an EESM benefits from a mature, well-established, and economically efficient recycling supply chain, thereby improving the overall sustainability performance of the product from a Life Cycle Assessment perspective.
- **Field-weakening performance:** as discussed in Section 2.1.2, direct control of the excitation flux enables EESMs to achieve a wider and more efficient constant-power operating range than PMSMs. This aspect is becoming increasingly relevant for high-performance automotive applications, in which the motor must operate over a broad speed range—from standstill to highway speeds—while maintaining high efficiency even at high-speed operating points.
- **Safety under fault conditions:** the absence of a permanent magnetic flux source eliminates the phenomenon of uncontrolled braking torque (uncontrolled generator operation, or back-EMF) that may occur in PMSMs in the event of an inverter failure involving a short circuit of the stator phases. In an EESM, interrupting the power supply to the excitation winding is sufficient to reduce the residual torque to nearly zero almost instantaneously, representing a significant advantage for the functional safety certification (ISO 26262) of automotive traction systems.
- **Industrial adoption:** these advantages have found practical application in several mass-produced automotive products. Manufacturers such as Renault (with motors employed in the

Renault Zoe and the Megane E-Tech range) and, more recently, other European automotive groups have adopted EESMs for their traction systems, confirming the industrial, as well as academic, interest in this technology. This industrial adoption, together with the growing number of scientific publications devoted to the thermal and electromagnetic management of EESMs, supports the decision in this thesis to investigate rotor cooling in particular, which is recognized in the literature as one of the main factors limiting the further widespread adoption of this technology.

2.2 Rotor Structure

The rotor of an EESM consists of a laminated core—or, in some designs with a limited number of poles and relatively low rotational speeds, a solid forged or cast structure—shaped so as to form a number of pole expansions (or salient poles) corresponding to the number of poles of the machine, around which the excitation winding (field winding) is wound. Two main rotor configurations are generally distinguished in the literature:

- **Salient-pole rotor:** this is the most widely adopted configuration in low- to medium-speed applications (including automotive traction applications) and is characterized by clearly defined pole expansions separated by large interpolar spaces in which the excitation winding is accommodated. This configuration, which is the subject of the present thesis, exhibits high magnetic anisotropy ($L_q >> L_d$), which enhances the reluctance torque contribution discussed in Section 2.1.1, but is less suitable for extremely high rotational speeds because of the centrifugal stresses concentrated on the protruding pole structures.
- **Cylindrical or non-salient rotor:** this configuration is employed in large synchronous machines operating at high rotational speeds (e.g., turbogenerators), in which the excitation winding is accommodated in slots machined directly into the cylindrical surface of the rotor, thereby providing greater mechanical integrity at high peripheral speeds at the expense of lower magnetic anisotropy.

In the present thesis, reference is made to the salient-pole configuration, representative of low- to medium-power automotive traction applications, in which the interpolar ducts—bounded by the lateral surfaces of the excitation windings of two adjacent pole expansions, by the rotor yoke surface towards the shaft axis and, on the outer side, by the air-gap region—represent a natural location for the integration of the cooling channels investigated through the CFD analysis developed in the following chapters.

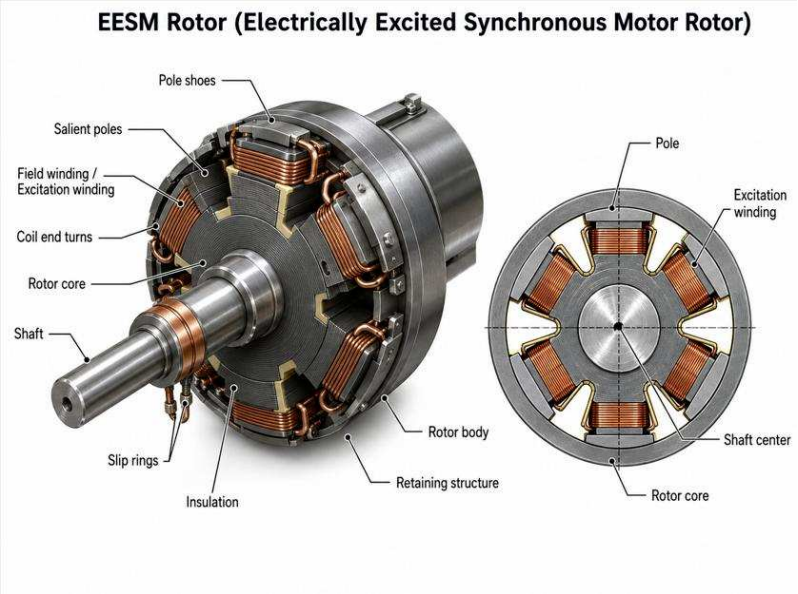


Figure 5 – EESM Rotor

2.2.1 The Excitation Winding

The excitation winding is typically made of copper conductors, either in the form of round enameled wire (for wound-coil windings) or rectangular/square-section conductors (for preformed-coil windings, or "pancake windings"), with the conductor cross-section and number of turns sized to provide the magnetomotive force $N_f \cdot I_f$ required to generate the excitation flux at the rated operating point, while complying with the thermal constraints imposed by the adopted insulation class (typically Class H, with a temperature limit of 180 °C, or Class C for more demanding applications) and with the space constraints imposed by the geometry of the rotor core.

The resistance of the excitation winding, a key parameter in the rotor Joule loss balance discussed in Section 2.4, depends on the total conductor length (which is a function of the number of turns and the mean length per turn, itself related to the perimeter of the pole expansion), the conductor cross-sectional area, and the electrical resistivity of copper. The latter is strongly temperature-dependent according to the relationship

$$\rho_{Cu}(T) = \rho_{Cu,20°C} \cdot [1 + \alpha_{Cu} \cdot (T - 20°C)]$$

with a temperature coefficient $\alpha_{Cu} \approx 0.00393\ K^{-1}$,, confirming that even a moderate increase in operating temperature results in a non-negligible increase in the generated Joule losses, according to the feedback mechanism already introduced in Section 2.3 and discussed in detail in Section 2.4.

The excitation winding is generally surrounded by an electrical insulation system (resin impregnation, composite-material tape insulation, or encapsulation) whose function, in addition to providing dielectric insulation to ground and between adjacent turns, is to ensure the mechanical integrity and consolidation of the winding against the centrifugal forces generated during rotation. These forces tend to load the turns radially outwards, particularly in the end-winding regions, which are not directly supported by the laminated core.

2.2.2 Rotor Excitation Supply Systems

The DC supply of the excitation winding, which represents both a distinctive feature and the main design complexity of EESMs compared with other synchronous machines, can be implemented according to two main approaches, each characterized by specific implications for the thermal balance and overall system reliability.

- **Supply through slip rings and brushes:** this is a relatively simple and low-cost design solution in which direct current is transferred from the stator (or from a stationary power circuit) to the rotor by means of two (or more) conductive rings mounted on and rotating with the rotor shaft. Brushes made of graphite-based composite material slide against these rings and are maintained in contact by preload springs. This solution presents several disadvantages related to the progressive wear of the brushes, the need for periodic maintenance and replacement, the generation of conductive graphite dust that may compromise the electrical insulation of other machine components, mechanical friction losses proportional to the peripheral speed of the rings, and a practical limitation on the maximum allowable rotational speed, beyond which there is a risk of brush lift-off due to centrifugal forces.
- **Supply through a brushless exciter with a rotating rectifier bridge:** this solution is more complex from a design and manufacturing standpoint but is free from maintenance requirements and wear associated with sliding contacts. In this configuration, a small auxiliary electrical machine—the exciter, typically designed as an inverted-pole synchronous machine (three-phase wound rotor, with either a permanent-magnet stator or a separately excited stator)—generates a set of three-phase alternating voltages in its rotor winding, which is mechanically coupled to and rotates with the shaft of the main machine. These voltages are rectified either by a diode bridge (the simpler solution, which does not provide direct active control of the excitation current except through indirect control of the exciter stator current) or by a bridge employing actively controlled semiconductor devices. Both types of rectifier are also mounted on the rotating shaft, thereby supplying the direct current required by the main excitation winding without any sliding electrical contacts.

In both cases, the excitation supply system introduces additional components into the rotor—slip rings and brushes in the first case, and the exciter and rectifier bridge in the second—which, together with the main excitation winding, contribute to the overall thermal balance of the rotor, occupy axial space that would otherwise be available for the active core of the machine, and impose constraints on the geometry and accessibility of the cooling channels. This aspect must be taken into account when defining the simplified computational domain adopted in the present work (see Chapter 4).

2.3 Advantages

The main advantages of EESMs, already discussed in the technological comparison presented in Section 2.1, can be summarized and further detailed as follows:

- **Controllability of the excitation flux:** this makes it possible to optimize the relationship between the torque produced and the losses generated over a wide operating range, to achieve efficient field weakening at high speeds, and to regulate the machine power factor by directly adjusting the excitation current, thereby providing an additional degree of control freedom that is unavailable in a fixed-flux PMSM.
- **Absence of permanent magnets:** this results in independence from the rare-earth supply chain, reduced exposure to the risks associated with the price volatility of these materials, and greater inherent safety under fault conditions, since there is no non-switchable magnetic flux source capable of generating uncontrolled braking torque.
- **Robustness at high temperatures:** the rotor is not subject to the risk of thermal demagnetization associated with permanent magnets, whose remanent magnetic flux density irreversibly deteriorates when critical temperatures are exceeded (typically 150–200 °C, depending on the magnet grade). This constraint is absent in EESMs, whose main thermal limit is represented by the insulation class of the excitation winding, which is generally more permissive.
- **Sustainability and raw material costs:** as discussed in Section 2.1.4, EESMs exhibit an overall more favorable life-cycle balance compared with PMSMs based on rare-earth permanent magnets.

2.3.1 Critical Issues

Despite these advantages, EESMs also present significant critical issues that strongly influence their design choices and, more specifically, provide the motivation for the present thesis:

- **Rotor Joule losses:** these are generated by the Joule effect in the excitation winding carrying the direct current required to produce the rotor magnetic flux. These losses, which are absent in PMSMs, directly penalize the overall efficiency of the machine—typically by several percentage points compared with an equivalent PMSM, particularly under partial-load conditions, where optimization of the excitation current is not sufficient to completely eliminate this disadvantage—and, more importantly, represent a localized heat source within the rotor, a region that is inherently more difficult to cool than the stator for the reasons discussed in detail in Section 2.5.
- **Complexity of the rotor excitation supply system:** as discussed in Section 2.2.3, this requires the introduction of additional components, with associated costs and axial space requirements and, in the case of brushes, maintenance requirements that are incompatible with the maintenance-free objective typically pursued in automotive applications.
- **Lower torque density** for a given machine volume compared with the highest-performing PMSMs, as a direct consequence of the lower energy density achievable with an

electromagnetic winding compared with a high-performance permanent magnet. Therefore, for a given torque requirement, a larger active volume or more intensive cooling is required in order to sustain higher current densities.

- **Thermal constraints as a limiting factor in machine sizing:** the maximum allowable current density in the excitation winding—and, consequently, the maximum continuous torque that can be delivered by the machine—is ultimately determined by the ability of the cooling system to maintain the winding temperature within the limits imposed by the thermal class of the insulation system employed. This constraint becomes increasingly stringent as the machine is pushed towards higher power densities in an attempt to bridge the performance gap with PMSMs.

This last point represents the logical connection between the theoretical discussion presented here and the case study developed in the following chapters: improving rotor cooling performance does not merely represent an incremental refinement, but rather an enabling requirement for EESM technology to compete effectively with PMSMs in terms of power density, in addition to its already favorable characteristics in terms of raw material costs and safety under fault conditions.

2.4 Heat Sources in the Rotor: Joule Losses in the Windings, Iron Losses, and Mechanical Losses

The thermal balance of an EESM rotor is determined by the combined contribution of several loss-generation mechanisms, characterized by different magnitudes and spatial distributions. These mechanisms should be analyzed separately in order to understand their relative contribution and their relevance to the design and sizing of the cooling system.

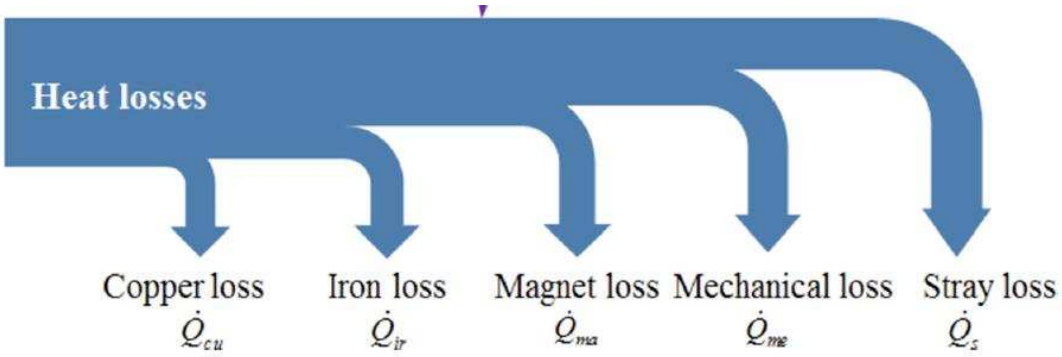


Figure 6 – Losses in an Electric Motor

2.4.1 Joule Losses in the Excitation Winding

In most EESM traction applications, Joule losses represent the dominant component of rotor losses, typically accounting for more than 80–90% of the total heat generated within the rotor under full-load operating conditions. As a first approximation, these losses can be described by the relationship

$$P_{Cu,rot} = R_f(T) \cdot {I_f}^2$$

where $R_f(T)$ is the electrical resistance of the excitation winding at the operating temperature, as defined in Section 2.2.2, and I_f is the DC excitation current, whose value may reach several tens of amperes under full-torque operating conditions, depending on the machine size and the number of winding turns. Since the resistance R_f increases with temperature, a positive feedback mechanism between heating and losses is established, as previously anticipated: an increase in temperature raises the winding resistance and, for a given current imposed by the control system, increases the generated losses, thereby further intensifying the heating process. In the absence of an adequate cooling margin, this mechanism may lead to a condition of thermal instability (thermal runaway) if, as the temperature increases, the generated thermal power exceeds the heat removal capability of the cooling system. This highlights the importance of robust cooling-system design and sizing under transient overload conditions as well as under steady-state operation.

2.4.2 Rotor Iron Losses

Rotor iron losses include magnetic hysteresis losses and eddy-current losses (induced currents) within the ferromagnetic material of the rotor, generated by the time variation of the magnetic field experienced by a given point of the rotor itself, according to the classical Steinmetz formulation for iron losses:

$$P_{Fe} \approx k_h \cdot f^{\alpha} \cdot B^{\beta} + k_e \cdot f^2 \cdot B^2$$

where k_h and k_e are material-specific coefficients (associated with hysteresis losses and eddy-current losses, respectively), f is the frequency of the magnetic field variations, and B is the amplitude of the magnetic flux density. Under ideal synchronous steady-state operation, in the absence of magnetic field harmonics, the rotor experiences a magnetic field that is constant over time in its own rotating reference frame, and rotor iron losses would therefore be zero ($f = 0$ in the rotor reference frame). In practice, however, the presence of spatial permeance harmonics caused by stator slotting, local magnetic saturation at the tips of the pole expansions, and possible manufacturing asymmetries induces time-varying magnetic field components even in the rotor reference frame. These components generate iron losses that are generally limited compared with excitation Joule losses (typically accounting for less than 10–15% of the total losses), but become non-negligible at high rotational speeds, since the eddy-current component is proportional to the square of the frequency.

2.4.3 Mechanical Losses

Mechanical losses include bearing friction losses (which depend on rotational speed, preload, and lubricant viscosity) and aerodynamic losses (windage) resulting from the interaction between the rotating surfaces of the rotor and the fluid—air or, in the case considered here, the cooling fluid—present within the air gap and the interpolar ducts. Although these losses do not directly contribute to the heating of the windings, they contribute to the overall thermal balance of the machine and, with regard to the aerodynamic component, increase approximately with the cube of the rotor peripheral speed under turbulent flow conditions, according to empirical relationships of the form

$$P_{windage} \propto C_f \cdot \rho_{fluido} \cdot \omega^3 \cdot r^4$$

where C_f is a friction coefficient depending on the local Reynolds number and the air-gap geometry, ρ_{fluido} is the density of the fluid present within the air gap, ω is the angular velocity, and r is the outer radius of the rotor. In the case of direct rotor cooling using a liquid dielectric fluid, as considered in the present thesis, this contribution becomes particularly relevant because the density and viscosity of the cooling fluid (which are several orders of magnitude higher than those of air) may significantly increase windage losses compared with air cooling. Although this aspect is not directly modelled in the present study, which focuses exclusively on the thermal balance of the cooling channel, it should be acknowledged as a potential direction for future development of the analysis.

For the purposes of the present thesis, the dominant contribution and the one of direct relevance to the design and sizing of the cooling system is represented by the Joule losses in the excitation winding (Section 2.4.1). In the numerical model developed in the following chapters, these losses are introduced as a volumetric heat generation condition applied to the winding domain, in accordance with widely adopted practices in the literature for the simplified CFD modelling of electromagnetic components. In such approaches, the detailed electromagnetic calculation—typically performed separately through magnetic Finite Element Method (FEM) simulations—provides the dissipated power input that is subsequently used as a heat source in the fluid-dynamic model.

2.4.4 Industrial Context: Adoption of EESMs

In recent years, the automotive electric traction sector has shown growing interest in motor technologies that do not rely on rare-earth permanent magnets, due to the strong geographical concentration of the production of these materials and the resulting exposure of automotive manufacturers to supply risks and cost volatility. Within this context, the Electrically Excited Synchronous Motor (EESM), in which the rotor magnetic field is generated by a current-carrying winding rather than by permanent magnets, represents one of the main technological alternatives currently undergoing industrial adoption, together with other rare-earth-free solutions such as induction motors and reluctance motors.

Several manufacturers have already introduced EESMs, or conceptually similar solutions, in production vehicles. The Renault Group has adopted this technology across its E-Tech range, including the Renault Mégane E-Tech Electric, whose motor is identified by the manufacturer as an EESM [18]; the AmpR Medium platform on which the Mégane E-Tech is based is also shared by the Renault Scénic E-Tech [19] and the Nissan Ariya [20], the latter also being equipped with an

electrically excited motor whose rotor is cooled by means of an oil-spray system that delivers the fluid both to the stator and axially onto the rotor coils, with ventilation holes in the end ring to facilitate oil evacuation [21]. BMW has also recently introduced an electrically excited synchronous motor ("externally-excited synchronous") on the rear axle of the new iX3, the first model built on the Neue Klasse platform, which entered production in October 2025 [22]. In parallel, automotive component suppliers such as ZF have developed magnet-free doubly excited synchronous motors, with the stated objective of providing performance comparable to that of the permanent-magnet synchronous motors currently more widely adopted [23].

The Stellantis Group, within which the present thesis work was developed in collaboration with the company's research center, is also active in this technological field. Together with General Motors, Stellantis has entered into a collaboration with Niron Magnetics for the development of rare-earth-free magnets, an initiative partially supported by funding from the U.S. Department of Energy [24]. Within the framework of this collaboration, communications issued by Niron Magnetics also indicate a joint project aimed at developing induction and electrically excited motors for vehicles belonging to the Chrysler and Jeep brands, as part of a broader research effort focused on "variable-flux" technologies for rotor excitation [25]. These initiatives confirm that rare-earth-free motors, and the EESM family in particular, represent an area of technological development actively pursued by Stellantis as well, consistently with the broader industry trend described above.

A common feature of most EESM solutions currently reported in the technical literature, and one that is of direct relevance to the present thesis, concerns the rotor cooling strategy. Recent scientific literature highlights that, unlike direct stator cooling, which is already widely established, direct shaft-side rotor cooling has not yet been documented for EESMs, either in traction applications or in high-power electrical machines. The solutions currently employed are predominantly based on oil-spray systems directed towards the rotor end-winding regions, sometimes supplemented, in higher-power machines, by evaporative spray cooling systems applied to the stator end-winding regions and the rotor air gap [26]. Within this framework, the two-phase cooling approach investigated in the present thesis, based on the evaporation of a refrigerant directly within channels integrated into the rotor body, represents a research direction that remains relatively unexplored in the literature for this type of electrical machine and may potentially complement the oil-spray solutions currently more widely adopted in production applications.

2.5 Why Rotor Cooling Represents the Main Critical Issue in This Technology

The previous sections have outlined a framework in which the overall performance of an EESM—in terms of maximum deliverable torque, power density, and efficiency—is intrinsically constrained by the ability of the cooling system to thermally manage the Joule losses generated within the rotor excitation winding. To conclude this chapter, it is appropriate to examine in greater detail the reasons why this constraint is more critical in the rotor than in the corresponding stator thermal management problem, thereby justifying the particular attention devoted to this topic in the present thesis.

2.5.1 Heat Dissipation Path

The heat dissipation path from the rotor to the external environment is inherently more complex than that of the stator. While the stator is in direct contact, through the machine housing, with the external cooling systems (water jackets, cooling fins, forced ventilation)—solutions that are relatively straightforward to implement since they are applied to a stationary component—the heat generated within the rotor must necessarily be dissipated through one of the following two paths: it may cross the air gap, a thin region (typically less than one millimeter in small- to medium-sized machines) occupied by air or by the cooling fluid, characterized by a non-negligible thermal resistance and by a convective heat transfer coefficient that strongly depends on the flow regime (often laminar or transitional, resulting in relatively modest heat transfer coefficients); alternatively, the heat may be removed locally by a coolant that directly reaches the rotor, with the resulting system-level complexities associated with the need to introduce and/or extract the fluid through a rotating component by means of rotating seals or rotating unions. These components are subject to wear, additional pressure losses, and particularly demanding sealing requirements when low-viscosity and highly volatile fluids are employed, such as the two-phase dielectric refrigerants considered in the present work.

2.5.2 Centrifugal Force Constraints

The presence of centrifugal forces, which are proportional to the square of the rotational speed and to the radius, imposes additional constraints on the mechanical design of rotor cooling channels and on the materials that can be employed. These constraints may prevent or complicate the adoption of solutions—such as small-cross-section channels, additional components with limited mechanical strength, or phase change material (PCM) interfaces sometimes employed in stator applications—that could otherwise be readily implemented in a stationary component. The design and sizing of rotor cooling channels must therefore simultaneously satisfy thermofluid-dynamic requirements (an adequate flow cross-sectional area to limit pressure losses) and structural requirements (fatigue resistance to the cyclic centrifugal stresses associated with speed variations during the operating cycle). In practical design, this trade-off significantly limits the available geometric degrees of freedom compared with an equivalent stator cooling channel.

2.5.3 Thermal Feedback of Joule Losses

The temperature dependence of the excitation winding resistance, discussed in detail in Section 2.4.1, introduces a positive feedback mechanism between temperature and the generated losses. This effect is absent—or at least of lower practical significance, owing to the greater cooling capacity generally available for the stator—in the stator winding. Inadequate rotor cooling therefore does not merely result in a higher steady-state temperature for a given imposed excitation current, but may also trigger, during overload transients or in the event of a partial failure of the cooling system, a progressive thermal runaway that can potentially cause destructive damage to the winding insulation. In practical machine design, this phenomenon requires the adoption of adequate thermal safety margins and, in

some cases, monitoring systems and active limitation of the excitation current based on the estimated or measured winding temperature.

2.5.4 Direct Relationship Between Thermal and Electromechanical Performance

From a purely performance-oriented perspective, since the maximum torque that can be continuously delivered by an EESM is, as previously discussed, limited by the maximum allowable current density in the excitation winding while ensuring compliance with the thermal class of the insulation system, any improvement in the rotor heat removal capability directly translates into an increase in the torque and power that can be continuously delivered by the machine, for a given overall size and active mass. This relationship can be qualitatively expressed by observing that, for a given winding geometry, the maximum allowable excitation current $I_{f,max}$ is constrained by the thermal equilibrium condition

$$R_f(T_{max}) \cdot {I_{f,max}}^2 \leq Q_{removable}$$

where $Q_{removable}$ represents the maximum thermal power that can be dissipated by the cooling system under the considered operating conditions while maintaining the winding temperature below the limit T_{max} imposed by the insulation class. It can therefore be readily observed that an increase in $Q_{removable}$ —which, as discussed in Chapter 3, can be achieved by transitioning from single-phase cooling to two-phase cooling capable of exploiting the latent heat of vaporization of the refrigerant— allows $I_{f,max}$ to be directly increased and, consequently, increases the excitation flux and the maximum torque that can be delivered by the machine.

This direct relationship between thermal performance and electromechanical performance represents the fundamental technical and scientific motivation underlying the present thesis. Through the numerical analysis presented in the following chapters, this work aims to investigate the extent to which the transition from single-phase rotor cooling to two-phase cooling, implemented using different candidate dielectric fluids, can contribute to overcoming this limitation, thereby at least partially bridging the performance gap that currently separates EESMs from the highest-performing PMSMs while simultaneously preserving the inherent advantages of EESM technology in terms of cost, sustainability, and safety discussed throughout this chapter.

Chapter 3

Cooling Techniques for Electric Motors

3.1 Indirect Cooling (Housing, Water Jacket)

Indirect cooling represents the simplest and most widely adopted cooling solution in electrical machines for automotive traction applications. It is characterized by the fact that the coolant never comes into direct contact with the internal heat sources of the machine (windings and laminated core), but instead removes heat through one or more intermediate solid interfaces, typically the external motor housing. The most common configuration involves the integration of a cooling jacket (water jacket) into the housing, consisting of helical, axial, or serpentine channels manufactured by casting, subtractive machining, or the assembly of separate components (for example, a double-wall housing with a welded or brazed spiral insert). A single-phase coolant flows through these channels—typically a mixture of water and ethylene or propylene glycol in varying proportions (commonly 50/50), selected for its antifreeze and anticorrosion properties and its high specific heat capacity, typically ranging from 3.3 to 3.8 kJ/(kg·K), depending on the glycol concentration.

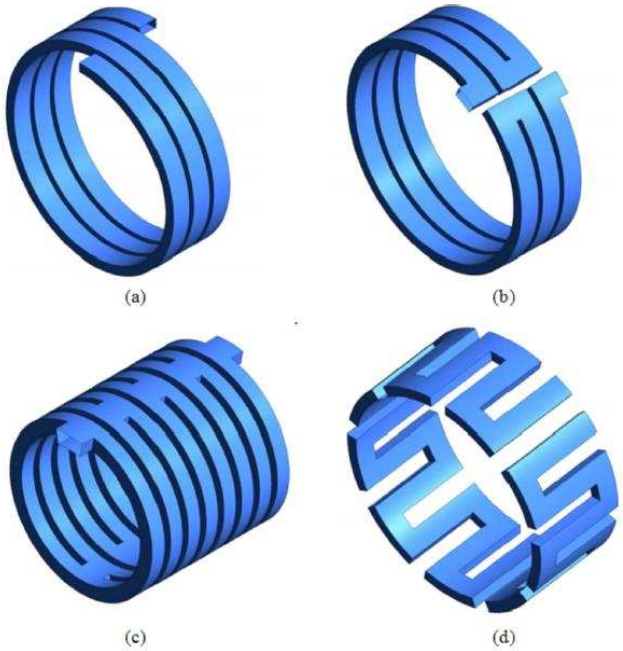


Figure 7 – Water Jacket

In the case of indirect stator cooling, the heat dissipation path sequentially passes through the heat source (typically Joule losses in the stator winding and stator iron losses), the electrical insulation system of the slots (slot liner), the stator laminated core, any resin impregnation interface, the housing, and finally the coolant circulating through the cooling jacket. Each of these interfaces introduces a thermal resistance in series, the sum of which determines the overall thermal resistance of the heat transfer path:

$$R_{tot} = R_{iinsulation} + R_{stack} + R_{housing\ contact} + R_{convectiv\ fluid}$$

where the term $R_{housing\ contact}$, associated with the mechanical contact between the laminated core and the housing (often achieved through an interference fit or shrink-fitting process), may represent a non-negligible thermal resistance—on the order of $10^{-4} - 10^{-3}$ m²K/W, depending on surface roughness, contact pressure, and the possible presence of microscopic air gaps—unless specific measures are adopted, such as the use of thermally conductive pastes or gap fillers, or optimization of the interference-fit profile during the design stage. The term $R_{convectiv\ fluid}$, in turn, depends on the flow regime established within the cooling jacket (laminar, transitional, or turbulent) and is typically estimated using empirical correlations such as the Dittus–Boelter or Gnielinski correlations, which are discussed in greater detail in Section 3.3.1.

3.1.1 Advantages and Limitations

Indirect cooling through a water jacket offers the advantage of considerable system simplicity: the coolant is confined within a circuit completely external to the electromagnetically active volume of the machine, without the need for seals separating the coolant from electrically active components, without the risk of fluid contamination of the windings, and without dielectric compatibility constraints on the coolant itself. Consequently, electrically conductive fluids such as conventional water-glycol mixtures can be employed. This solution is also mature, reliable, and economically advantageous, since it is already widely used in the automotive industry for internal combustion engine cooling and can be readily adapted, including in terms of circuit components (pumps, radiators, and expansion tanks), to electric traction systems, thereby reducing development and industrialization costs.

The main limitations of this solution, however, concern precisely its effectiveness in cooling the most critical heat sources. The rotor, in particular, is the component most adversely affected by this architecture, since the heat generated within it (excitation Joule losses in EESMs, magnet losses in PMSMs, or rotor-bar losses in IMs) must necessarily cross the air gap—a low-thermal-conductivity region typically filled with air, with a limited radial thickness (on the order of 0.5–1.5 mm in small- to medium-sized machines) but characterized by relatively low convective heat transfer coefficients due to the often laminar or transitional flow regime induced by rotation (Taylor–Couette flow within the annular gap, characterized by relatively low Nusselt numbers until the onset of Taylor instability at high Reynolds numbers)—before reaching the stator and, from there, the cooling jacket. This heat transfer path results in a high overall thermal resistance between the rotor heat source and the coolant, typically one or two orders of magnitude greater than that of the corresponding stator heat transfer path, thereby severely limiting the maximum sustainable rotor power density with this cooling architecture. This limitation is particularly significant for EESMs, in which, as discussed in Chapter 2, excitation Joule losses represent the dominant contribution to the overall thermal balance and for

which, consequently, indirect cooling is often insufficient to sustain the power densities required by the most demanding automotive traction applications, thus motivating the adoption of direct cooling solutions.

3.2 Direct Cooling: Hairpin/Oil Spray and Rotor Channel Cooling

To overcome the inherent limitations of indirect cooling, particularly in the management of rotor losses discussed in the previous section, recent scientific literature and the most advanced industrial applications have progressively adopted direct cooling solutions, in which the coolant comes into direct or nearly direct contact with the heat sources to be cooled, minimizing the number of intermediate solid interfaces and, consequently, the overall thermal resistance of the heat dissipation path. For a given thermal margin of the insulation system, direct cooling makes it possible to sustain significantly higher current densities in the windings, resulting—as discussed in Section 2.5.4—in a direct increase in the torque and continuous power that can be delivered by the machine.

3.2.1 Direct Stator Cooling: Hairpin Cooling and Oil Spray

On the stator side, two direct cooling techniques have become widely adopted in the latest generation of automotive traction machines. Oil spray cooling involves the direct injection, through calibrated nozzles distributed circumferentially, of a low-viscosity dielectric oil (typically an Automatic Transmission Fluid, ATF, appropriately selected for its dielectric properties and compatibility with electrical insulation materials) onto the stator end-windings. These regions are particularly critical because they are not in direct contact with the laminated core and, consequently, do not benefit from the primary heat dissipation path towards the housing. Upon impinging on the surface of the end-windings, the oil jet removes heat by forced convection and, to some extent, by surface evaporation in the case of oils containing volatile components. The oil subsequently drains by gravity into a collection sump and is recirculated through a dedicated circuit comprising a pump, filter, and heat exchanger.

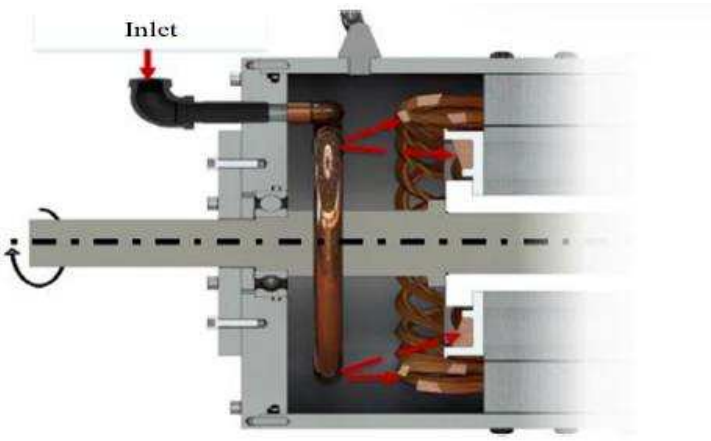


Figure 8 – Oil Spray Technique

Hairpin cooling, associated with the increasing adoption of preformed rigid hairpin-conductor windings (hairpin windings) in place of conventional round-wire windings, exploits their greater compactness and improved slot fill factor (typically above 70%, compared with 40–50% for round-wire windings) to enhance direct heat conduction from the conductor to the laminated core, reducing the thickness and, consequently, the thermal resistance of the insulation layer located between the conductor and the slot.

3.2.2 Direct Rotor Cooling: Internal Rotor Channels

On the rotor side—which is of direct relevance to the present thesis—direct cooling is implemented through cooling channels integrated within the rotor volume itself and positioned as close as possible to the heat sources to be cooled. In the salient-pole configuration typically adopted in EESMs (see Section 2.2.1), these channels can be implemented according to two main architectures:

- **Channels within the rotor shaft:** axial passages machined along the shaft axis, combined with radial branch passages that convey the coolant from the inlet—typically coaxial with the shaft itself and supplied through a rotating union located at one end—towards the peripheral regions of the rotor, often through radial holes intersecting the interpolar spaces. This architecture offers the advantage of concentrating the rotating sealing interface at a single axial location, but requires a larger shaft diameter to accommodate the internal passages, with potential implications for the torsional and flexural stiffness of the rotor.
- **Channels within the interpolar spaces:** channels located in the regions between adjacent pole expansions and bounded by the lateral surfaces of the excitation windings. This is the configuration adopted in the present thesis, in which the coolant flows directly (or nearly directly, in the presence of a thin electrical insulation layer or a polymeric partition, as in the model developed in Chapter 4) along the outer surfaces of the excitation winding, thereby minimizing the thermal resistance between the heat source and the fluid. This configuration is particularly effective because the coolant removes heat in the immediate vicinity of the dominant heat source (the excitation Joule losses; see Section 2.4.1). However, it requires the coolant to be distributed among multiple parallel channels (one for each interpolar space) through an appropriate supply manifold and provides less geometric design freedom because of the constraints imposed by the shape of the pole expansions and the space occupied by the winding.

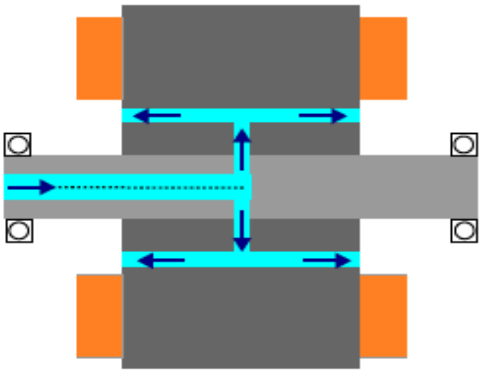


Figure 9 – Direct Oil Cooling

Introducing the coolant into a rotating region requires, as discussed in Section 2.5.1, the use of rotating seals or rotating unions at the interface between the stationary coolant supply circuit and the rotating shaft. These components represent one of the main system-level challenges associated with this architecture, particularly when low-viscosity, high-vapor-pressure two-phase fluids are employed, for which sealing against external leakage is more critical than for the high-viscosity single-phase fluids typically employed in indirect cooling. Additional system-level challenges include the need to balance the mass flow rate among the different parallel channels (especially in multipole configurations, where flow-distribution asymmetries may generate thermal non-uniformities among the different poles) and the management of the centrifugal forces acting on the fluid itself. In the case of a two-phase fluid, these forces may influence the radial distribution of the liquid and vapor phases within the channel, a phenomenon of potential interest for future developments of the present study but not explicitly modelled within the simplified two-dimensional axisymmetric domain adopted in Chapter 4.

3.2.3 Electrical Compatibility of the Coolant

A design aspect of primary importance in direct cooling, for both stator and rotor applications, is the need to employ dielectric coolants, i.e., electrically non-conductive fluids, whenever the coolant comes into direct contact with components at electrical potential (uninsulated windings or windings with a thin insulation layer). The use of a conductive fluid, such as a conventional water-glycol mixture (whose electrical conductivity, even in the absence of ionic additives, is typically on the order of $10^{-4} - 10^{-2}$ S/m and therefore entirely unsuitable for direct contact with energized conductors), in a non-insulated direct cooling configuration would result in the risk of short circuits to ground or between adjacent turns, thereby preventing its use except in configurations where an additional electrical insulation layer is interposed between the conductor and the fluid. Such a solution would reintroduce, although to a lesser extent than in indirect cooling, the thermal resistance that direct cooling is specifically intended to minimize.

This constraint motivates the selection of engineered dielectric fluids—such as hydrofluoroethers (HFEs) of the Novec family, fluorocarbons (FCs), and synthetic refrigerants of the Opteon family—as preferred candidates for direct rotor cooling. These fluids are characterized by dielectric strengths typically exceeding 10–15 kV/mm and electrical resistivities on the order of $10^{12} - 10^{16}$ Ω·cm, values fully compatible with direct or nearly direct contact with energized conductors. These aspects are discussed in greater detail in Section 3.5.

3.3 Single-Phase Cooling

In single-phase cooling, the coolant flows through the cooling channel while remaining entirely in the liquid phase (or, less commonly, in the gaseous phase), removing the heat generated by the thermal sources exclusively through sensible forced convection. The thermal power that can be removed by a single-phase fluid, within a section of the channel where the wall temperature T_w is higher than the fluid temperature T_f, is governed by Newton's law of cooling:

$$q'' = h \cdot (T_w - T_f)$$

where h is the convective heat transfer coefficient, which depends on the thermophysical properties of the fluid (thermal conductivity k, dynamic viscosity μ, density ρ, and specific heat c_p), the channel geometry (hydraulic diameter, D_h), and the flow regime, through dimensionless correlations of the form $Nu = f(Re, Pr)$, where the Nusselt number $Nu = \frac{h \cdot D_h}{k}$, the Reynolds number $Re = \frac{\rho \cdot v \cdot D_h}{\mu}$, and the Prandtl number $Pr = \frac{\mu \cdot c_p}{k}$ are involved. For fully developed turbulent flow in circular ducts or ducts with comparable cross-sectional geometries, correlations such as the Dittus–Boelter correlation are typically employed:

$$Nu = 0.023 \cdot Re^{0.8} \cdot Pr^{n}$$

with $n = 0.4$ for a fluid being heated and $n = 0.3$ for a fluid being cooled. Alternatively, for a broader range of validity in terms of Reynolds number, the Gnielinski correlation, based on the Darcy friction factor, may be employed. Under laminar flow conditions, which are not uncommon in small-cross-section rotor cooling channels operating at low flow rates, the Nusselt number asymptotically approaches a constant value depending on the cross-sectional geometry and the imposed thermal boundary condition (constant wall temperature or constant heat flux), resulting in heat transfer coefficients significantly lower than those achieved under turbulent flow conditions.

In the absence of phase change, the temperature rise of the fluid along the channel is governed by the overall energy balance:

$$Q_{tot} = \dot{m} \cdot c_p \cdot (T_{out} - T_{in})$$

where $\dot{m}$ is the mass flow rate, c_p is the specific heat capacity of the fluid at constant pressure, and Q_{tot} is the total thermal power removed. Since the specific heat capacity of fluids commonly employed in the liquid phase (water, oils, glycols, and dielectric refrigerants) is limited—typically ranging from 1.0 to 4.2 kJ/(kg·K)—the removal of high thermal loads requires significant mass flow rates, resulting in increased pressure losses (which, under turbulent flow conditions, are approximately proportional to the square of the velocity and therefore of the flow rate) and increased pumping power requirements. Alternatively, a significant increase in fluid temperature along the channel must be accepted, which may become excessive in applications characterized by high thermal power densities, generating pronounced axial thermal gradients and the occurrence of hot spots in regions farther downstream from the channel inlet—or, as observed in the preliminary simulations conducted in the present thesis, even within the inlet regions themselves, when heat generation is particularly concentrated and the fluid, although still relatively cold, has not yet developed fully established velocity and temperature profiles (the hydrodynamic and thermal entrance region,

characterized by local heat transfer coefficients higher than the average but by a lower heat removal effectiveness when the local heat generation exceeds the instantaneous heat removal capacity).

3.3.1 Two-Phase Cooling and Latent Heat

In two-phase cooling, the refrigerant is introduced into the channel under subcooled liquid conditions (at a temperature below the saturation temperature T_{sat} at the operating pressure) and, as a result of the heat absorbed from the hot surfaces, undergoes a liquid-to-vapor phase transition within the channel itself, typically through a flow-boiling mechanism that develops through successive regimes along the axial direction of the channel:

- **Subcooled boiling:** in regions close to the inlet, where the wall temperature locally exceeds the saturation temperature (thereby allowing vapor bubble nucleation near the wall), while the bulk fluid temperature is still below T_{sat}; the bubbles generated during this stage tend to rapidly recondense in the surrounding subcooled liquid;
- **Saturated boiling:** characterized by the formation, growth, and detachment of vapor bubbles that remain stable within the fluid, which has by then reached the saturation temperature, progressively increasing the vapor quality along the channel;
- **Flow regimes with progressively increasing vapor content:** according to conventional two-phase flow pattern maps, these typically evolve from bubbly flow to intermittent/slug or plug flow and, at higher vapor qualities, to annular flow, characterized by a thin liquid film in contact with the heated wall and a central vapor core. This regime is generally associated with the highest heat transfer coefficients due to convective evaporation of the liquid film.

The fundamental advantage of two-phase cooling lies in exploiting the latent heat of vaporization h_{fg} of the fluid, namely the energy required to convert a unit mass of saturated liquid into saturated vapor at the same temperature. This quantity is typically one or two orders of magnitude greater than the sensible heat that can be exchanged by the same fluid in the liquid phase over a temperature range representative of electromechanical applications (on the order of 10–30 °C of subcooling). The energy balance for a section of the channel in which the fluid undergoes partial vaporization can be expressed as

$$Q_{tot} = \dot{m} \cdot \left[c_{p,liq} \cdot (T_{sat} - T_{in}) + x_{out} \cdot h_{fg} \right]$$

where x_{out} is the outlet vapor quality of the considered channel section, defined as the mass fraction of vapor relative to the total mass flow rate. The second term, associated with latent heat, makes it possible to remove significantly greater thermal loads for a given mass flow rate compared with sensible cooling alone or, for a given thermal load to be removed, to substantially reduce the required mass flow rate and, consequently, the pressure losses and pumping power. This advantage is particularly relevant for rotor cooling, where, as discussed in Section 3.2.3, reducing the flow rate circulating through the rotating channels simplifies the design and sizing of the rotating seals and reduces the mechanical stresses acting on them.

A further advantage of two-phase cooling, which is particularly relevant for limiting hot spots, lies in the nearly isothermal nature of the saturated boiling process: as long as the fluid remains under saturation conditions (coexistence of liquid and vapor in thermodynamic equilibrium), its temperature

remains substantially constant and equal to T_{sat} throughout the entire section of the channel undergoing phase change, regardless of the amount of heat absorbed (for a given operating pressure), unlike the progressive temperature rise that characterizes single-phase cooling according to the energy balance discussed in Section 3.3.1. In principle, this characteristic makes it possible to achieve a more uniform temperature distribution over the cooled surface, thereby reducing axial and circumferential thermal gradients and, consequently, the cyclic thermomechanical stresses experienced by the windings, the insulation system, and any intermediate polymeric components (such as the insulating partition modelled in the present study; see Chapter 4).

3.3.2 Critical Issues in Two-Phase Cooling: Critical Heat Flux and Instabilities

Two-phase cooling, however, also presents specific critical issues that make its design and numerical modelling more complex than in the single-phase case. These aspects should be discussed in detail, as they are directly relevant to the interpretation of the results presented in Chapter 6.

- **Critical Heat Flux and dryout:** when the heat flux imposed at the wall exceeds a critical value, referred to as the Critical Heat Flux (CHF), a sudden deterioration of the heat transfer mechanism occurs. This behavior can be attributed to two distinct phenomena, depending on the local vapor quality. At low vapor fractions, the phenomenon known as Departure from Nucleate Boiling (DNB) is associated with the formation of an unstable vapor film that locally and temporarily isolates the wall from the liquid. At high vapor fractions, typical of the annular flow regime, the dryout phenomenon is associated with the progressive thinning and eventual rupture of the liquid film at the wall, resulting in direct exposure of the heated surface to the central vapor core, which is considerably less effective in removing heat. In both cases, exceeding the CHF results in a sudden increase in wall temperature for a given imposed heat flux (or, under imposed wall-temperature conditions, a sudden decrease in the heat flux that can be removed). This phenomenon must be carefully avoided when designing the cooling system by maintaining an adequate safety margin relative to the CHF estimated for the geometry and operating conditions under consideration.
- **Two-phase flow instabilities:** the compressibility of the vapor phase and the nonlinear relationship between pressure drop and mass flow rate under two-phase conditions can give rise to dynamic flow instabilities, such as density-wave oscillations. These phenomena are particularly relevant in systems comprising multiple parallel channels supplied by a common manifold—a configuration directly applicable to the rotor interpolar channels discussed in Section 3.2.3—in which a local perturbation of the vapor quality within one channel may induce a redistribution of the mass flow rate among the parallel channels, potentially leading to the onset of self-sustained oscillations.
- **Greater complexity of numerical modelling:** CFD simulation of two-phase cooling requires the adoption of dedicated multiphase approaches—such as the Mixture model coupled with appropriate interphase mass-transfer mechanisms (for example, the Lee model), as discussed in detail in Chapter 4—which are significantly more computationally demanding and more challenging in terms of numerical convergence than single-phase flow simulations. Such simulations typically require smaller time steps (in transient analyses), relatively low under-

relaxation factors, and careful verification of the overall energy balance as an indicator of the reliability of the numerical solution obtained.

- **Increased system-level complexity:** finally, the vapor generated within the cooling channel must be condensed and recirculated through a closed-loop circuit, requiring the introduction of a condenser (typically air-cooled or coupled to a secondary liquid cooling circuit) and, depending on the circuit architecture, either a compressor (in an active vapor-compression refrigeration cycle configuration) or a simple pump for the condensed liquid (in a two-phase forced-circulation configuration, which is more commonly adopted for the direct cooling of electromechanical components). This increases the overall system complexity compared with a conventional single-phase forced-circulation cooling loop.

3.4 State of the Art: Review of Scientific Studies on Two-Phase Cooling of Rotors and Stators

The scientific literature devoted to the direct two-phase cooling of electrical machines for traction applications has grown significantly over the past decade, in parallel with the drive towards increasingly higher power densities imposed by the requirements of automotive and aerospace electrification [2]. For the purposes of the present review, the available studies can be classified according to two main research directions: studies focused on two-phase stator cooling, which are more numerous and mature from both a numerical and experimental perspective, and studies devoted to two-phase rotor cooling, a more recent and less extensively investigated field within which the present thesis is positioned.

3.4.1 Two-Phase Stator Cooling

Regarding stator cooling, numerous studies have investigated the direct immersion of windings in two-phase dielectric fluids (immersion cooling), demonstrating significantly higher equivalent heat transfer coefficients compared with conventional indirect cooling techniques [15]. Particular attention has been devoted to the experimental characterization of nucleate boiling on geometries representative of stator windings (individual wires or bundles of horizontal/vertical conductors representative of end-windings or slot regions), the assessment of the impact of dielectric fluids on the electrical properties and long-term durability of the insulation system, and the experimental determination of boiling curves for different candidate dielectric fluids as a function of subcooling and operating pressure [9].

3.4.2 Two-Phase Rotor Cooling

With regard to two-phase rotor cooling, which represents the more specific focus of the present work, the available literature is comparatively more limited, reflecting the greater system-level challenges discussed in Section 3.2.3 (the need for rotating seals, centrifugal constraints, and mass flow distribution among multiple parallel channels), which have so far limited its experimental and industrial adoption compared with stator cooling. Existing studies are predominantly focused on CFD-based numerical analyses performed on simplified computational domains—similarly to the approach adopted in the present work, in which the domain is reduced to a single interpolar channel by exploiting the geometric and electromagnetic periodicity of the machine—aimed at characterizing the behavior of flow boiling in channels with geometries representative of rotor interpolar spaces, the sensitivity of thermal performance to the mass flow rate and inlet fluid subcooling, and the comparison of different candidate dielectric refrigerants in terms of the maximum winding temperature and the associated pressure drop [1].

A further research direction, complementary to CFD-based investigations, concerns coupled electromagnetic-thermal modelling studies (electromagnetic FEM and CFD/thermal co-simulation), in which the calculation of the losses (Joule, iron, and eddy-current losses; see Section 2.4) is performed separately within a magnetic finite-element environment and subsequently transferred to the fluid-dynamic model as a volumetric heat generation condition [6]. A simplified version of this approach is also adopted in the present thesis, in which the power dissipated within the excitation windings is imposed as an input to the CFD model developed in Ansys Fluent (see Chapter 4), without performing a direct electromagnetic co-simulation.

3.4.3 Positioning of the Present Work with Respect to the State of the Art

This second research direction, owing to its still-emerging nature and the limited availability of experimental validation data in the open literature—a direct consequence of the experimental complexity inherent in developing test rigs representative of a rotating rotor with two-phase fluid circulation—represents an area of scientific interest. The present thesis aims to contribute to this field through a systematic numerical analysis of the behavior of a rotor cooling channel representative of an EESM. More specifically, the present study is distinguished by the direct and systematic comparison, using the same computational domain, boundary conditions, and numerical methodology, between a reference single-phase configuration and four two-phase configurations employing four different candidate dielectric fluids (Novec 7000, R134a, Opteon SF-33, and FC-72). In light of the literature review presented above, such a multi-fluid comparison appears to be only partially addressed by the existing literature on rotor cooling and is conducted according to the selection methodology described in the following section.

3.5 Working Fluid Selection Criteria

The selection of the refrigerant represents a design decision of primary importance in the direct two-phase cooling of electromechanical components, as it must simultaneously satisfy thermofluid-dynamic, electrical, chemical, regulatory, and safety requirements. The main criteria adopted for the selection of the candidate fluids considered in the present thesis are as follows:

- **Dielectric properties:** since the cooling system involves direct or nearly direct contact with the excitation windings, the fluid must exhibit high dielectric strength (typically greater than 10–15 kV/mm) and very high volume electrical resistivity (on the order of $10^{12} - 10^{16}$ Ω·cm), sufficient to prevent electrical conduction or discharge phenomena in the vicinity of energized conductors. This requirement inherently excludes electrically conductive fluids (water and water-glycol mixtures) and directs the selection towards fluorinated organic fluids, which are well known for their excellent dielectric properties.

- **Saturation temperature and pressure:** the fluid must exhibit a saturation temperature T_{sat}, at the operating pressure of the circuit, compatible with the thermal operating level required by the application—in the present work, a target of $T_{sat} \approx 60$ °C, as discussed in Chapter 4—in order to ensure the onset of phase change within the temperature range allowable for the winding insulation. At the same time, this condition should be achieved without requiring excessively high operating pressures, which would impose additional requirements on the mechanical design of the circuit, rotating seals, and pressure-containing components, or excessively low pressures, i.e., pressures below atmospheric pressure, which could lead to the ingress of air and moisture through potential micro-leaks in the circuit, a particularly undesirable phenomenon in a dielectric system.

- **Latent heat of vaporization and transport properties:** a high latent heat h_{fg} maximizes the heat removal capacity for a given mass flow rate, as discussed in Section 3.3.2. The thermal conductivity, viscosity, and density of the liquid directly influence the convective heat transfer coefficient under both single-phase conditions (inlet and possible subcooled regions) and two-phase conditions (nucleation and convective evaporation), as well as the pressure losses associated with fluid flow through the channel, according to the Darcy–Weisbach equation appropriately extended to two-phase flow through homogeneous or two-fluid models (e.g., the Lockhart–Martinelli correlation).

- **Material compatibility:** the fluid must not chemically degrade the polymeric electrical insulation materials, protective sheaths, impregnation resins, or any rubber or elastomeric components present within the cooling circuit (static and dynamic seals and rotating unions). This aspect is assessed using chemical compatibility data provided by fluid manufacturers and, ideally, through long-term immersion testing of the candidate materials.

- **Environmental impact and current regulations:** the Global Warming Potential (GWP) and Ozone Depletion Potential (ODP) of the fluid must be considered in light of the European F-gas Regulation (EU Regulation 2024/573) and its progressively stricter restrictions on the use and commercialization of high-GWP refrigerants. This criterion is becoming increasingly relevant in the selection of refrigerants for new industrial and transportation applications and is driving research towards new-generation, low-GWP fluids.

- **Flammability and safety classification:** the safety classification of the fluid according to ASHRAE Standard 34 (or the equivalent ISO 817 standard), which defines its toxicity

(Classes A/B) and flammability (Classes 1/2L/2/3), represents an additional selection criterion of particular relevance for automotive applications, favoring Class A1 fluids (low toxicity and non-flammable).

3.5.1 Candidate Fluids Analyzed in the Present Work

Based on the criteria outlined above, the present thesis compares the two-phase cooling performance of four candidate dielectric fluids: Novec 7000, R134a, Opteon SF-33, and FC-72, belonging to three different chemical families of fluorinated refrigerants. Their main characteristics are summarized below.

- **Novec 7000** (methoxy-nonafluorobutane, $C_4F_9OCH_3$) [28], belonging to the hydrofluoroether (HFE) family developed by 3M, is a dielectric fluid widely employed for the direct cooling of electronic and electromechanical components. It is characterized by a boiling temperature at atmospheric pressure of approximately 34 °C, a relatively low GWP compared with other earlier-generation fluorinated fluids (on the order of hundreds and significantly lower than that of conventional HFCs), excellent dielectric properties (dielectric strength exceeding 30 kV with a 2.5 mm electrode gap according to ASTM D877), and excellent compatibility with most polymeric and elastomeric materials employed in electrical insulation. Novec 7000 was therefore selected as the reference fluid during the preliminary phase of the present thesis and serves as the primary benchmark for comparison with the other three candidate fluids.
- **R134a** (1,1,1,2-tetrafluoroethane, CH_2FCF_3) [29], a refrigerant widely used in HVAC&R applications and, historically, in automotive air-conditioning systems (where it is currently being progressively replaced by lower-GWP refrigerants as a result of F-gas regulations), has a particularly low boiling temperature at atmospheric pressure (approximately −26 °C). It therefore requires operating pressures significantly above atmospheric pressure (on the order of several bar) to achieve the target saturation temperature considered in the present study (~60 °C). Although it was not specifically designed for the direct cooling of electrical components—and, unlike the other three candidates, exhibits a lower volume electrical resistivity, although still compatible with low-voltage applications—its widespread commercial availability, low cost, extensive database of experimentally determined thermophysical properties available in the literature, and widespread use in the automotive sector (potentially facilitating integration with existing air-conditioning systems) justify its inclusion in the comparison as a benchmark against dedicated dielectric fluids.
- **Opteon SF-33** (developed by Chemours) [31], representative of the latest generation of lower-GWP synthetic refrigerants (hydrofluoroolefins, HFOs, or their blends), was specifically developed as a dielectric fluid for the direct cooling (“immersion cooling”) of high-power-density electronic and electrical components. Its thermophysical and dielectric properties are optimized for this application, including a boiling temperature close to the target values of direct cooling applications, high dielectric strength, and low viscosity, together with a more favorable environmental profile than earlier-generation fluorinated fluids, in line with the regulatory trends discussed in Section 3.5.1. These characteristics make it a candidate of particular interest for emerging applications such as the one investigated in the present thesis.

- **FC-72** (perfluorohexane, C_6F_{14}) [30], belonging to the fluorocarbon (FC, or perfluorocarbon, PFC) family also developed by 3M, is a dielectric fluid that has historically been employed—since the 1990s—for the direct cooling of high-power-density electronic components (e.g., processors and power electronic modules). It is characterized by a boiling temperature at atmospheric pressure of approximately 56 °C, close to the target thermal level of the present study (thereby allowing operation at pressures close to atmospheric pressure), as well as by dielectric properties and chemical inertness that are well characterized in the literature, making it one of the most widely employed fluids in experimental research on dielectric-fluid boiling. However, it exhibits a lower latent heat of vaporization than other fluids of the same family and a relatively high GWP due to the complete perfluorination of the molecule, an aspect that limits its attractiveness for future large-scale industrial applications despite its well-established experimental characterization.

3.5.2 Methodology for the Verification of Thermophysical Properties

For each of the four fluids considered, the methodology adopted in the present thesis (described in detail in Chapter 4) involves verifying the thermophysical properties implemented in the numerical model—density, viscosity, thermal conductivity, and specific heat capacity of both phases, latent heat of vaporization, and surface tension—against saturation curves and reference tables obtained from manufacturers' technical datasheets and established thermophysical property databases. The objective is to ensure that the target saturation temperature adopted for the comparison (~60 °C) corresponds, for each fluid, to the appropriate saturation pressure, correctly specified as the operating condition of the simulation (Fluent operating pressure). This procedure ensures full comparability of the results obtained in Chapter 6 for the four two-phase configurations analyzed, under identical thermal boundary conditions (dissipated power) imposed on the computational domain.

Chapter 4

Numerical Modelling and CFD Setup

4.1 Geometry and Computational Domain

This chapter describes the numerical methodology adopted for the CFD modelling of the rotor cooling channel investigated in the present study. The discussion begins with the definition of the geometry and computational domain, obtained by reducing the complete rotor to a representative portion by exploiting its geometric and electromagnetic symmetries. The solid and fluid components included in the model, the materials employed, and the method used to introduce the heat source associated with Joule losses in the windings are then described. The subsequent sections present the computational mesh generation process and the corresponding grid-independence assessment, the physical models adopted for the single-phase and two-phase simulations, the imposed boundary conditions, and the numerical solution strategy followed for each of the two types of case study.

4.1.1 Domain Reduction Criteria

Performing a Computational Fluid Dynamics simulation of the entire rotor of an EESM, including all poles, the corresponding excitation windings, and the interpolar cooling channels, would involve a computational cost that is not justified by the objectives of the present thesis, due to the large number of computational cells required to discretize a complete three-dimensional domain while providing adequate resolution in regions characterized by high temperature and vapor-fraction gradients (thermal boundary layers and wall liquid films under flow-boiling conditions). In the literature, it is well-established practice in the CFD simulation of electromechanical components characterized by geometric and electromagnetic periodicity to reduce the computational domain to a representative portion of the complete machine by exploiting cyclic (rotational) periodicity conditions and any mirror symmetries present in the geometry and boundary conditions.

In the present thesis, the computational domain was reduced to one-eighth of the complete rotor, identified on the basis of the geometric and electromagnetic symmetries of the machine. The rotational periodicity associated with the number of pole pairs makes it possible to isolate a repeating angular sector, which is further halved by exploiting the mirror symmetry of each pole with respect to its median plane, owing to the symmetrical distribution of Joule losses between the two half-windings facing the same interpolar channel. This reduction significantly decreases the number of

computational cells required and, consequently, the simulation time, without compromising solution accuracy, provided that appropriate periodic or symmetry boundary conditions are correctly imposed on the cut surfaces of the computational domain in order to represent the behavior of the portion of the machine that is not explicitly modelled (an aspect further discussed in Section 4.4).

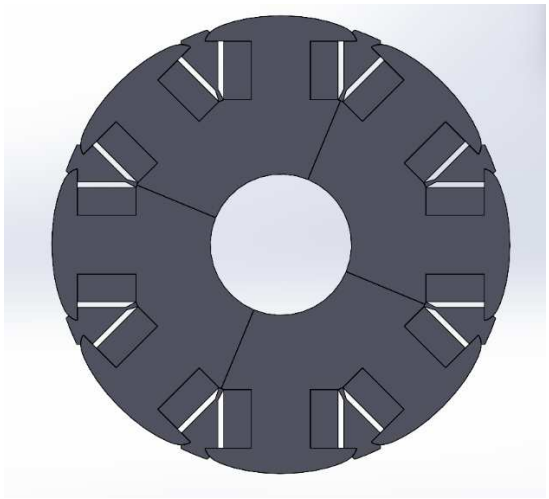


Figure 10 – 3D SolidWorks Model of the Rotor, Front View

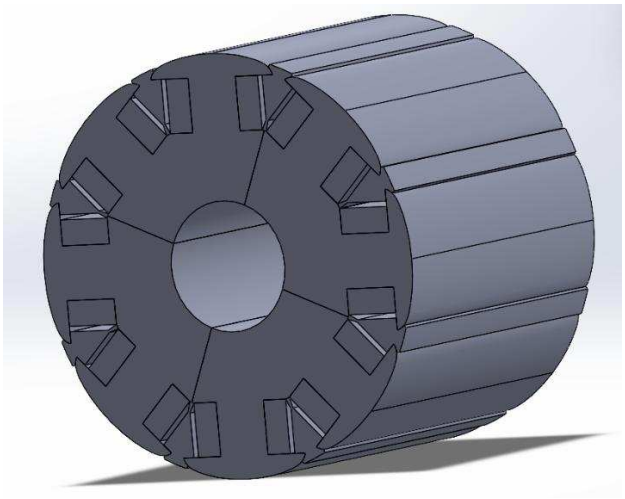
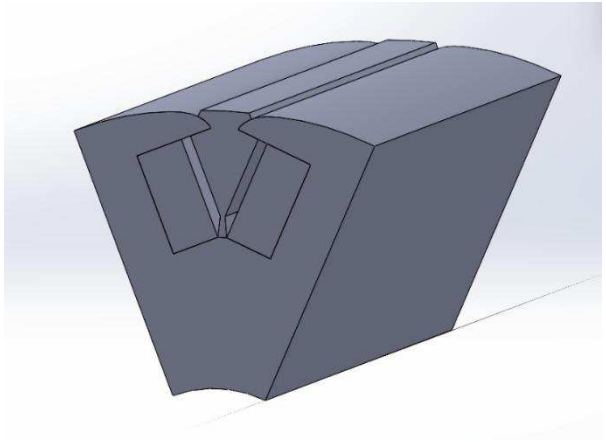


Figure 11 – 3D SolidWorks Model of the Rotor isometric view

4.1.2 Description of the Modelled Geometry

The computational domain, developed in SolidWorks and subsequently imported into Ansys Fluent for computational mesh generation and numerical solution, reproduces the rotor portion defined above and consists of the following solid components, which are shown in the isometric and front views in the following figures:



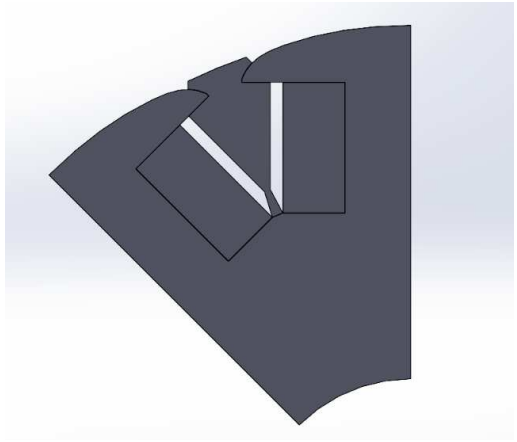
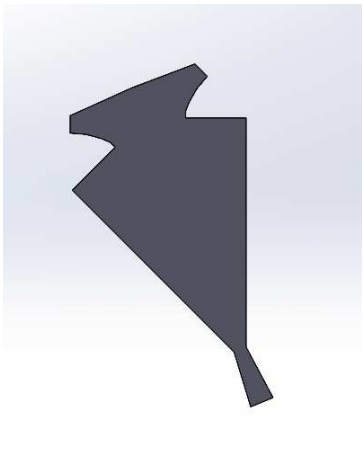


Figure 12 – Developed Rotor Computational Domain, Isometric View
Figure 13 – Analyzed Rotor Computational Domain, Front View

- **Ferromagnetic rotor structure:** corresponding to the largest component in the model, including the rotor yoke and the pole expansions that laterally bound the modelled angular

sector. This component constitutes the solid structural support of the rotor and provides the path for the magnetic flux towards the air gap. Furthermore, for the purposes of the present thermal analysis, it represents a heat conduction path towards the peripheral regions of the machine that are not directly exposed to the coolant.

- **Two copper excitation half-windings:** modelled as blocks with an essentially parallelepipedal cross-section, symmetrically arranged on the two sides of the central cooling channel and representing, according to the adopted symmetry cut (see Section 4.1.1), half of the complete excitation winding wound around each pole expansion. As discussed in Section 4.4, a volumetric heat generation term representative of the excitation Joule losses (see Section 2.4.1) is applied to these components.
- **Polymeric partition:** hereinafter referred to, due to its characteristic toothed profile visible in cross-section, as the "Christmas-tree" partition, positioned between the two copper half-windings and closing the cooling channels at the top. This component performs a dual function: on the one hand, it provides electrical insulation between the coolant and the energized windings, allowing the cooling channels to be positioned in close proximity to the copper without the risk of undesirable electrical conduction, even if the fluid is not perfectly dielectric or in the presence of potential impurities; on the other hand, it geometrically defines the flow cross-section of the two cooling channels, confining them between the surface of the partition itself, the lateral walls of the copper windings, and the inner surface of the ferromagnetic structure.



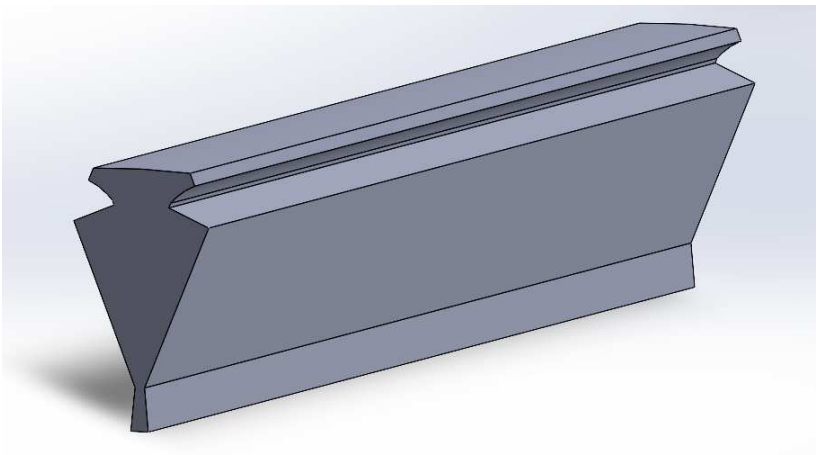


Figure 14 – Plastic Partition, Front View
Figure 15 – Plastic Partition, Isometric View

4.1.3 The Two Cooling Channels

As highlighted by the double-cavity geometry visible in cross-section (Figure 12-13), the polymeric partition defines two distinct and symmetrical cooling channels, each laterally bounded by the surface of one of the two copper half-windings and by the corresponding portion of the partition, arranged in a converging ("V-shaped") profile towards the symmetry axis of the domain. This dual-channel configuration reflects the need to symmetrically cool both half-windings facing the same interpolar space and, from the perspective of the boundary-condition setup (see Section 4.4), requires the definition of two separate inlet sections and two separate outlet sections, one for each channel. The two channels are supplied under nominally identical operating conditions (the same mass flow rate, inlet temperature, and inlet pressure) owing to the geometric and thermal symmetry of the domain.

In plan view (Figure 13), the overall geometry of the modelled sector has the shape of a circular wedge, bounded by the two radial cutting planes that define the angular extent of the isolated sector (one-eighth of the complete rotor; see Section 4.1.1) and by the inner and outer cylindrical surfaces corresponding, respectively, to the shaft surface and the rotor air-gap surface. Within this sector, the two rectangular-cross-section copper windings and the double-cusp geometry of the intermediate polymeric partition are clearly visible in projection, with the two cooling channels arranged symmetrically with respect to the bisector plane of the sector.

The decision to explicitly model the polymeric partition, rather than merely defining the fluid flow cross-section through a fluid control volume without an intermediate solid component, makes it possible to correctly account, in the subsequent thermal analysis (Chapters 5 and 6), for the conductive thermal resistance introduced by this component along the heat dissipation path between the copper windings and the coolant. The relevant thermophysical properties of this component (thermal conductivity of the polymeric material and partition thickness) are discussed in detail in Section 4.3, together with the definition of the physical models adopted for the simulation of conjugate heat transfer between the solid and fluid domains of the model.

4.1.4 Component Materials

For the purpose of modelling conjugate heat transfer between the solid domains and the coolant (see Section 4.3), each of the solid components described in Section 4.1.2 was assigned a specific set of thermophysical properties, summarized in Table 2, including thermal conductivity, density, and specific heat capacity at constant pressure.

- **Copper windings:** the copper excitation windings were modelled as a material with anisotropic thermal conductivity, with an axial thermal conductivity (along the direction of the conductor, k_{ax} = 129.5 W/(m·K)) significantly lower than the value expected for isotropic bulk copper (typically approximately 385–400 W/(m·K)), and an even lower radial thermal conductivity (k_{rad} = 1.07 W/(m·K)). This modelling choice reflects the inherently non-homogeneous nature of the actual winding, which consists of a bundle of individual enamelled conductors arranged in successive turns and separated, in the transverse direction, by thin layers of electrical insulation (enamel) and by possible air gaps or impregnation resin. While in the axial direction (along an individual conductor) heat conduction occurs through solid copper, resulting in relatively low thermal resistance, in the radial direction (transverse to the winding turns) heat must instead pass successively through multiple layers of low-thermal-

conductivity insulation, thereby significantly reducing heat transfer in this direction. The density (ρ = 1935 kg/m³, significantly lower than that of pure copper, 8960 kg/m³, due to the homogenization of the winding volume including the insulation regions and any voids) and specific heat capacity (c_p = 1035 J/(kg·K)) complete the characterization of the equivalent homogenized material.

- **Ferromagnetic structure (lamination):** the rotor laminated core, manufactured from M250-35A grade electrical steel (non-grain-oriented electrical steel with reference specific iron losses of approximately 2.5 W/kg at 1.5 T and 50 Hz, and a nominal sheet thickness of 0.35 mm), was also modelled with anisotropic thermal conductivity, consistently with the laminated nature of the component. The radial thermal conductivity (k_{rad} = 30 W/(m·K), representative of heat conduction within the plane of the laminations) is significantly higher than the axial thermal conductivity (k_{ax} = 5.73 W/(m·K), representative of heat conduction in the direction perpendicular to the laminations, i.e., through the stack of electrical steel sheets electrically insulated from one another by a thin insulating coating). This behavior is analogous to that observed for the copper windings, although with a less pronounced degree of anisotropy. The density (ρ = 7650 kg/m³) and specific heat capacity (c_p = 460 J/(kg·K)) are representative of typical Fe-Si alloys employed in electromagnetic applications.
- **Plastic insulation:** the polymeric partition described in Section 4.1.2 was modelled as an isotropic material with low thermal conductivity (k = 0.2 W/(m·K), a value representative of engineering polymeric materials such as polyamides or filled composite materials employed in electrical insulation applications), with a density of ρ = 1000 kg/m³ and a specific heat capacity of c_p = 1200 J/(kg·K). The low thermal conductivity of this component, together with its thickness, results in a non-negligible conductive thermal resistance along the heat dissipation path between the copper windings and the coolant. This aspect will be taken into account when interpreting the results presented in Chapters 5 and 6, particularly when comparing the maximum temperature recorded in the windings with that recorded on the outer surface of the partition itself in contact with the fluid.

Table 2 – Thermophysical Properties of the Materials Used in the Model

Material	**Conductivity**	**Density [kg/m³]**	**c_p [J/(kg·K)]**
Copper winding	k_{ax} = 129.5 W/(m·K); k_{rad} = 1.07 W/(m·K)	1935	1035
Lamination (M250-35A)	k_{ax} = 5.73 W/(m·K); k_{rad} = 30 W/(m·K)	7650	460
Plastic insulation	k = 0.2 W/(m·K)	1000	1200

4.1.5 Heat Source in the Copper Windings

The volumetric heat generation condition applied to the copper winding domain, representing the excitation Joule losses discussed in Section 2.4.1, was defined starting from the total loss value of the entire rotor, provided as 1420 W, corresponding to machine operation at 4200 rpm and half load under continuous operating conditions (S1 duty, according to the classification of duty types for rotating electrical machines).

Owing to the reduction of the computational domain to one-eighth of the complete rotor, including two half-windings (see Section 4.1.1), the share of Joule losses attributable to each modelled half-winding is equal to 88.75 W, obtained by proportionally distributing the total rotor losses over the number of equivalent half-windings present in the complete machine.

However, the value of 88.75 W represents the total Joule loss of each half-winding, including both the portion of the conductor located within the active stack (directly facing the cooling channel and considered in the present CFD model) and the end-winding portions, which extend axially beyond the laminated stack and in which a higher-than-average share of losses per unit length is concentrated due to the greater local electrical resistance and the generally less effective cooling conditions in these regions. Since the computational domain modelled in the present work (see Section 4.1.2) represents only the portion of the winding facing the cooling channel, excluding the end-windings, it was necessary to scale the 88.75 W loss value in order to determine the loss effectively attributable solely to the modelled portion of the conductor, resulting in a value of 72.7 W for each half-winding.

This scaling procedure was carried out by assuming, as a simplification, constant and uniform heat generation along the entire length of the modelled conductor (L = 136 mm; see Table 4.2), rather than a non-uniform distribution with a concentration of losses in the end-winding regions. Although this simplifying assumption does not accurately reproduce the actual distribution of losses along the conductor, it represents a reasonable and conservative approximation for the purposes of the present comparative study among the different cooling configurations (the single-phase case and the four two-phase cases; see Chapters 5 and 6), since it is applied identically to all the configurations analyzed, thereby ensuring full comparability of the results obtained.

The volumetric heat generation value imposed as a condition within the winding domain, obtained by dividing the scaled power of 72.7 W by the volume of the individual modelled half-winding (V = 27796.7 mm³; see Table 4.2), is equal to

$$\dot{q_{vol}} = 72.7\ W\ /\ 27796.7\ mm^3 \approx 2\,615\,808\ W/m^3$$

This value was imposed in Ansys Fluent as a uniform volumetric heat source (source term) in the energy equation solved within the solid winding domain, according to the formulation discussed in detail in Section 4.4.

Table 3 – Geometric and Thermal Data of the Winding (for Each Modelled Half-Winding)

Parameter	Value
Length [mm]	136
Thickness [mm]	9.85
Heigth [mm]	20.75
Volume [mm^3]	27796.7
End-winding length ($L_{end\ winding}$) [mm]	15
Total rotor Joue losses (4200 rpm, half load) [W]	1420
Half winding losses [W]	88.75
Distributed losses (model) [W]	72.7
Heat generation [W/m³]	2 615 808

4.1.6 Geometry Preparation in Ansys SpaceClaim

Once the three-dimensional modelling of the component had been completed in SolidWorks, the geometry was imported into Ansys SpaceClaim to undergo a preliminary preparation process required to make it suitable for the subsequent spatial discretization (meshing) stage. This operation is essential in CFD applications, since a geometry originally developed for mechanical design purposes is generally not directly compatible with the topological requirements of a finite-element or finite-volume solver. The preparation process carried out in SpaceClaim involved the following main operations.

- **Check Geometry:** an automatic geometry diagnostic tool was first used to identify and correct topological anomalies such as disconnected surfaces, small gaps, duplicate or degenerate faces, and geometric features of negligible size, whose presence could compromise the quality of the computational mesh or, in the most critical cases, prevent its generation altogether;
- **Share Topology:** topology sharing was subsequently applied among the different solid bodies forming the assembly (windings, lamination, plastic insulation, and cooling channel) in order to ensure perfect correspondence (conformity) between the meshes generated on the interface surfaces of adjacent volumes. This procedure is essential for the correct resolution of conjugate heat transfer between the different solid and fluid domains, avoiding the need for non-conformal interfaces and the associated numerical interpolation between independent meshes;
- **Creation of the Fluid Flow Volume:** using the Volume Extract command, the solid volume corresponding to the internal cavity of the cooling channel was generated, representing the domain actually occupied by the refrigerant fluid. This volume was obtained from the cavity

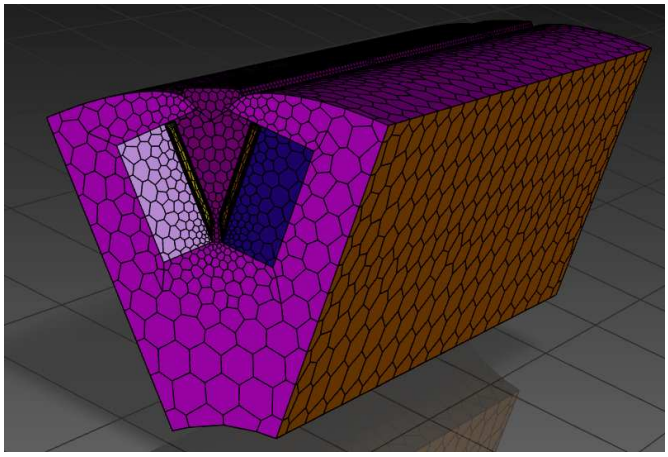
surface present in the original CAD model, which had been developed in SolidWorks as a solid body without an explicitly defined fluid volume;

- **Naming of Bodies and Surfaces:** finally, an appropriate designation (Named Selection) was assigned to each solid body and to each surface of interest within the domain (fluid inlet and outlet sections, walls exposed to the external environment, solid–solid interfaces, and solid–fluid interfaces), in order to uniquely identify the different regions of the model during the subsequent mesh-generation and boundary-condition setup stages in Fluent.

At the end of this preparation process, the geometry consists of a set of topologically conformal and correctly identified solid and fluid bodies, providing the necessary conditions for proceeding with the generation of the computational mesh described in the following section.

4.2 Mesh and Computational Grid Generation Strategy

The computational grid was generated in Ansys Fluent Meshing using a poly-hexcore approach, following the sequence of operations typically adopted in the Watertight Geometry workflow, including the definition of local mesh refinement (local sizing), surface mesh generation (surface mesh), the insertion of prismatic cell layers in the near-wall regions (boundary layers), and, finally, volume mesh generation. Each of these stages is governed by a set of control parameters whose influence on the quality and resolution of the resulting grid was investigated through the grid-independence study described in Section 4.2.3.



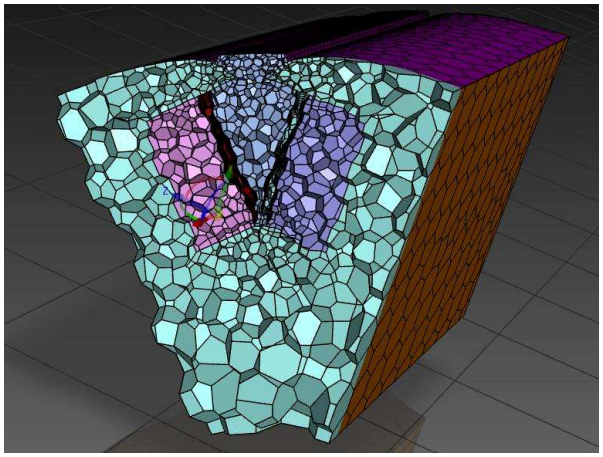


Figure 16 – Surface Mesh Generation, Isometric View
Figure 17 – Volume Mesh Generation Using Poly-Hexcore, Isometric View

During the local sizing stage, a growth rate of 1.2 was imposed for all the grid configurations analyzed in order to ensure a gradual transition in cell size between the more highly refined regions (in the vicinity of the walls and in regions where high thermal and fluid-dynamic gradients are expected, such as the interface between the fluid and the surface of the polymeric partition) and the coarser-grid regions, thereby avoiding excessive variations in cell size that could introduce numerical interpolation errors.

During the surface mesh generation stage, a target mesh size, a minimum allowable cell size, and a maximum allowable cell size were defined, together with the cells per gap parameter, which specifies the minimum number of cells to be generated across narrow regions of the domain (such as the fluid flow cross-section within the cooling channels), thereby ensuring adequate resolution even in the most geometrically confined regions.

During the generation of the near-wall layers (boundary layers), performed using the smooth transition approach, the number of layers to be generated in the vicinity of each solid wall, the transition ratio between the thicknesses of successive layers, and the number of transition layers to be introduced in correspondence with acute geometric angles (acute angle buffer layers) were defined. The latter parameter is particularly relevant because of the toothed-profile geometry of the polymeric partition (see Section 4.1.2), which is characterized by edges and fillet regions that may generate poor-quality cells in the absence of adequate local refinement.

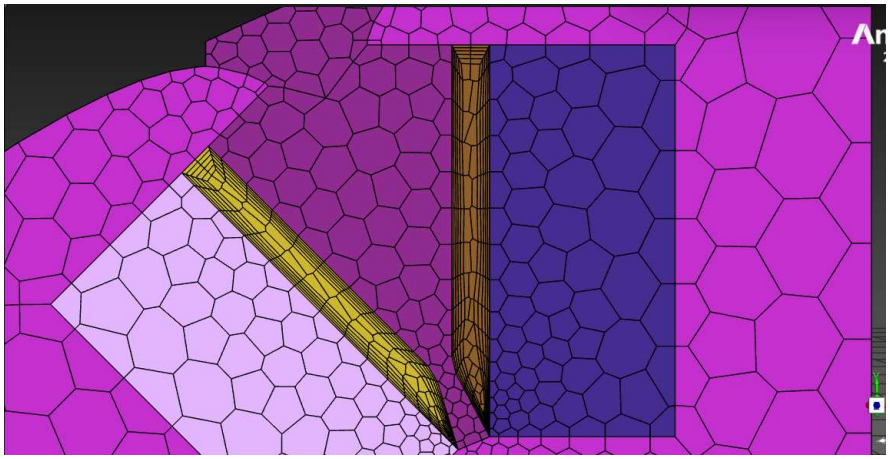


Figure 18 – Detail of the Boundary-Layer Mesh in the Fluid Domains, Front View

Finally, during the volume mesh generation stage, a maximum cell size (max cell length) was specified for the polyhedral cells generated within the core of the domain, outside the near-wall layers.

4.2.1 Grid Configurations Analyzed

In order to conduct the grid-independence study, three mesh configurations with progressively increasing resolution were generated by simultaneously varying the target size, the minimum and maximum surface mesh sizes, the number of near-wall layers, and the maximum cell size of the volume mesh, while keeping the growth rate (growth rate = 1.2) and the near-wall layer transition ratio (transition ratio = 0.272) unchanged across all three configurations. The distinguishing parameters of the three meshes, reported in Table 4.3, were designated, in order of increasing resolution, as follows:

- **Mesh 2 (coarse):** the coarsest configuration, with a target cell size of 0.13 mm, one cell per gap (cells per gap = 1), and 4 near-wall layers;

- **Mesh 1 (medium):** the intermediate configuration, with a target cell size of 0.1 mm, one cell per gap, and 5 near-wall layers;
- **Mesh 3 (fine):** the finest configuration, with a target cell size of 0.075 mm, two cells per gap (cells per gap = 2), and 7 near-wall layers.

Table 4 – Mesh Generation Parameters for the Three Mesh Configurations

Parameter	Mesh 2 (coarse)	Mesh 1 (medium)	Mesh 3 (fine)
Growth rate	1.2	1.2	1.2
Target mesh size [mm]	0.13	0.1	0.075
Minimum size [mm]	0.000173	0.00013281	0.0001
Maximum size [mm]	0.013	0.01	0.0075
Cells per gap	1	1	2
Boundary layers	4	5	7
Transition ratio	0.272	0.272	0.272
Acute angle buffer layers	2	2	2
Max cell length [mm]	0.0117	0.00902895	0.00695

The results of the meshing process, in terms of the number of nodes, cells, and faces generated for each configuration, are reported in Table 4.4. It can be observed that, while the transition from Mesh 2 to Mesh 1 involves only a moderate increase in surface and volume mesh resolution (+3.8% in the number of nodes and +8.0% in the number of cells, respectively), the transition to Mesh 3 results in a significantly greater increase in grid density, with more than 200% additional cells compared with Mesh 1. This increase is a direct consequence of the combined effect of reducing the target cell size and doubling the cells per gap parameter, which significantly refines the discretization within the narrowest flow passages of the cooling channel.

Table 5 – Number of Nodes, Cells, and Faces for the Three Mesh Configurations

Mesh	N° nodes	N° cells	N° faces
Mesh 2 (coarse)	155 826	38 833	214 371
Mesh 1 (medium)	161 582	41 938	226 930
Mesh 3 (fine)	497 425	129 488	695 123

4.2.2 Methodology and Results of the Grid-Independence Assessment

The independence of the numerical solution from the computational grid was assessed using the single-phase cooling configuration with refrigerant oil, described in detail in Chapter 5, as the reference case, with a fluid inlet temperature of 40 °C. This choice makes it possible to conduct the grid-sensitivity study on a computationally less demanding case than the two-phase configurations, since no phase-change phenomena are involved, whose accurate resolution would, in principle, require a dedicated grid-independence assessment. The analysis is therefore based on the assumption—commonly adopted in CFD modelling practice for components with unchanged geometry—that the grid resolution identified as adequate for the single-phase case is also representative of the two-phase configurations analyzed in Chapter 6, since the latter share the same geometry and the same spatial discretization of the computational domain.

For each of the three mesh configurations, the following quantities of interest, reported in Table 4.5, were compared: the maximum and average temperatures within the cooling channel ($T_{\max duct}$, $T_{av\ duct}$), the fluid outlet temperature (T_{out}), the maximum and average temperatures of the copper winding ($T_{\max winding}$, $T_{av\ winding}$), the pressure drop between the channel inlet and outlet (pressure drop), the maximum and average temperatures of the polymeric partition ($T_{\max plastic}$, $T_{av\ plastic}$), and the maximum and average temperatures of the ferromagnetic structure ($T_{\max lamiantion}$, $T_{av\ lamination}$).

Table 6 – Results of the Grid-Independence Assessment (Oil Inlet Temperature of 40 °C)

Parameter	Mesh 2 (coarse)	Mesh 1 (medium)	Mesh 3 (fine)
$T_{\max duct}$ [°C]	97.11	97.4	104.21
$T_{av\ duct}$ [°C]	43.24	43.26	43.3
T_{out} [°C]	42.1	42.12	42.1
$T_{\max winding}$ [°C]	122.62	123.17	124.35
$T_{av\ winding}$ [°C]	106.77	106.54	106.9
Pressure drop [Pa]	1168.78	1185.1	1193.05
$T_{\max plastic}$ [°C]	99.09	99.47	104.62
$T_{av\ plastic}$ [°C]	60.99	60.92	60.98
$T_{\max lamiantion}$ [°C]	111.96	112.02	112.22
$T_{av\ lamination}$ [°C]	106.85	106.74	106.62

The analysis of the data reported in Table 4.5 allows the following observations to be made.

First, the domain-averaged quantities ($T_{av\ duct}$, T_{out}, $T_{av\ winding}$, $T_{av\ plastic}$, $T_{av\ lamination}$) are essentially independent of grid resolution, with differences among the three configurations consistently limited to a few tenths of a degree Celsius. For example, T_{out} varies between 42.1 °C

and 42.12 °C, corresponding to a difference of 0.05%, while $T_{av\,winding}$ varies between 106.54 °C and 106.9 °C, corresponding to a difference of less than 0.4%. These results confirm the adequacy of all three grid configurations analyzed in correctly predicting the overall energy balance of the computational domain and the integral quantities of interest.

Second, the pressure drop also shows good convergence among the three meshes, with a maximum difference of approximately 2% between Mesh 2 (1168.78 Pa) and Mesh 3 (1193.05 Pa). This variation is consistent with the numerical uncertainty typically associated with this type of simulation and does not indicate any significant systematic dependence on grid resolution.

Third, the local maximum temperatures ($T_{\max duct}$ and $T_{\max plastic}$, in particular) exhibit a different behavior between the two coarser meshes (Mesh 2 and Mesh 1, which are in good agreement with each other, with differences below 1%) and the finest mesh (Mesh 3), which predicts significantly higher values of $T_{\max duct}$ and $T_{\max plastic}$ (104.21 °C and 104.62 °C, respectively, compared with values of approximately 97–99 °C for the other two meshes, corresponding to a difference of approximately 7%). This behavior can be attributed to the inherently local nature of maximum-value quantities, whose estimation is more sensitive to the ability of the computational grid to correctly resolve local gradients in the immediate vicinity of edges and geometric discontinuities of the polymeric partition (see the "Christmas-tree" profile described in Section 4.1.2). In these regions, Mesh 3, owing to its greater local refinement (cells per gap = 2, 7 near-wall layers), is able to capture localized temperature peaks that are not adequately resolved by the coarser meshes. It should also be noted that this sensitivity does not significantly affect the maximum temperature of the copper winding ($T_{\max winding}$), which represents the quantity of greatest engineering interest for the purposes of the present thesis (see Section 2.5.4) and exhibits a difference of no more than 1.4% among the three configurations.

Based on the considerations outlined above, Mesh 1 (medium) was considered to provide the best compromise between result accuracy and computational cost. Indeed, this configuration reproduces the domain-averaged quantities and the maximum winding temperature obtained with the finest mesh (Mesh 3) with almost identical values, while requiring more than 65% fewer cells (41,938 compared with 129,488). This results in significant savings in terms of computational time and resources required to perform the subsequent parametric simulations involving the four two-phase fluids (see Chapter 6), which require a large number of simulations at different mass flow rates for each fluid. Mesh 1 was therefore adopted as the reference grid for all the simulations presented in Chapters 5 and 6 of the present thesis, with the only caveat that the local maximum temperature values in correspondence with the edges of the polymeric partition ($T_{\max duct}$, $T_{\max plastic}$) should be interpreted with slightly greater caution than the domain-averaged quantities, owing to the grid sensitivity highlighted by the comparison with Mesh 3.

4.3 Physical Model for the Single-Phase Case

The simulation of the reference single-phase cooling configuration (described in detail in Chapter 5) was performed by solving, in addition to the mass and momentum conservation equations governing any fluid-dynamic simulation, the energy equation (with the Energy equation enabled in Ansys Fluent), which is required to determine the temperature field within both the fluid domain and the coupled solid domains (windings, polymeric partition, and ferromagnetic structure). This was

achieved according to a Conjugate Heat Transfer (CHT) approach, which enables the simultaneous solution of heat conduction within the solids and convective heat transfer within the fluid, with continuity of temperature and heat flux automatically enforced at the solid-fluid interfaces of the computational domain.

For turbulence closure, the k-ω SST model (Shear Stress Transport, formulated by Menter) was adopted. This choice was motivated by its well-established ability to provide accurate predictions both within the near-wall boundary-layer regions—where the model employs a formulation analogous to the standard k-ω model, which is particularly effective for directly resolving the viscous sublayer without requiring wall functions when the grid is sufficiently refined near the wall—and within the free-stream regions away from the walls, where the model instead employs a k-ε-type formulation, which is less sensitive to the boundary conditions imposed on the inlet turbulence quantities. This hybrid characteristic makes the k-ω SST model particularly suitable for simulating confined flows within complex geometries such as the irregular-cross-section cooling channel investigated in the present study (see Section 4.1.2), where developed boundary-layer regions in the vicinity of the winding and polymeric-partition walls coexist with regions of relatively unrestricted flow within the channel core. Furthermore, the model provides an accurate prediction of the wall convective heat transfer coefficient, a quantity of primary importance for determining the maximum temperatures reached by the windings (see Chapter 5).

4.3.1 Physical Model for the Two-Phase Case

For the simulation of the two-phase cooling configurations (described in detail in Chapter 6), employing the four candidate refrigerants (see Section 3.5.2), a multiphase model was added to the k-ω SST turbulence model and the energy equation, both of which were retained unchanged from the single-phase case. This multiphase model is required to describe the coexistence and interaction of the liquid and vapor phases of the refrigerant within the channel, as well as the mass-transfer mechanism associated with phase change during flow boiling (see Section 3.3.2).

Among the Eulerian multiphase approaches available in Ansys Fluent—the Volume of Fluid (VOF) model, the Mixture model, and the full Eulerian-Eulerian model—the Mixture model was adopted. This approach is simplified compared with the full Eulerian-Eulerian model, as the two phases (liquid and vapor) are treated as a single “mixture” fluid characterized by average properties weighted according to the local volume fraction of each phase. A single momentum equation is solved for the mixture, rather than a separate equation for each phase as in the full Eulerian-Eulerian model, together with a transport equation for the volume fraction of each phase, which describes its spatial distribution within the computational domain. Although this approach introduces a simplification in the description of the relative motion between the phases (slip velocity), which is modelled through an algebraic local-equilibrium relationship rather than through separate momentum equations for each phase, it is computationally more efficient than the full Eulerian-Eulerian model while maintaining adequate accuracy in representing dispersed bubbly and transitional two-phase flow regimes (see Section 3.3.2), which are typical of flow boiling in small-scale channels such as those investigated in the present study. Its adoption is therefore justified as a compromise between physical accuracy and computational cost, the latter being particularly relevant given the large number of parametric simulations required to compare the four candidate fluids at different mass flow rates (see Chapter 6).

The mass-transfer mechanism between the liquid and vapor phases associated with the evaporation—and, where relevant, condensation—of the refrigerant was modelled using the Lee model, a semi-empirical model widely employed in CFD simulations of boiling and condensation. The model is based on the assumption that the interphase mass-transfer rate is proportional to the deviation of the local mixture temperature from the saturation temperature T_{sat} of the fluid at the operating pressure, according to a relationship of the form

$$\dot{m}_{lv} = coeff \cdot \alpha_l \cdot \rho_l \cdot \frac{(T - T_{sat})}{T_{sat}} \qquad (se\ T > T_{sat}\ , evaporation)$$

$$\dot{m}_{vl} = coeff \cdot \alpha_v \cdot \rho_v \cdot \frac{(T_{sat} - T)}{T_{sat}} \qquad (se\ T < T_{sat}\ , condensation)$$

where $\dot{m}_{lv}$ and $\dot{m}_{vl}$ represent the volumetric mass-transfer rates associated with evaporation and condensation, respectively, α_l and α_v are the liquid and vapor volume fractions, ρ_l and ρ_v are the corresponding densities, and coeff is a mass-transfer coefficient (Lee coefficient, or relaxation frequency), with dimensions of inverse time (s^{-1}), which governs the rate at which the system tends to restore thermodynamic equilibrium conditions ($T = T_{sat}$) when a deviation from equilibrium occurs. In the present thesis, consistently with the most commonly adopted practice in the literature for the application of the Lee model to dielectric refrigerants in the absence of specific experimental data for coefficient calibration, a Lee coefficient of 1 s^{-1} was adopted for all two-phase simulations. This value represents a compromise between the need to ensure a sufficiently rapid system response when the saturation temperature is exceeded, thereby avoiding excessive local numerical subcooling or superheating of the fluid relative to T_{sat}, and the need to preserve the numerical stability of the solution. Indeed, it is well established in the literature that excessively high values of the mass-transfer coefficient may introduce stiffness into the system of governing equations, resulting in convergence difficulties.

The introduction of the Mixture model coupled with the Lee mass-transfer mechanism substantially affects the behavior of the two-phase simulation compared with the single-phase case discussed in Section 4.3.1, since it allows the fluid to absorb a significant fraction of the thermal power generated within the windings in the form of latent heat rather than exclusively as sensible heat. Consequently, in regions where an appreciable vapor fraction is present, the fluid temperature tends to remain close to the imposed saturation temperature (T_{sat} = 60 °C for all the fluids analyzed; see Section 4.4.2), consistently with the nearly isothermal nature of the saturated boiling process discussed in Section 3.3.2. Furthermore, the model introduces a dependence of the thermophysical properties of the mixture (density, viscosity, and thermal conductivity) on the local vapor volume fraction, resulting in a nonlinear coupling between the thermal and fluid-dynamic fields. This coupling is responsible for the greater complexity and computational cost of the two-phase simulations compared with the single-phase case, as also reflected in the different numerical solution strategies adopted for the two types of simulation (see Section 4.5).

4.4 Boundary Conditions for the Single-Phase Case

The boundary conditions imposed for the simulation of the reference single-phase cooling configuration are summarized in Table 4.6 and can be described as follows.

- **Inlet:** a mass-flow-inlet boundary condition was imposed, with a total mass flow rate of 15 L/min for all the rotor (equivalent, for the density of the fluid considered, to 0.205 kg/s), equally distributed in the eight pole and after between the two cooling channels present within the domain (see Section 4.1.3), corresponding to a mass flow rate of 0.0128125 kg/s for each individual channel. As part of the parametric study conducted for the single-phase case (see Chapter 5), the fluid inlet temperature was varied among four values—20, 40, 60, and 80 °C—in order to characterize the sensitivity of the cooling performance to the thermal level of the fluid entering the channels.
- **Outlet:** a pressure-outlet boundary condition was imposed, with a relative gauge pressure of 0 Pa, corresponding to an absolute outlet pressure equal to the reference operating pressure of the computational domain (101325 Pa), and with the prevent backflow option enabled. This setting was adopted in order to prevent the occurrence of numerically unstable solutions associated with potential local reverse-flow phenomena at the outlet section, ensuring that the fluid crosses this boundary exclusively in the direction leaving the computational domain.
- **External walls subjected to convection:** the outer surface of the ferromagnetic structure (lamination) and the outer surface of the plastic partition, corresponding to portions of the domain that are not directly exposed to the coolant but that, in the actual machine, are exposed to the surrounding environment (air gap and internal rotor air cavities), were modelled as walls subjected to a convection boundary condition (convective heat transfer coefficient). A convective heat transfer coefficient of h_{conv} = 40 W/(m²·K) and an external fluid reference temperature (free-stream temperature) of $T_{free\ stream}$ = 50 °C were imposed, representing the average thermal conditions expected within these regions of the machine in the absence of their explicit CFD modelling.
- **Adiabatic walls:** the surfaces bounding the computational domain as a consequence of the geometric reduction to one-eighth of the complete rotor (see Section 4.1.1), namely the radial cut surfaces and the surfaces associated with the symmetry planes of the modelled sector, were defined as adiabatic walls (zero heat flux), consistently with the absence of net heat transfer across these surfaces owing to the periodicity and symmetry of the physical problem being represented.
- **Solid-solid and solid-fluid interface walls:** all remaining surfaces of the domain, corresponding to the internal interfaces between the different solid components (winding-partition and partition-ferromagnetic structure interfaces) and between the solid components and the coolant, were defined with a zero heat flux condition in coupled mode. In Ansys Fluent, this option does not impose a predefined heat-flux value but instead determines the interfacial heat transfer through the coupled (conjugate) solution of the energy equations in the adjacent domains, thereby ensuring continuity of temperature and heat flux across the interface according to the Conjugate Heat Transfer approach described in Section 4.3.1.

Table 7 – Boundary Conditions for the Single-Phase Case

Boundary conditions	Value
Total mass flow rate	15 L/min (0.205 kg/s)
Single channel mass flow rate	0.0128125 kg/s
Inlet temperature (T_{in})	20 / 40 / 60 / 80 °C
Outlet gauge pressure	0 Pa (prevent backflow)
h_{conv} external walls (lamination, plastic)	40 W/(m²·K)
$T_{free\ stream}$ external walls	50 °C
Domain Cut Surfaces (Periodicity/Symmetry)	Adiabatiche
Solid–Solid and Solid–Fluid Interfaces	Heat flux 0, coupled

4.4.1 Boundary Conditions for the Multiphase Case

For the two-phase cooling configurations employing the four candidate refrigerants (see Section 3.5.2), the inlet mass flow rate and temperature boundary conditions were redefined with respect to the single-phase case, owing to the different operating conditions required to promote phase change within the cooling channel, and are summarized in Table 4.7.

A total inlet mass flow rate of 0.001136 kg/s (explain in section 6.6.1) was imposed, significantly lower than that used in the single-phase case owing to the ability of the two-phase fluid to remove a substantial fraction of the thermal power in the form of latent heat (see Sections 3.3.2 and 4.3.2). This flow rate was equally distributed between the two cooling channels, corresponding to 0.000568 kg/s (equivalent to 0.568 g/s, or 34.1 g/min) for each individual channel, with a cross-sectional area of 3.77×10^{-5} m². The fluid inlet temperature was set to 50 °C, corresponding to a subcooled liquid condition relative to the target saturation temperature adopted for the comparison.

To ensure full comparability of the results obtained with the four candidate fluids (see Section 3.5.3), an operating pressure corresponding to a saturation temperature of T_{sat} = 60 °C was imposed for each fluid, with the same saturation temperature being adopted for all four fluids analyzed. In this way, the different two-phase configurations are compared at the same phase-change temperature, thereby isolating the effect of the different thermophysical properties of the fluids (latent heat, thermal conductivity, viscosity, and density) as the only discriminating variables among the configurations.

Consistently with the discussion presented in Section 4.3.2, a Lee coefficient of 1 s^{-1} was adopted for all two-phase simulations and applied identically to all four candidate fluids.

The outlet conditions (pressure-outlet, gauge pressure of 0 Pa, with prevent backflow enabled), the external walls subjected to convection (h_{conv} = 40 W/(m²·K), $T_{free\ stream}$ = 50 °C), the adiabatic walls corresponding to the cut surfaces of the computational domain, and the solid-solid/solid-fluid

interfaces defined in coupled mode were kept unchanged with respect to the single-phase case described in Section 4.4.1.

Table 8 – Boundary Conditions for the Multiphase Case

Boundary conditions	Value
Total mass flow rate	0.001136 kg/s
Single channel mass flow rate	0.000568 kg/s (34.1 g/min)
Single channel area	3.77×10^{-5} m²
Inlet temperature	50 °C
Saturation temperature (for all fluids)	60 °C
Lee coefficients	1 s^{-1}

The specific thermophysical properties of each of the four fluids analyzed (Novec 7000, R134a, Opteon SF-33, and FC-72), required to fully characterize their respective liquid and vapor phases within the Mixture model, will be presented in detail in Chapter 6, together with the results obtained for each two-phase configuration.

4.5 Under-Relaxation Factors and Solver Setup

Within the numerical setup of the CFD simulations, particular attention was devoted to the definition of the solver relaxation parameters, as they directly influence the numerical stability and iterative behavior of the solution. The pressure-based coupled solver employed in Ansys Fluent allows two distinct categories of control parameters to be adjusted. On the one hand, the Under-Relaxation Factors in the strict sense limit the update of a generic solution variable φ between two successive iterations according to the relationship $\varphi_{new} = \varphi_{old} + \alpha \cdot \Delta\varphi$, where α is an under-relaxation coefficient ranging between 0 and 1. On the other hand, for the momentum and pressure equations, an alternative formulation based on Pseudo-Transient Explicit Relaxation Factors is available, in which stabilization of the iterative solution process is achieved by introducing a fictitious time step (pseudo-time step), whose magnitude is governed by the Flow Courant Number, rather than by acting directly as a multiplicative factor on the variable increment.

For the single-phase case, pressure-velocity coupling was handled using the pseudo-transient explicit relaxation formulation, with the Flow Courant Number set to 200 and the explicit relaxation factors set to 0.50 for both momentum and pressure. This setting, which is typical of the Coupled solver approach with pseudo-transient stabilization, enables faster convergence compared with the traditional explicit under-relaxation of the momentum and pressure equations, while maintaining adequate control over numerical stability. A Courant Number of 200, although relatively high, was considered compatible with the stability of the problem owing to the steady-state and purely single-

phase nature of the simulation, which does not involve phase-change phenomena or the associated additional nonlinearities.

For the remaining solution variables, the conventional under-relaxation factors reported in the table were adopted.

Table 9 – Under-Relaxation Factors for the Single-Phase Case

Parameter	Value
Flow Courant Number	200
Momentum (Explicit Relaxation Factor)	0.5
Pressure (Explicit Relaxation Factor)	0.5
Density	1
Body forces	0.7
Vaporization Mass	0.3
Slip velocity	0.1
Volume fraction	0.3
Turbulent kinetic energy (k)	0.4
Specific dissipation rate (ω)	0.8
Turbolent viscosity	1
Energy	1

It should be noted that, since the solver control panel is configured for use with the Mixture multiphase model, the under-relaxation factors also include the entries Vaporization Mass, Slip Velocity, and Volume Fraction, which are specific to the two-phase formulation and have no physical significance in the single-phase configuration. In the latter case, the corresponding equations are not solved and, consequently, the associated factors have no influence on the behavior of the solver.

For the two-phase case, the relaxation strategy for the pressure-velocity coupling was instead based on the conventional formulation using direct Under-Relaxation Factors, with values of 0.30 for pressure and 0.40 for momentum, in place of the pseudo-transient explicit relaxation approach employed for the single-phase case. This choice was dictated by the greater complexity and more pronounced nonlinear nature of the two-phase problem, in which the momentum and continuity equations are strongly coupled, through the source terms introduced by the Lee model, with the volume-fraction transport equations and the interphase mass-transfer equations. A more conservative relaxation of pressure and momentum (0.30 and 0.40, respectively, lower than the values implicit in the pseudo-transient formulation adopted for the single-phase case) was therefore considered necessary to reduce the risk of numerical instabilities induced by this coupling, at the expense of a slower convergence rate, which was considered acceptable in view of the greater numerical complexity of the problem.

Table 10 – Under-Relaxation Factors for the Multiphase Case

Parameter	Value
Momentum (Under Relaxation Factor)	0.4
Pressure (Under Relaxation Factor)	0.3
Density	1
Body forces	0.7
Vaporization Mass	0.3
Slip velocity	0.1
Volume fraction	0.3
Turbulent kinetic energy (k)	0.4
Specific dissipation rate (ω)	0.8
Turbolent viscosity	1
Energy	1

The comparison between the two configurations highlights that, with the exception of the relaxation strategy adopted for the pressure-velocity coupling, all the other parameters were kept strictly unchanged between the single-phase and two-phase cases. Density (1.00) and body forces (0.70) were assigned the same values in both configurations, as were the turbulence-model parameters (turbulent kinetic energy equal to 0.40 and specific dissipation rate equal to 0.80, both representing intermediate values selected to limit potential oscillations in the turbulent field without excessively compromising the convergence rate), turbulent viscosity (1.00, with no under-relaxation), and the energy factor (1.00). The latter value, equal to unity, indicates that, in the present configuration, the update of the thermal field was not subjected to any additional explicit under-relaxation. This choice was made possible by the good numerical stability observed for the energy equation in both physical configurations analyzed, including in the presence of the energy-exchange source terms associated with phase change in the two-phase case.

The factors associated with Vaporization Mass (0.30), Slip Velocity (0.10), and Volume Fraction (0.30), which are active and physically meaningful only in the two-phase configuration, govern, respectively, the update of the interphase mass-transfer source term between the liquid and vapor phases calculated by the Lee model, the update of the relative slip velocity between the phases predicted by the algebraic formulation of the Mixture model (Algebraic Slip Mixture), and the update of the volume fraction of each phase within each computational cell. The selection of relatively low values for these parameters, particularly 0.10 for the slip velocity and 0.30 for both vaporization mass and volume fraction, reflects the need to strongly limit the rate of update of quantities that are inherently prone to numerical instability, such as the phase-change rate and the phase distribution in the vicinity of the liquid-vapor interface. An excessively rapid evolution of these quantities between successive iterations could lead to solution divergence or to the onset of non-physical oscillations in the temperature and vapor-fraction distributions within the computational domain.

Overall, the adopted relaxation-parameter configuration reflects a compromise between numerical stability and convergence rate, calibrated according to the physical complexity of each configuration. In the single-phase case, characterized by a less pronounced thermofluid-dynamic coupling, it was

possible to adopt a more aggressive pseudo-transient relaxation strategy for the pressure-velocity coupling. In the two-phase case, owing to the greater nonlinearity introduced by phase change and the resulting coupling among the flow field, thermal field, and phase distribution, a more conservative relaxation strategy was instead preferred. Nevertheless, all the other solver parameters were kept unchanged in order to preserve, as far as possible, consistency in the numerical setup between the two configurations and to limit the influence of purely numerical parameters on the comparison of the physical results obtained.

Finally, it should be emphasized that the under-relaxation factors and pseudo-transient explicit relaxation parameters do not represent physical quantities of the system under investigation, but rather purely numerical parameters intended to control the iterative solution procedure. Their appropriate setting was therefore assessed a posteriori on the basis of the behavior of the numerical solution by verifying residual stability, the absence of non-physical numerical oscillations in the monitored quantities, and the achievement of adequate convergence conditions for the main variables of interest in both the single-phase and two-phase cases.

4.5.1 Solver Setup for the Single-Phase Case

For both types of simulation (single-phase and two-phase), the pressure-based solver available in Ansys Fluent was adopted, with the velocity formulation set to Absolute, which represents the standard choice for the simulation of low-velocity, incompressible or weakly compressible flows such as those investigated in the present study, rather than the density-based solver, which is typically employed for high-speed flows characterized by dominant compressibility effects.

The single-phase simulation was performed under steady-state conditions, consistently with the objective of characterizing the system under steady operating conditions, without the need to resolve the temporal evolution of the flow and temperature fields.

For pressure-velocity coupling, the Coupled scheme was adopted, which simultaneously solves the continuity and momentum equations as a single coupled algebraic system, rather than using segregated schemes such as SIMPLE or SIMPLEC. This approach generally provides a higher convergence rate, particularly for high-quality grids and problems characterized by strong coupling between the pressure and velocity fields, such as the one considered in the present study. The Flux Type was set to Rhie-Chow momentum-based, with automatic selection.

For spatial discretization, the Least Squares Cell Based scheme was adopted for gradients, the Second Order scheme for pressure, and the Second Order Upwind scheme for the transported scalars (momentum, turbulent kinetic energy, specific dissipation rate of the k-ω SST model, and energy). This configuration ensures second-order spatial accuracy for all the equations solved, at the expense of a higher computational cost compared with first-order schemes, but with reduced numerical diffusion, which is particularly important for accurately resolving the thermal and velocity gradients present within the computational domain. The Warped-Face Gradient Correction option was also enabled to improve the accuracy of gradient calculations on grids containing cells with non-planar faces (warped faces), which may occur in the polyhedral cells generated in the vicinity of the complex geometry of the polymeric partition (see Section 4.1.2).

4.5.2 Solver Setup for the Two-Phase Case

The greater physical and numerical complexity introduced by the Mixture multiphase model coupled with the Lee mass-transfer mechanism (see Section 4.3.2) required the adoption of a different solution strategy compared with the single-phase case, consisting of two successive stages.

Stage 1 — Steady-State Initialization: the two-phase simulation was initially performed under steady-state conditions for 6000 iterations in order to obtain sufficiently developed and numerically stable flow, temperature, and vapor-fraction fields to be used as the initial condition for the subsequent transient stage. This two-stage approach, widely adopted in CFD practice for the simulation of complex multiphase phenomena, helps prevent the onset of numerical instabilities during the initial stages of the transient simulation, which would otherwise be initialized from an undeveloped flow field. This reduces the risk of solution divergence while simultaneously limiting the overall number of time steps required to reach a statistically or quasi-steady operating condition, since the transient simulation starts from a state already close to the solution of interest rather than from arbitrary uniform initial conditions.

Stage 2 — Transient Simulation with Gravity Enabled: starting from the steady-state solution obtained in Stage 1, the simulation was continued in transient mode for a total physical time of 10 s, with Gravity enabled, unlike in the single-phase case. The inclusion of gravity in the two-phase simulation is necessary because the density difference between the liquid and vapor phases of the refrigerant—typically one or two orders of magnitude, with the vapor being significantly less dense than the liquid—makes the gravitational field a non-negligible factor in determining the spatial distribution of the two phases within the channel. This includes buoyancy phenomena, stratification, and preferential migration of the vapor phase towards the upper regions of the channel, effects that are absent, or in any case not relevant to the solution, in the single-phase case characterized by an essentially uniform density. The transition to a transient formulation also makes it possible to correctly capture the inherently unsteady nature of the flow-boiling process, characterized by the cyclic formation, growth, and detachment of vapor bubbles (see Section 3.3.2), a phenomenon that cannot be adequately represented using a purely steady-state formulation.

Owing to the greater numerical stiffness introduced by the Lee-model mass-transfer mechanism and the transient two-phase formulation, the solution schemes adopted for the single-phase case (see Section 4.5.1) were appropriately modified for the two-phase simulations, as follows:

- The pressure-velocity coupling scheme was changed from Coupled to PISO (Pressure-Implicit with Splitting of Operators), a segregated scheme particularly suitable for transient simulations and multiphase problems owing to its greater numerical robustness and its ability to effectively handle the multiple pressure corrections required by the presence of multiple phases and the strong coupling among the pressure, velocity, and vapor volume-fraction fields.
- The spatial discretization scheme for pressure was changed from Second Order to PRESTO! (PREssure STaggering Option), a scheme recommended in Ansys Fluent for the simulation of multiphase flows characterized by strong pressure gradients, high streamline curvature, or large density variations, such as those associated with phase change within the cooling channel.

For the spatial discretization of the convective terms, upwind schemes of different orders were adopted depending on the variable considered. Momentum, turbulent kinetic energy (k), specific dissipation rate (ω), and energy were discretized using the Second Order Upwind scheme in order to

limit numerical diffusion and ensure a consistent level of accuracy among the solved fields, particularly for the thermal field. The volume fraction, which is fundamental to the formulation of the Mixture model, was instead discretized using the First Order Upwind scheme. This choice, commonly adopted in multiphase simulations, prioritizes numerical robustness, since higher-order schemes for this variable tend to generate non-physical oscillations (undershoots/overshoots) in the vicinity of the liquid-vapor interface, with the potential occurrence of vapor-fraction values outside the physically admissible range [0, 1].

For temporal discretization, a First Order Implicit formulation (implicit Euler scheme) was adopted. This scheme is unconditionally stable from a numerical standpoint and was considered adequate, despite its first-order temporal accuracy, for the characteristic time scales of the thermofluid-dynamic quantities of interest in the present study.

Although this solution strategy involves a higher overall computational cost compared with the single-phase case—owing both to the larger number of equations being solved (vapor volume fraction and interphase mass transfer) and to the need to perform a transient simulation following the steady-state initialization—it was applied identically to all the two-phase configurations analyzed in Chapter 6, including the four candidate fluids and the different mass flow rates investigated in the parametric study, thereby ensuring full methodological comparability of the results obtained.

Chapter 5

Case Study: Single-Phase Cooling

This chapter presents and discusses the results of the simulation campaign conducted on the reference single-phase cooling configuration, using AmpCool AC-110 dielectric oil as the coolant. After reviewing the thermophysical properties of the fluid employed and the operating conditions of the case study, the results obtained by varying the fluid inlet temperature are presented and analyzed, with the aim of characterizing the thermal and fluid-dynamic performance of the rotor cooling channel under single-phase conditions and identifying its inherent limitations, thereby providing the basis for comparison with the two-phase configurations presented in Chapter 6.

5.1 Coolant: AmpCool AC-110 Dielectric Oil

The fluid employed in the reference single-phase cooling configuration is AmpCool AC-110, a dielectric oil belonging to the AC-100 series developed by Engineered Fluids, specifically designed for immersion cooling applications involving batteries, battery chargers, inverters, and electronic and electromechanical components in general. According to the classification discussed in Section 3.5.2, it is a dielectric fluid characterized by the absence of phase change (boiling point: none) and composed of synthetic hydrocarbons. It therefore represents a natural reference fluid for assessing, in the subsequent chapters, the benefits associated with the transition to two-phase cooling.

The main thermophysical and dielectric properties of the fluid, reported in the manufacturer's technical datasheet [27] and summarized in Table 5.1, justify its use in the present study as the single-phase reference case:

- **Rheological behavior and physical state:** the fluid is specified by the manufacturer as incompressible, isotropic, and Newtonian, with a clear appearance and no boiling point within the temperature range of interest for electrical applications. This therefore excludes any phase-change phenomenon over the operating range investigated in the present chapter (20–80 °C), confirming the purely single-phase nature of the configuration analyzed.
- **Dielectric properties:** the fluid exhibits a dielectric strength greater than 60 kV (measured according to the ASTM D1816 method using electrodes with a 2 mm gap) and a volume electrical resistivity greater than 1×10^{14} Ω·cm. These values are fully consistent with the electrical compatibility requirements discussed in Section 3.2.4 for the direct cooling of energized components, making the fluid suitable for direct contact with the excitation windings through the polymeric partition described in Section 4.1.2.

- **Kinematic viscosity:** the fluid exhibits a pronounced dependence of kinematic viscosity on temperature, typical of mineral and synthetic oils, with values of 41.10 cSt at 0 °C, 8.11 cSt at 40 °C, and 2.22 cSt at 100 °C. The viscosity therefore decreases significantly with increasing temperature, an aspect of primary importance for interpreting the pressure-drop results presented in Section 5.2.
- **Thermal conductivity and specific heat capacity:** the thermal conductivity, approximately 0.136–0.138 W/(m·K) over the temperature range considered, and the specific heat capacity, ranging from approximately 2.06 to 2.44 kJ/(kg·K), are representative of typical values for dielectric oils used in electrical cooling applications. Although their specific heat capacity is significantly lower than that of water-based fluids, these properties are fully compatible with the use of AmpCool AC-110 as a sensible-heat coolant in a single-phase cooling configuration.
- **Safety and sustainability:** the fluid has a flash point of 193 °C, a pour point of −57 °C, a Global Warming Potential (GWP) of zero, and is characterized by biodegradability and non-toxicity. These aspects are relevant to compliance with the environmental regulations discussed in Section 3.5.1, although they do not directly affect the thermofluid-dynamic performance investigated in the present analysis.

Table 11 – Main Thermophysical and Dielectric Properties of AmpCool AC-110 Oil

Property	Value
Appearance / Behavior	Clear; incompressible, isotropic, Newtonian
Density @ 15.6 °C	0.82 g/cm³
Kinematic viscosity @ 0 / 40 / 100 °C	41.10 / 8.11 / 2.22 cSt
Thermal conductivity @ 0 / 40 / 100 °C	0.1382 / 0.1359 / 0.1325 W/(m·K)
Specific heat @ 0 / 40 / 100 °C	2.0608 / 2.2121 / 2.4390 kJ/(kg·K)
Dielectric Strength	> 60 kV
Volume Resistivity	$> 1\times10^{14}$ Ω·cm
Flash Point	193 °C
Pour Point	−57 °C
Boiling Point	Absent
GWP	0

5.1.1 Operating Conditions of the Case Study

The simulations presented in this chapter were performed using the boundary conditions defined for the single-phase case in Section 4.4.1: a total mass flow rate of 15 L/min (0.205 kg/s), equally distributed between the two cooling channels of the computational domain (0.0128125 kg/s per channel); zero gauge pressure at the outlet; external surfaces of the ferromagnetic structure and polymeric partition subjected to convection (h_{conv} = 40 W/(m²·K), $T_{free\ stream}$ = 50 °C); and volumetric heat generation within the windings corresponding to 72.7 W for each modelled half-winding (see Section 4.1.5), resulting in a total thermal power input to the computational domain (including both channels) of 145.42 W.

In order to characterize the sensitivity of the cooling-system performance to the thermal level of the fluid at the inlet, a parametric study was conducted by varying the oil inlet temperature over four values: T_{in} = 20, 40, 60, and 80 °C. These values represent a broad range of plausible operating conditions for the machine cooling circuit, ranging from start-up conditions with a relatively cold fluid to steady operating conditions in which the fluid has already been partially heated by the overall vehicle cooling circuit.

5.2 Simulation Results

The results of the four simulations performed are presented below, showing the trends of the various quantities investigated up to the convergence achieved in Ansys Fluent.

5.2.1 Simulation with $T_{in} = 20°C$

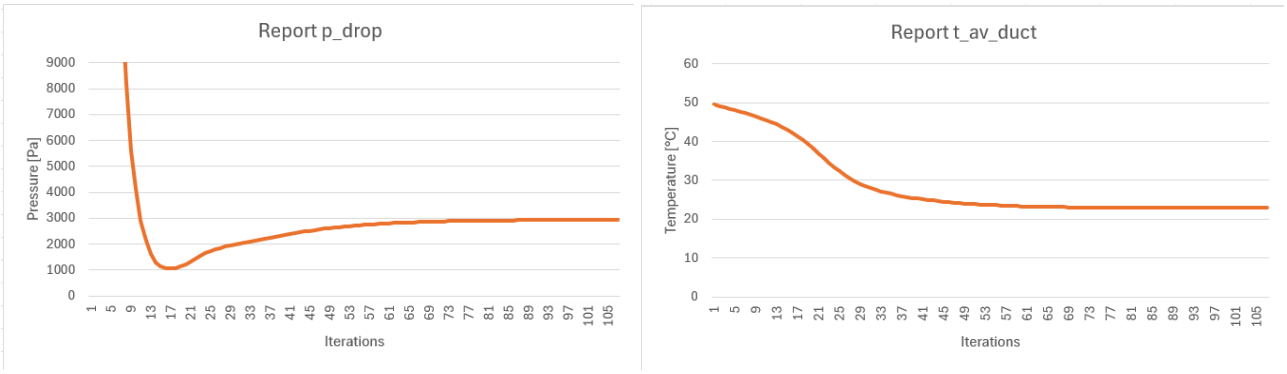
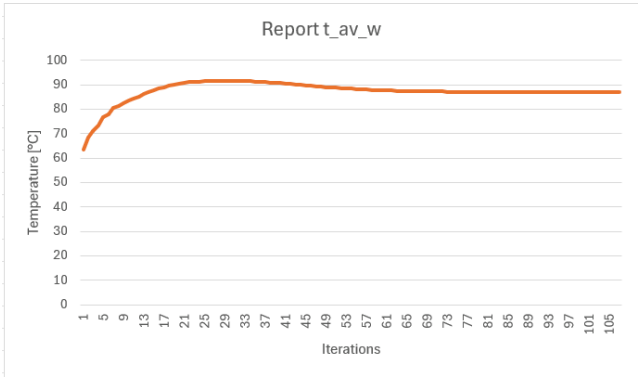


Figure 19 – Pressure Drop Report, Inlet Temperature 20 °C
Figure 20 – Average Cooling-Channel Temperature Report, Inlet Temperature 20 °C



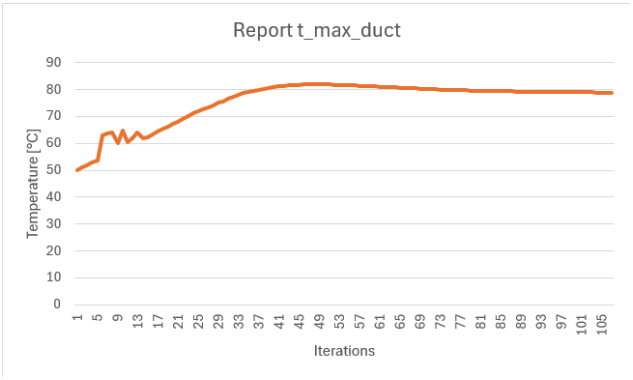
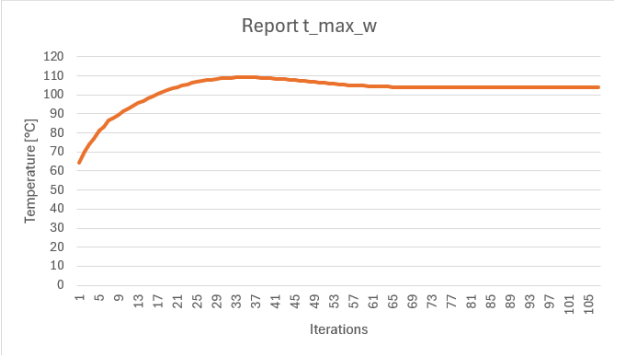


Figure 21 – Average Winding Temperature Report, Inlet Temperature 20 °C
Figure 22 – Maximum Cooling-Channel Temperature Report, Inlet Temperature 20 °C



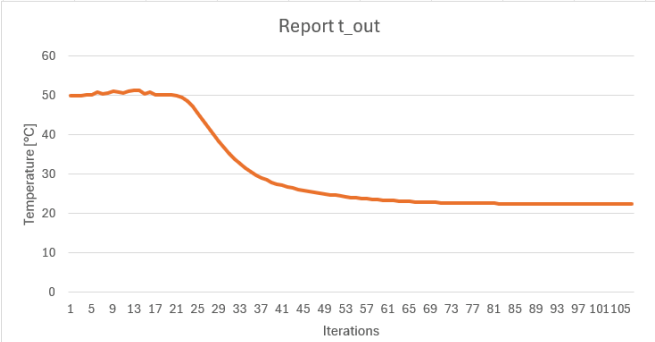
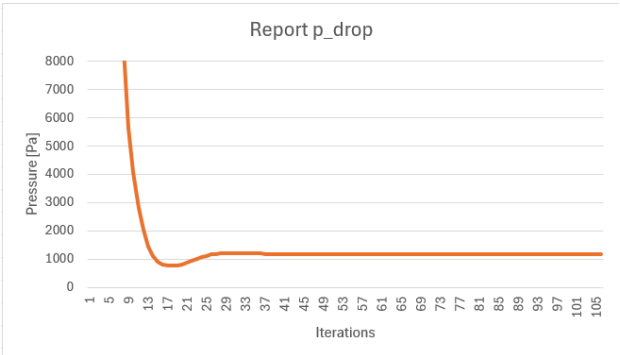


Figure 23 – Maximum Winding Temperature Report, Inlet Temperature 20 °C
Figure 24 – Cooling-Channel Outlet Temperature Report, Inlet Temperature 20 °C

5.2.2 Simulation with $T_{in} = 40°C$



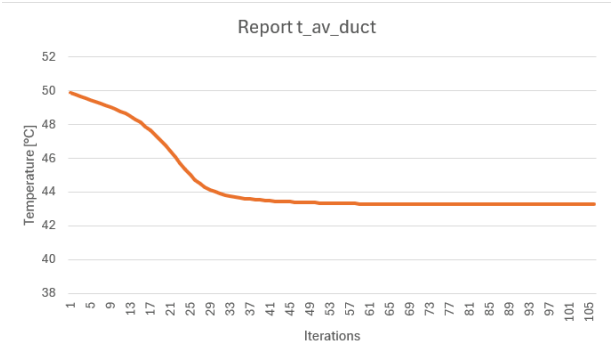
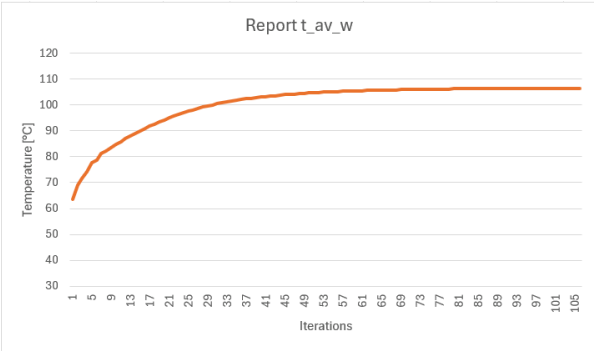


Figure 25 – Pressure Drop Report, Inlet Temperature 40 °C
Figure 26 – Average Cooling-Channel Temperature Report, Inlet Temperature 40 °C



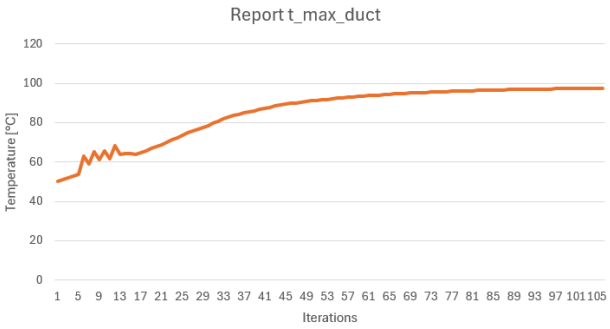
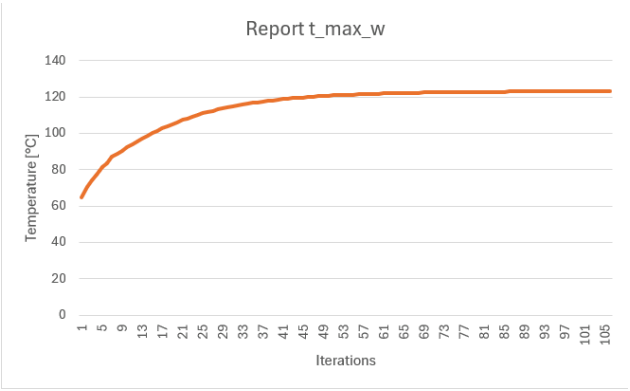


Figure 27 – Average Winding Temperature Report, Inlet Temperature 40 °C
Figure 28 – Maximum Cooling-Channel Temperature Report, Inlet Temperature 40 °C



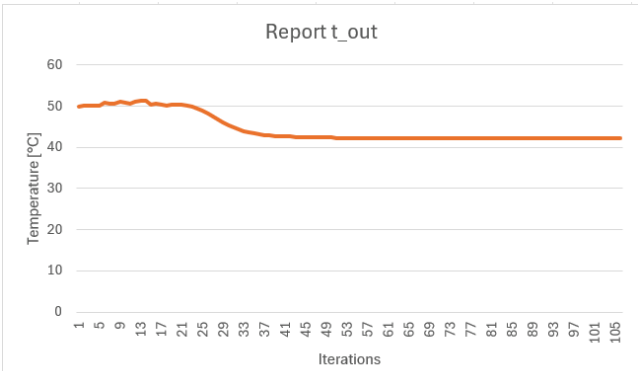
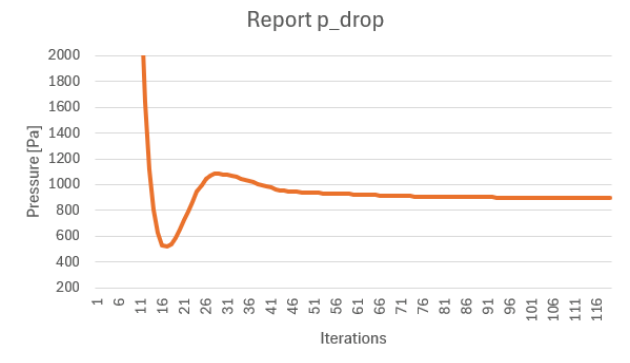


Figure 29 – Maximum Winding Temperature Report, Inlet Temperature 40 °C
Figure 30 – Cooling-Channel Outlet Temperature Report, Inlet Temperature 40 °C

5.2.3 Simulation with $T_{in} = 60°C$



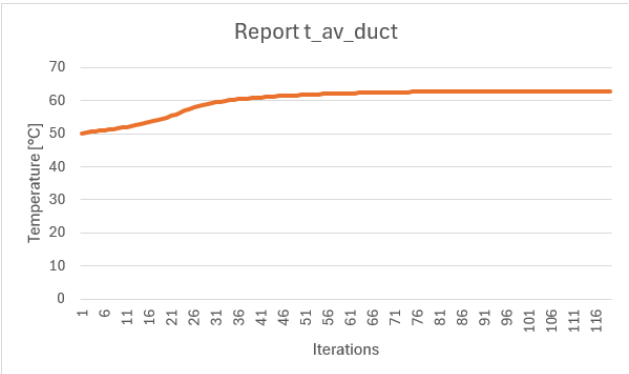
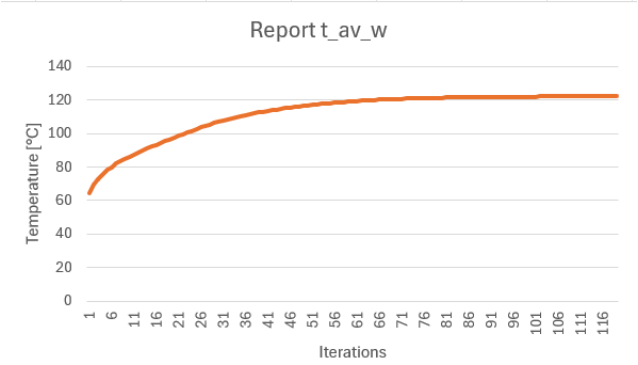


Figure 31 – Pressure Drop Report, Inlet Temperature 60 °C
Figure 32 – Average Cooling-Channel Temperature Report, Inlet Temperature 60 °C



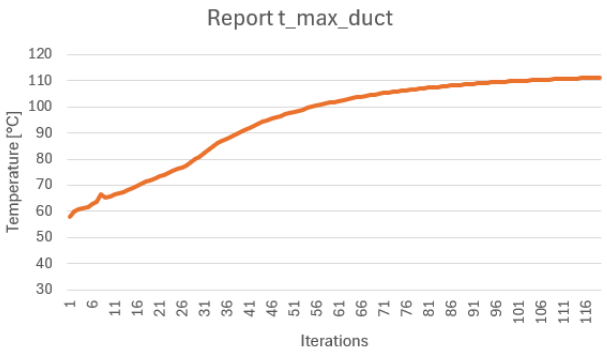
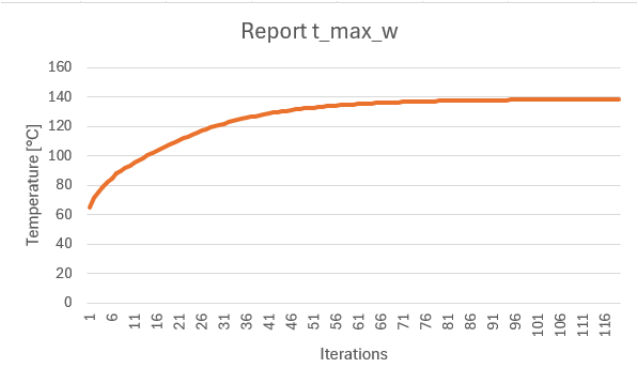


Figure 33 – Average Winding Temperature Report, Inlet Temperature 60 °C
Figure 34 – Maximum Cooling-Channel Temperature Report, Inlet Temperature 60 °C



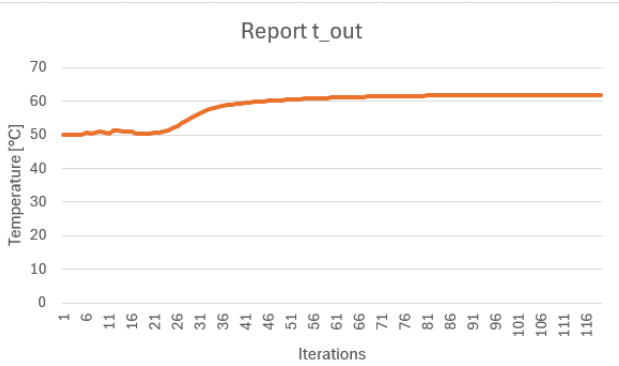
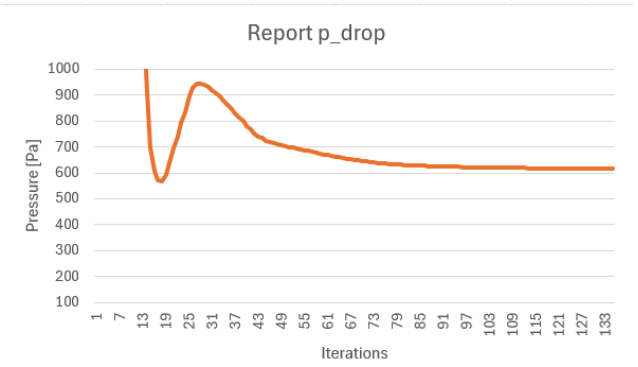


Figure 35 – Maximum Winding Temperature Report, Inlet Temperature 60 °C
Figure 36 – Cooling-Channel Outlet Temperature Report, Inlet Temperature 60 °C

5.2.4 Simulation with $T_{in} = 80°C$



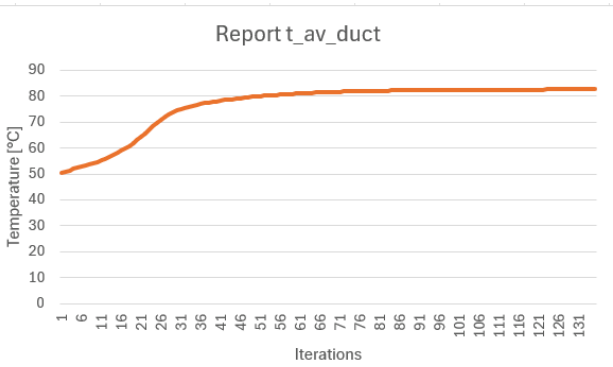
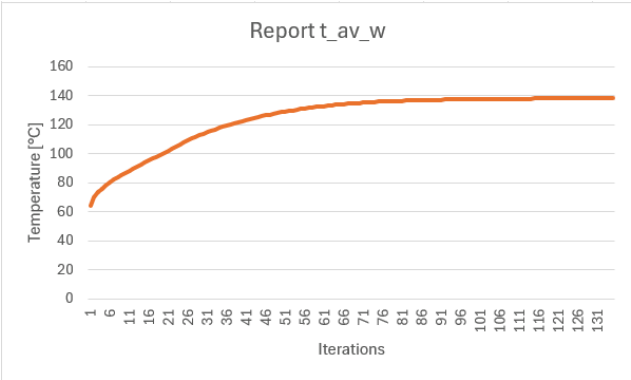


Figure 37 – Pressure Drop Report, Inlet Temperature 80 °C
Figure 38 – Average Cooling-Channel Temperature Report, Inlet Temperature 80 °C



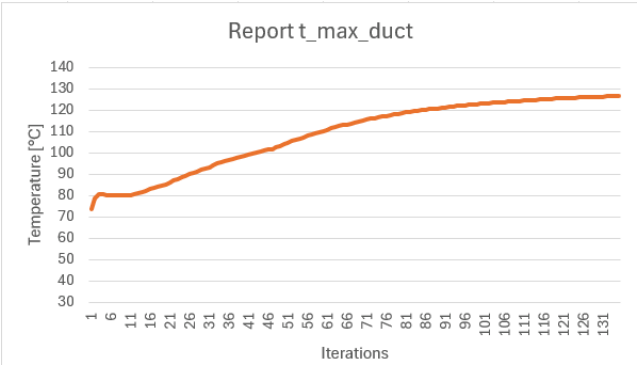
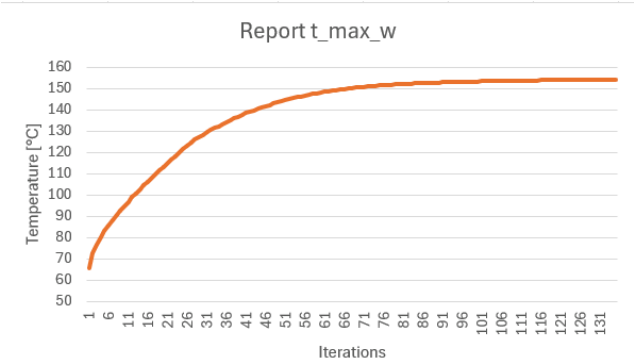


Figure 39 – Average Winding Temperature Report, Inlet Temperature 80 °C
Figure 40 – Maximum Cooling-Channel Temperature Report, Inlet Temperature 80 °C



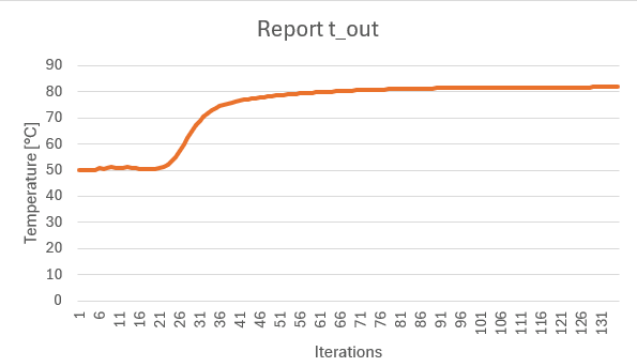


Figure 41 – Maximum Winding Temperature Report, Inlet Temperature 80 °C
Figure 42 – Cooling-Channel Outlet Temperature Report, Inlet Temperature 80 °C

5.3 Presentation of the Results

The results of the four simulations performed, in terms of the maximum and average temperatures of the different components within the computational domain, the fluid outlet temperature, and the pressure drop between the inlet and outlet of the cooling channel, are summarized in Table 5.2.

Table 12 – Results of the Single-Phase Simulations at Different Inlet Temperatures

Parameter	**T_{in} = 20 °C**	**T_{in} = 40 °C**	**T_{in} = 60 °C**	**T_{in} = 80 °C**
$T_{max\ duct}$ [°C]	78.91	97.47	111.12	126.62
$T_{av\ duct}$ [°C]	22.87	43.26	62.88	82.58
T_{out} [°C]	22.37	42.12	61.92	81.75
$T_{max\ winding}$ [°C]	104.08	123.18	138.42	154.18
$T_{av\ winding}$ [°C]	87.07	106.53	122.2	138.23
Pressure drop [Pa]	2936.91	1185.06	898.26	616.89
$T_{max\ plastic}$ [°C]	81.63	99.49	113.15	128.64
$T_{av\ plastic}$ [°C]	42.78	60.92	77.57	94.4
$T_{max\ lamination}$ [°C]	93.84	112.01	126.31	141.21
$T_{av\ lamination}$ [°C]	88.83	106.73	120.87	135.45

For each of the four configurations, the overall energy balance of the computational domain was also verified, as reported in Table 5.3, as an indicator of the reliability and convergence of the numerical solution obtained. Indeed, the imposed heat source (user source, equal to 145.42 W in all cases and corresponding to the volumetric heat generation applied to the two half-windings of the computational domain; see Section 4.1.5) must be balanced, within a limited residual (net result), by the sum of the heat fluxes exchanged by the fluid through the inlet and outlet sections of each channel (Q_{inlet}, Q_{outlet}) and the heat fluxes dissipated by convection from the external surfaces of the ferromagnetic structure and polymeric partition ($Q_{conv\ lam}$, $Q_{conv\ plastic}$).

Table 13 – Overall Energy Balance of the Computational Domain at Different Inlet Temperatures

Parameter	T_{in} = 20 °C	T_{in} = 40 °C	T_{in} = 60 °C	T_{in} = 80 °C
User source [W]	145.42	145.42	145.42	145.42
$Q_{inlet\,1}$ **[W]**	−137.47	419.68	996.22	1592.148
$Q_{inlet\,2}$ **[W]**	−137.47	419.68	996.22	1592.148
$Q_{outlet\,1}$ **[W]**	71.54	−480.57	−1053.28	−1645.82
$Q_{outlet\,2}$ **[W]**	71.88	−480.27	−1053.16	−1645.73
$Q_{conv\,lam}$ **[W]**	−14.098	−21.34	−27.06	−32.96
$Q_{conv\,plastic}$ **[W]**	−1.03	−2	−2.8	−3.62
Net result [W]	−1.228	0.599	1.559	1.575

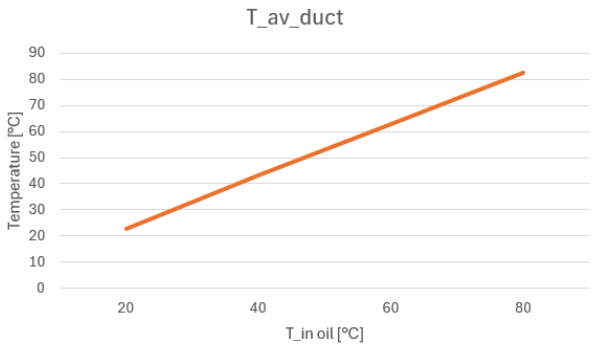
The residual of the energy balance (net result) remains, for all four configurations analyzed, within a maximum absolute value of approximately 1.6 W, corresponding to less than 1.1% of the total imposed thermal power (145.42 W). This value was considered fully acceptable for the numerical validation of the simulations and confirms the adequate convergence of the solution under all the operating conditions investigated.

5.3.1 Analysis of Temperature Fields and Hot Spots

The analysis of the results reported in Table 5.2 first highlights that the maximum temperature of the copper winding ($T_{\max winding}$), which represents a quantity of primary engineering interest for assessing the thermal margin with respect to the adopted insulation class (see Section 2.5.4), increases almost linearly with the fluid inlet temperature, rising from 104.08 °C for T_{in} = 20 °C to 154.18 °C for T_{in} = 80 °C, corresponding to an overall increase of approximately 50 °C for a 60 °C increase in inlet temperature. This behavior is consistent with the physics of single-phase cooling discussed in Section 3.3.1, according to which, for a given thermal power generation and convective heat transfer coefficient (which remains substantially constant when only the inlet temperature is varied, since the mass flow rate is kept unchanged), the wall temperature remains approximately shifted relative to the fluid temperature according to Newton's law of cooling (q" = h·($T_w - T_f$)), thereby explaining the essentially linear trend observed.

Furthermore, the difference between the maximum and average winding temperatures (ΔT = $T_{\max winding} - T_{\text{av}\,winding}$) remains relatively stable as T_{in} varies, ranging between 15.9 °C (T_{in} = 60 °C) and 17.5 °C (T_{in} = 60 °C, minimum value), up to approximately 17.0 °C (T_{in} = 20 °C) and 16.0 °C (T_{in} = 80 °C). This indicates that the relative temperature distribution within the winding, and consequently the location and relative intensity of the hot spot with respect to the average temperature, is not significantly affected by the absolute thermal level of the coolant, but is instead primarily determined by the channel geometry and the distribution of heat generation, both of which remain unchanged across the four configurations analyzed.

Another noteworthy aspect concerns the comparison between the maximum temperature of the polymeric partition ($T_{\max plastic}$) and the maximum winding temperature. In all the configurations analyzed, $T_{\max plastic}$ is approximately 22–26 °C lower than $T_{\max winding}$, a difference attributable to the conductive thermal resistance introduced by the partition itself (see Section 4.1.4), which lies along the heat dissipation path between the heat source (the winding) and the coolant flowing through the channel. This observation quantitatively confirms the importance of the thermal resistance of the insulating partition in the overall thermal balance of the system. This aspect is also relevant to the comparison with the two-phase configurations presented in Chapter 6, in which the same conductive thermal resistance is present, while the convective heat transfer coefficient on the fluid side is expected to be significantly higher.



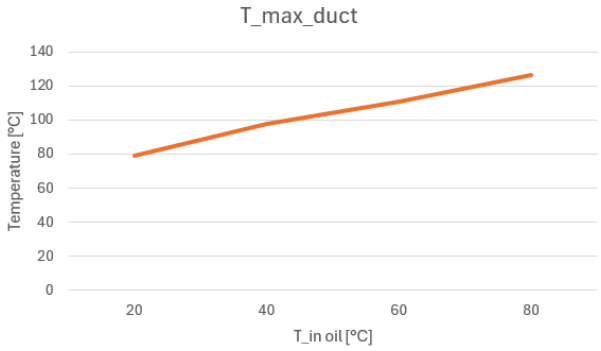
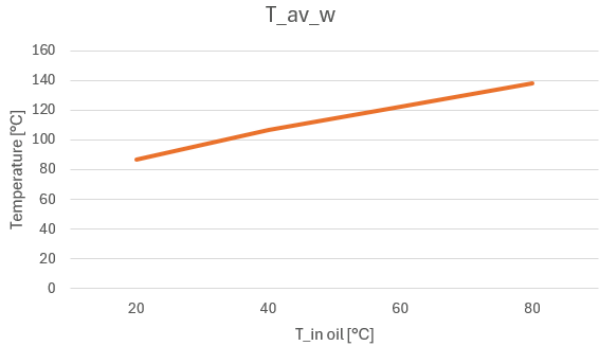


Figure 43 – Average Cooling-Channel Temperature as a Function of Refrigerant Inlet Temperature
Figure 44 – Maximum Cooling-Channel Temperature as a Function of Refrigerant Inlet Temperature



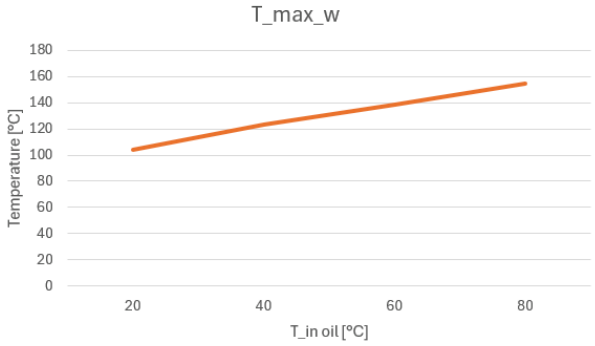
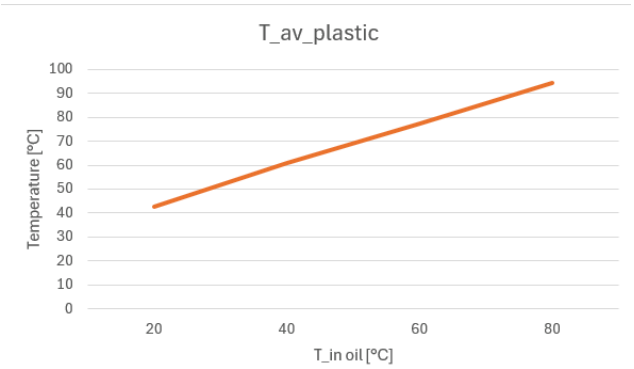


Figure 45 – Average Winding Temperature as a Function of Refrigerant Inlet Temperature
Figure 46 – Maximum Winding Temperature as a Function of Refrigerant Inlet Temperature



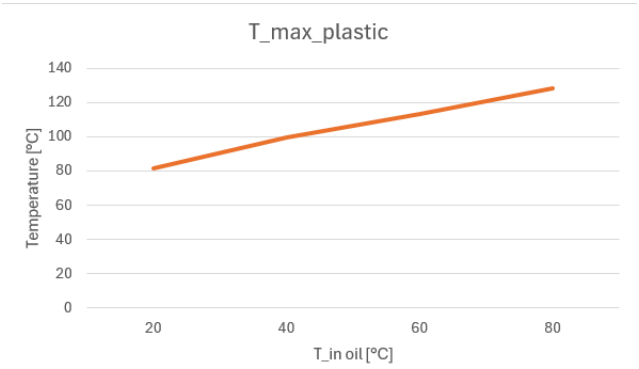
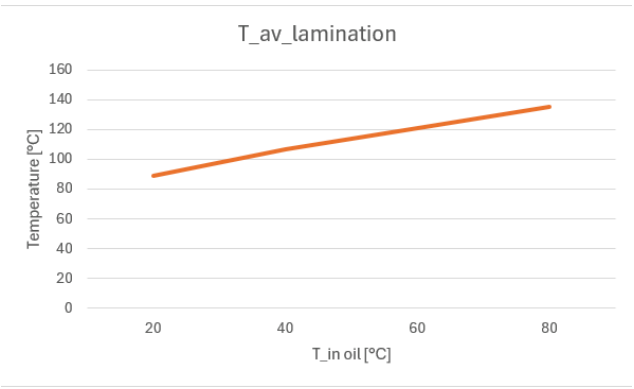


Figure 47 – Average Plastic-Partition Temperature as a Function of Refrigerant Inlet Temperature
Figure 48 – Maximum Plastic-Partition Temperature as a Function of Refrigerant Inlet Temperature



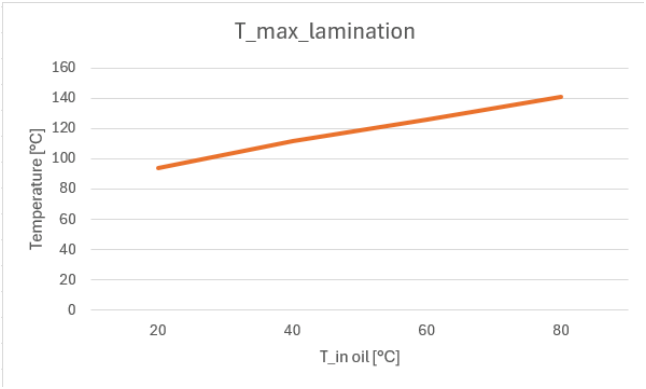


Figure 49 – Average Rotor Ferromagnetic-Structure Temperature as a Function of Refrigerant Inlet Temperature
Figure 50 – Maximum Rotor Ferromagnetic-Structure Temperature as a Function of Refrigerant Inlet Temperature

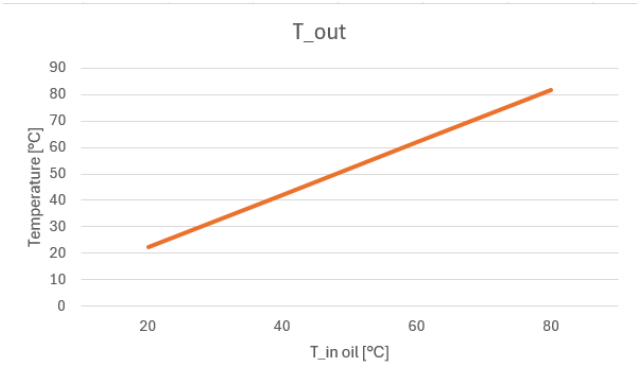


Figure 51 – Refrigerant Outlet Temperature as a Function of Refrigerant Inlet Temperature

5.3.2 Pressure Drop Analysis

The most significant fluid-dynamic result emerging from the analysis of the four cases concerns the trend of the pressure drop between the inlet and outlet of the cooling channel, which decreases markedly as the fluid inlet temperature increases, falling from 2936.91 Pa at T_{in} = 20 °C to only 616.89 Pa at T_{in} = 80 °C, corresponding to a reduction of almost 79% over the entire investigated range, while the mass flow rate was kept constant in all configurations. This behavior, which may appear counterintuitive if the temperature dependence of the fluid properties is neglected, can be fully explained by the strong dependence of the kinematic viscosity of AC-110 oil on temperature, as discussed in Section 5.1.1 (41.10 cSt at 0 °C compared with 2.22 cSt at 100 °C, corresponding to a reduction of almost one order of magnitude). Since distributed pressure losses in a duct are, particularly under laminar or transitional flow conditions, directly related to fluid viscosity (according to the Hagen–Poiseuille relation for laminar flow, or through the Darcy friction factor, which is itself a function of the Reynolds number and therefore of viscosity, for turbulent flow), a fluid with significantly lower viscosity at higher inlet temperatures results in a substantially lower pressure drop for the same imposed volumetric flow rate.

This result has significant design implications. On the one hand, the use of a higher fluid inlet temperature is advantageous in terms of the pumping power required by the cooling circuit, owing to the lower pressure drop. On the other hand, as observed in Section 5.2.2, an increase in inlet temperature leads to a reduction in the thermal margin of the winding, highlighting a trade-off between fluid-dynamic and thermal performance that must be appropriately considered during the design of the complete machine cooling circuit, depending on the expected operating conditions and the coolant supply temperature provided by the overall vehicle cooling system.

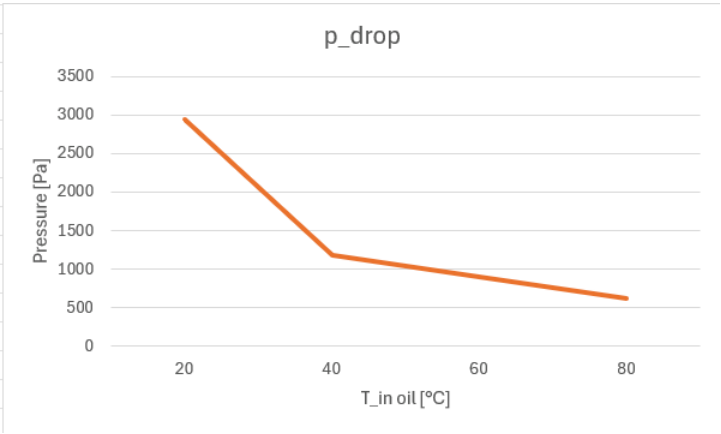


Figure 52 – Pressure Drop as a Function of Refrigerant Inlet Temperature

5.4 Critical Analysis of the Limitations of Single-Phase Cooling

The results presented in the previous sections make it possible to formulate several critical considerations regarding the inherent limitations of single-phase cooling applied to the rotor cooling channel investigated in the present study, providing the basis for assessing the benefits achievable through the transition to two-phase cooling, as presented in Chapter 6.

In the configuration with T_{in} = 80 °C, the maximum winding temperature reaches 154.18 °C, a value that, although not necessarily exceeding the thermal-class limit of the insulation employed (typically 180 °C for Class H; see Section 2.2.2), significantly reduces the available safety margin for potential overload transients or future increases in machine power density. This highlights that single-phase cooling, although satisfactory under the nominal operating conditions analyzed, provides limited margins for further performance improvement.

As discussed in Section 3.3.1, single-phase cooling removes the generated heat exclusively in the form of sensible heat, resulting in an approximately linear increase in fluid temperature along the channel and, consequently, in the winding wall temperature. The absence of the contribution provided by the latent heat of vaporization, which is instead available in the two-phase configurations analyzed in Chapter 6, represents the most significant physical limitation of the single-phase configuration and motivates the interest of the present thesis in two-phase cooling technologies as a potential evolution of the rotor cooling system.

The pressure drop recorded for the configuration with T_{in} = 20 °C (2936.91 Pa), although not necessarily prohibitive for the machine cooling circuit, is significantly higher than that obtained for the other configurations analyzed, owing to the high viscosity of the oil at low temperatures (see Section 5.2.3). This result suggests that, if the machine cooling circuit is required to operate under cold-start conditions or at low ambient temperatures, the required pumping power may represent a non-negligible design constraint. This is an aspect that the two-phase refrigerants considered in Chapter 6—characterized by significantly lower viscosities than dielectric oils—could potentially help mitigate.

The approximately linear relationship between the maximum winding temperature and the fluid inlet temperature, observed in Section 5.2.2, implies that the performance of the single-phase cooling system is strongly dependent on the thermal conditions of the cooling circuit upstream of the rotor channel (for example, on the coolant supply temperature provided by the overall vehicle cooling system). Consequently, the channel itself has a limited ability to “absorb” variations in inlet temperature without a corresponding variation in winding temperature. As will be shown in Chapter 6, this behavior is partially mitigated in two-phase configurations owing to the nearly isothermal nature of the saturated boiling process (see Section 3.3.2), which tends to “anchor” the fluid temperature, and consequently the temperature of the cooled walls, to the saturation temperature imposed by the operating pressure of the circuit, within certain limits independently of the inlet temperature of the fluid itself.

Overall, the results presented in this chapter confirm that single-phase cooling using AC-110 dielectric oil represents a technically valid and reliable solution for the nominal operating conditions of the EESM investigated in the present study. Nevertheless, significant margins for improvement remain in terms of the maximum temperature reached by the windings and the sensitivity of the cooling performance to the fluid inlet conditions. These limitations are addressed by the two-phase cooling approach investigated in Chapter 6, which aims to overcome them by exploiting the latent heat of vaporization of the four candidate refrigerants.

Chapter 6

Case Study: Two-Phase Cooling — Comparison of Four Fluids

This chapter presents and discusses the results of the simulation campaign conducted on the two-phase cooling configurations, using four candidate dielectric refrigerants: Novec 7000, R134a, Opteon SF-33, and FC-72. After reviewing the thermophysical properties of the fluids implemented in the numerical model and the methodology adopted for their retrieval, the following sections present the comparison methodology, the results of the parametric mass-flow-rate study conducted for each fluid, and the overall comparative assessment of the four two-phase configurations against the reference single-phase configuration presented in Chapter 5.

6.1 Methodology for the Retrieval of Thermophysical Properties

The correct implementation of the Mixture model coupled with the Lee mass-transfer mechanism (see Section 4.3.2) requires, for each of the four candidate fluids, the definition of a complete set of thermophysical properties for both the liquid and vapor phases, evaluated under saturation conditions at the target comparison temperature (T_{sat} = 60 °C; see Section 4.4.2), together with the corresponding operating pressure and the latent heat of vaporization at the same temperature.

Since not all of these fluids are natively available in the material database included with Ansys Fluent with complete and accurate properties at the conditions of interest, the thermophysical properties employed in the present work were retrieved and verified using multiple independent sources: saturation charts and curves (pressure-temperature diagrams and transport-property diagrams as a function of temperature) published in the fluids' manufacturers' technical datasheets; commercial catalogs provided by the respective suppliers (3M for Novec 7000 and FC-72, Chemours for Opteon SF-33, and HVAC&R refrigerant manufacturers for R134a); reference tables and handbooks of thermophysical properties for refrigerants; as well as thermodynamic databases commonly used in engineering practice to determine saturated properties not directly available in commercial technical datasheets. For each property, cross-validation among the different available sources was performed whenever possible, giving preference to values obtained from saturation charts and curves providing data directly at the target saturation temperature (60 °C), rather than to values interpolated or extrapolated from temperature ranges that did not directly include the point of interest.

6.2 Thermophysical Properties of the Liquid and Vapor Phases

The thermophysical properties of density, specific heat capacity, thermal conductivity, kinematic and dynamic viscosity, and molecular weight for both the liquid and vapor phases of each of the four fluids, evaluated under saturation conditions, are summarized in Table 6.1.

Table 14 – Thermophysical Properties of the Four Two-Phase Fluids (Liquid and Vapor Phases under Saturation Conditions)

Fluid	Phase	Density [kg/m³]	c_p [J/(kg·K)]	Conductibility [W/(m·K)]	Kinematic viscosity [m²/s]	Dinamic viscosity [kg/(m·s)]	Molecular weigth [kg/kmol]
FC-72	Liquid	1680	1100	0.057	3.8×10^{-7}	6.384×10^{-4}	338
FC-72	Vapor	14.79	894.08	0.00909	5.399×10^{-7}	7.985×10^{-6}	338
Novec 7000	Liquid	1400	1300	0.075	3.214×10^{-7}	4.5×10^{-4}	200
Novec 7000	Vapor	16.2	980	0.012	6.790×10^{-7}	1.1×10^{-4}	200
Opteon SF-33	Liquid	1360	1200	0.077	2.794×10^{-7}	3.8×10^{-4}	164
Opteon SF-33	Vapor	16.5	920	0.011	6.667×10^{-7}	1.1×10^{-4}	164
R134a	Liquid	1207	1425	0.081	1.616×10^{-7}	1.95×10^{-4}	102.03
R134a	vapor	91.2	1150	0.016	1.425×10^{-7}	1.3×10^{-5}	102.03

The comparative analysis of the data reported in Table 6.1 highlights several preliminary observations that are relevant to the interpretation of the results presented in the following sections.

For all four fluids, the liquid phase exhibits a density significantly higher than that of water (ranging from approximately 1200 to 1680 kg/m³), consistently with the fluorinated nature of the compounds considered (see Section 3.5.2). FC-72 exhibits the highest liquid density (1680 kg/m³), whereas R134a has the lowest value (1207 kg/m³). Regarding the vapor phase, R134a exhibits a markedly higher density (91.2 kg/m³) than the other three fluids (ranging from approximately 14.8 to 16.5 kg/m³). This is a direct consequence of the high operating pressure required for R134a to reach the target saturation temperature (see Section 6.1.3), which results in a significantly denser vapor phase compared with the other candidate fluids, which operate at pressures close to atmospheric pressure.

The specific heat capacity of the liquid phase is highest for R134a (1425 J/(kg·K)) and lowest for FC-72 (1100 J/(kg·K)), with Novec 7000 and Opteon SF-33 exhibiting intermediate values (1300 and

1200 J/(kg·K), respectively). This parameter influences the ability of each fluid to remove sensible heat within the subcooled-liquid preheating region upstream of the onset of flow boiling (see Section 3.3.2).

The thermal conductivity of the liquid phase is relatively similar for Novec 7000, Opteon SF-33, and R134a (ranging from 0.075 to 0.081 W/(m·K)), whereas FC-72 exhibits a significantly lower value (0.057 W/(m·K)). This characteristic may adversely affect the single-phase convective heat transfer coefficient in the inlet region of the channel before the onset of boiling.

The dynamic viscosity of the liquid phase varies significantly among the four fluids, with FC-72 exhibiting the highest value (6.384×10^{-4} kg/(m·s)) and R134a the lowest value (1.95×10^{-4} kg/(m·s)). Together with density, this parameter contributes to determining the expected pressure losses within the cooling channel for each fluid, an aspect that will be further discussed in the comparative analysis presented in Section 6.5.

The four fluids exhibit significantly different molecular weights, ranging from the lowest value for R134a (102.03 kg/kmol) to the highest value for FC-72 (338 kg/kmol). This parameter is relevant to the determination of the vapor-phase density under the operating conditions of each simulation through the ideal gas equation of state (or more accurate formulations for real fluids).

6.3 Operating Conditions and Latent Heat of Vaporization

In order to ensure that each of the four fluids reaches saturation at the same target temperature (T_{sat}= 60 °C; see Section 4.4.2), the absolute operating pressure to be imposed within the computational domain was determined on the basis of the respective pressure-temperature saturation curves, together with the corresponding value of the latent heat of vaporization at the selected saturation temperature. These values, reported in Table 6.2 together with the pressure imposed at the outlet section of the computational domain (outlet pressure) and the standard state enthalpy value used for the correct definition of the material in Ansys Fluent, were also obtained from the same sources referenced in Section 6.1.1.

Table 15 – Saturation Conditions and Latent Heat of Vaporization of the Four Fluids (T_{sat} = 60 °C)

Fluid	Absolute pressure [Pa]	Outlet pressure [Pa]	Vaporization latent heat [kJ/kg]	Standard state enthalpy [J/kmol]
FC-72	113 500	12 000	83.54	28 236 520
Novec 7000	220 000	120 000	128	25 600 000
Opteon SF-33	240 000	140 000	153	25 092 000
R134a	1 680 000	1 580 000	139.1	14 192 373

The comparison among the operating pressures reported in Table 6.2 confirms the observations made in Section 3.5.2 regarding the atmospheric-pressure boiling temperatures of the four fluids. FC-72, whose normal boiling temperature (approximately 56 °C) is already close to the target value of 60 °C, requires an operating pressure only slightly above atmospheric pressure (113 500 Pa) to reach saturation at this temperature. Novec 7000 and Opteon SF-33, which have significantly lower normal boiling temperatures (approximately 34 °C and similar values, respectively), require moderately higher operating pressures (220 000 and 240 000 Pa, respectively). Finally, R134a, which has a normal boiling temperature well below 0 °C (approximately −26 °C), requires a markedly higher operating pressure than the other three candidates (1 680 000 Pa, corresponding to approximately 16.8 bar absolute). This represents a non-negligible design consideration in terms of the mechanical integrity of the cooling circuit and the rotating seals (see Section 3.2.3), should this fluid be selected for a potential real-world application.

With regard to the latent heat of vaporization, a quantity of primary importance for two-phase cooling performance (see Section 3.3.2), Opteon SF-33 exhibits the highest value among the four candidate fluids (153 kJ/kg), followed by Novec 7000 (128 kJ/kg) and R134a (139.1 kJ/kg), whereas FC-72 exhibits the markedly lowest value (83.54 kJ/kg), less than half that of Opteon SF-33. At the same mass flow rate employed in the simulations (see Section 4.4.2), this difference may result in a significantly different latent-heat removal capacity among the four fluids. This aspect will be quantitatively assessed through the comparison of the simulation results presented in the following sections of this chapter, in terms of the maximum temperature reached by the windings, outlet vapor fraction, and pressure drop for each two-phase configuration analyzed.

6.4 Rationale for the Selection of the Four Fluids

The selection of the four candidate fluids considered in the present comparison—Novec 7000, R134a, Opteon SF-33, and FC-72—was not arbitrary, but rather reflects the objective of exploring, within the range of dielectric fluids potentially suitable for direct rotor cooling (see the general selection criteria discussed in Section 3.5.1), a set of candidates sufficiently representative of the different chemical families available and their distinctive characteristics. This approach enables a broad comparative assessment rather than merely evaluating a single fluid in isolation.

As discussed in Section 3.5.2, the four selected fluids belong to different chemical families: hydrofluoroethers (Novec 7000), fluorocarbons/perfluorocarbons (FC-72), hydrofluoroolefins or next-generation blends (Opteon SF-33), and conventional hydrofluorocarbons (R134a). This selection makes it possible to determine whether the observed cooling performance can be attributed to specific characteristics of a particular chemical family or is instead governed more generally by the thermophysical properties of the individual fluid (latent heat, thermal conductivity, and viscosity), independently of the chemical family to which it belongs.

The inclusion of both FC-72, a dielectric fluid that has been used since the 1990s for the direct cooling of electronic components and has been extensively characterized in the experimental literature (see Section 3.4.1), and Opteon SF-33, a more recently developed fluid specifically designed for immersion-cooling applications with a more favorable environmental-impact profile, makes it possible to assess whether the environmental and regulatory advantages of next-generation fluids (see

Section 3.5.1) are accompanied by an actual improvement in thermal performance compared with historically established reference dielectric fluids.

Novec 7000, which has been widely employed in direct-cooling applications for electronic and electromechanical components and was adopted as the reference fluid during the preliminary phase of the present thesis work (see Section 3.5.2), represents the primary benchmark against which the performance of the other three candidates is evaluated, owing to its proven suitability for applications similar to that investigated in the present study.

The inclusion of R134a, a refrigerant not specifically designed for the direct cooling of electrical components but widely used in HVAC&R applications and, historically, in automotive air-conditioning systems, is motivated by two considerations. On the one hand, its widespread commercial availability, low cost, and extensive database of available experimental thermophysical properties (see Section 3.5.2) make it a useful "external" benchmark relative to dedicated dielectric fluids. On the other hand, its potential adoption could, in principle, enable integration with air-conditioning systems already installed on board the vehicle. Although this system-level consideration lies outside the thermofluid-dynamic objectives of the present work, it justifies the inclusion of R134a as a limiting comparison case, particularly in view of the high operating pressure required (see Section 6.1.3) and its lower electrical resistivity compared with the other three candidates (see Section 3.5.2), both of which make its suitability for direct cooling less straightforward than that of dedicated dielectric fluids.

Finally, as shown by the analysis of the data reported in Table 6.2, the four selected fluids cover a relatively broad range in terms of both latent heat of vaporization (from 83.54 kJ/kg for FC-72 to 153 kJ/kg for Opteon SF-33) and the operating pressure required to reach the target saturation temperature (from 113 500 Pa for FC-72 to 1 680 000 Pa for R134a). This makes it possible to investigate, within the comparison presented in the following sections, the influence of both parameters on the observed cooling performance, rather than limiting the analysis to fluids characterized by similar thermophysical and operating properties.

Overall, the selection of the four candidate fluids was therefore driven by the objective of maximizing the representativeness and general applicability of the conclusions that can be drawn from the present comparative study, compared with what could have been achieved by limiting the analysis to a smaller number of fluids or to fluids belonging to the same chemical family.

6.5 Comparison Methodology

In order to ensure that the performance differences observed among the four two-phase configurations analyzed can be attributed exclusively to the different thermophysical properties of the refrigerants employed, rather than to inconsistencies in the simulation methodology, the comparison among the four candidate fluids (Novec 7000, R134a, Opteon SF-33, and FC-72) was carried out using the same computational domain, structural boundary conditions, and numerical setup, as previously defined in Chapter 4. In particular, the following parameters were kept unchanged for all the two-phase configurations analyzed:

- the geometry of the computational domain, including the two cooling channels, the two copper half-windings, the ferromagnetic structure, and the polymeric partition (see Section 4.1);

- the computational grid selected following the mesh-independence assessment described in Section 4.2 (Mesh 1, medium);
- the volumetric heat generation applied to the winding domain, equal to 72.7 W for each modelled half-winding (see Section 4.1.5);
- the boundary conditions applied to the external walls subjected to convection (h_{conv} = 40 W/(m²·K), $T_{free\ stream}$ = 50 °C) and to the adiabatic cut surfaces of the computational domain (see Section 4.4.1);
- the fluid inlet temperature (T_{in} = 50 °C) and the target saturation temperature (T_{sat} = 60 °C), the latter being imposed for each fluid through the corresponding operating pressure determined from the respective saturation curves (see Section 6.1.3);
- the physical models (energy equation, k-ω SST turbulence model, Mixture multiphase model coupled with the Lee mass-transfer mechanism, with a Lee coefficient of 1 s^{-1}) and the numerical solution strategy (steady-state initialization followed by a transient simulation with gravity enabled) described in Sections 4.3 and 4.5.

The only discriminating variables among the four two-phase configurations are therefore the specific thermophysical properties of each fluid (density, specific heat capacity, thermal conductivity, liquid- and vapor-phase viscosities, and latent heat of vaporization), reported in Section 6.1.2. In this way, their influence on cooling performance can be isolated and correctly attributed.

6.5.1 Parametric Study of the Mass Flow Rate

In order to identify the most favorable operating condition for each of the candidate fluids, while simultaneously establishing a consistent basis for comparison among them, the mass flow rate parametric study was conducted on all four two-phase fluids considered in the present work. For each fluid, simulations were performed at progressively increasing mass flow rates, starting from the initial trial condition and gradually increasing the mass flow rate supplied to each individual cooling channel, while recording, for each operating point, the set of quantities of interest reported in Section 6.2.3.

The mass flow rate values investigated were kept identical for all four fluids, following the same sequence of increments. This methodological choice was based on a specific comparison criterion: by imposing the same fluid-dynamic inlet condition for all fluids, the only variable that effectively differentiates the simulations is the set of thermophysical properties of the refrigerant employed. Consequently, the observed differences in thermal performance can be directly attributed solely to the characteristics of the fluid—latent heat of vaporization, thermal conductivity and specific heat capacity of the liquid phase, density, and saturation pressure (see Section 6.1.2). Conversely, investigating different mass flow rate ranges for each fluid, although in principle aimed at separately optimizing each configuration, would have introduced a second independent variable, making the comparison among the fluids less straightforward and potentially ambiguous in terms of attributing the observed differences.

The adoption of a single set of mass flow rate values common to all fluids also makes it possible to determine whether the performance hierarchy among the four refrigerants remains consistent as the

mass flow rate varies or, conversely, whether it is sensitive to the specific operating condition considered. This provides an indication of the robustness of the conclusions drawn regarding the selection of the most suitable fluid for the application. The investigated mass flow rate range and the results obtained for each fluid are presented and discussed in the following sections, together with the identification of the optimal operating point and the physical rationale underlying its selection.

6.5.2 Criteria for the Selection of the Optimal Operating Point

The identification of the optimal operating point for each fluid within the parametric mass-flow-rate study described in Section 6.2.2 was carried out based on multiple competing criteria, rather than simply minimizing the maximum winding temperature, for the reasons outlined below.

An indiscriminate increase in mass flow rate, although progressively reducing the maximum winding temperature (owing to the higher convective heat transfer coefficient and the greater ability of the fluid to remove sensible heat before and during boiling), results in a corresponding reduction in the vapor fraction produced at the channel outlet (vapor fraction outlet). This progressively shifts the behavior of the system towards that of a single-phase configuration, in which the contribution of the latent heat of vaporization (see Section 3.3.2) would be only marginally exploited. Since the objective of the present thesis is specifically to assess the benefits of phase change as a heat-removal mechanism (see Chapter 3), the optimal operating point was identified by giving preference, all other conditions being equal, to operating points at which the outlet vapor fraction reaches a representative and consistent value for the flow-boiling regime, approximately 50%. This value was considered an appropriate compromise between substantial utilization of latent heat (a non-negligible vapor quality) and the preservation of an adequate safety margin against the critical heat-flux phenomena discussed in Section 3.3.3 (dryout, associated with high vapor qualities and progressive thinning of the liquid film at the wall).

For each operating point identified according to the previous criterion, it was verified that the maximum temperature recorded in the winding ($T_{\max winding}$) remained within an adequate safety margin with respect to the temperature limit imposed by the thermal class of the electrical insulation adopted (see Section 2.2.2). This condition is necessary to ensure that the identified operating point is effectively representative of a safe operating configuration, rather than merely a formal optimum based solely on the vapor-fraction criterion.

Although the pressure-drop values observed in the two-phase configurations analyzed were, as shown by the parametric study presented by way of example in this chapter, consistently limited (on the order of tens of pascals, compared with the hundreds to thousands of pascals observed in the reference single-phase case; see Chapter 5), this quantity was nevertheless monitored for each operating point in order to exclude the occurrence of significant fluid-dynamic penalties as the mass flow rate increased. However, as observed, this constraint was never found to be practically limiting within the range of mass flow rates investigated for the two-phase fluids.

Finally, the overall energy balance of the computational domain was verified for each operating point according to the same methodology adopted for the single-phase case (see Section 5.2.1), discarding any operating points characterized by a non-negligible energy-balance residual (**net result**) relative

to the total imposed thermal power, as this would indicate an inadequately converged numerical solution.

6.5.3 Performance Metrics Adopted for the Final Comparison

Following the identification of the optimal operating point for each of the four candidate fluids according to the criteria described in Section 6.2.3, the final comparative assessment among the four two-phase configurations, as well as between these configurations and the reference single-phase configuration presented in Chapter 5, was carried out on the basis of the following performance metrics, which are discussed in detail in Section 6.5:

- **maximum copper winding temperature ($T_{\max winding}$):** the primary indicator of the available thermal safety margin and, ultimately, of the maximum excitation current density that can be sustained by the machine (see Section 2.5.4);
- **maximum temperatures of the ferromagnetic structure and polymeric partition ($T_{\max lamination}$, $T_{\max plastic}$):** complementary indicators of the temperature-field distribution within the computational domain;
- **outlet vapor fraction (vapor fraction outlet):** an indicator of the degree to which the latent heat of vaporization is exploited by each fluid under the identified operating conditions;
- **pressure drop (pressure drop):** between the channel inlet and outlet, representing an indicator of the fluid-dynamic penalty associated with each configuration and, indirectly, of the pumping power required by the cooling circuit.

Taken together, these performance metrics, applied consistently to the four two-phase configurations and the reference single-phase configuration, provide the basis for the overall comparative assessment presented in Section 6.5 of this chapter, following the presentation of the fluid-specific results in Sections 6.3 and 6.4.

6.6 Parametric Study of the Mass Flow Rate: Identification of the Optimal Mass Flow Rate

The identification of the optimal refrigerant mass flow rate represents a necessary prerequisite for a meaningful comparison among the four candidate two-phase fluids. Unlike single-phase cooling systems, in which an increase in mass flow rate generally results in a substantially monotonic thermal benefit, in two-phase systems the mass flow rate plays a dual and partly conflicting role: on the one hand, it determines the overall heat removal capacity of the system; on the other hand, it governs the flow regime within the channel and, consequently, the extent to which the latent heat of vaporization is exploited. An excessively low mass flow rate leads to high vapor qualities and the consequent risk of dryout, whereas an excessively high mass flow rate reduces the evaporated fraction to values that cause the system behavior to approach that of single-phase cooling, thereby undermining its conceptual advantage. Consequently, the optimal mass flow rate cannot be

determined through a purely analytical approach, but requires a parametric investigation supported by numerical simulation.

In the present work, the parametric study was conducted on all four candidate fluids—Novec 7000, FC-72, Opteon SF-33, and R134a—by investigating the same sequence of mass flow rates per individual channel for each fluid: 0.568, 1.25, 2.50, 3.00, 3.50, and 4.00 g/s, for a total of twenty-four simulations. The adoption of a common set of mass flow rates for all fluids is based on a specific methodological criterion: by imposing the same fluid-dynamic inlet condition, the only variable that effectively differentiates the simulations is the set of thermophysical properties of the refrigerant employed, thus allowing the observed differences in thermal performance to be directly attributed solely to the characteristics of the fluid. An investigation conducted over fluid-specific mass flow rate ranges, although in principle aimed at separately optimizing each configuration, would have introduced a second independent variable, making the comparison less straightforward and potentially ambiguous in terms of attributing the observed differences.

A further advantage of this approach lies in the possibility of determining whether the performance hierarchy among the four refrigerants remains consistent as the mass flow rate varies or, conversely, whether it is sensitive to the specific operating condition considered. This represents relevant information for assessing the robustness of the conclusions drawn regarding the selection of the most suitable fluid for the application.

In order to ensure full comparability of the results, all twenty-four simulations were carried out using the same sequential solution strategy described in Section 4.5.2, consisting of a steady-state initialization of 6000 iterations, followed by a first transient phase with a time step of 0.0001 s for 1000 time steps and a second transient phase with a time step of 0.001 s for 10000 time steps, corresponding to a total simulated physical time of approximately 10.1 s. The thermal and fluid-dynamic boundary conditions (fluid inlet temperature of 50 °C, saturation temperature of 60 °C, and Lee model coefficients set to unity), the spatial discretization of the domain, and the adopted numerical schemes were also kept unchanged throughout the entire study, whereas the thermophysical properties of the two phases, the saturation pressure, and the static pressure imposed at the outlet were updated on a case-by-case basis according to the fluid considered (see Section 6.1.2).

6.6.1 First-Estimate Mass Flow Rate

The first mass flow rate investigated in the parametric study was determined through an initial analytical estimate obtained by assuming that the entire thermal power to be removed from the computational domain (User source = 145.42 W; see Section 4.1.5) was dissipated exclusively through the latent heat of vaporization, according to the relationship

$$m_{attempt\ 1} = \frac{Q_{tot}}{h_{fg,Novec\ 7000}} = \frac{145.42\ W}{128000\ {}^{J}/_{kg}} = 0.001136\ \frac{kg}{s}$$

$$\left(0.000568\ \frac{kg}{s}\ single\ channel\right)$$

The selection of Novec 7000 as the reference fluid for this preliminary estimate, common to all four cases, is motivated by several considerations. First, its latent heat of vaporization (128 kJ/kg) is close to the average value of the four candidate fluids, equal to 125.9 kJ/kg and calculated from the values of 83.54 kJ/kg for FC-72, 128 kJ/kg for Novec 7000, 139.1 kJ/kg for R134a, and 153 kJ/kg for Opteon SF-33. Novec 7000 therefore occupies an intermediate position within the investigated range, providing a balanced basis for the estimate that is neither systematically favorable nor detrimental to any of the fluids. Second, the adoption of a single initial mass flow rate for all fluids, rather than four separate estimates each based on the corresponding latent heat, is consistent with the methodological approach described in the previous section, which requires the same set of mass flow rates to be investigated for all candidates in order to preserve the consistency of the comparison. Third, Novec 7000 represents a reference fluid in the technical literature concerning the direct cooling of electromechanical components (see Section 3.5.2), and its adoption as the basis for the preliminary calculation allows the estimate to be related to a well-established benchmark.

It should be emphasized that this estimate is purely indicative, as it was obtained by neglecting the contribution of the sensible heat required to raise the fluid temperature from the inlet condition (subcooled liquid at T_{in} = 50 °C) to the saturation temperature (T_{sat} = 60 °C), and without taking into account the local two-phase flow regime. It therefore represents a theoretical lower limit for the required mass flow rate, ideally corresponding to a condition in which the entire mass of fluid introduced into the channel completely evaporates, thereby exploiting the full available latent heat.

6.6.2 Results and Limitations of the First-Estimate Mass Flow Rate

The simulations conducted using the initial trial mass flow rate (0.568 g/s per channel) demonstrated, for all four fluids, the inadequacy of this simplified estimate, yielding thermally unacceptable results. The maximum winding temperature reached 230.96 °C for Novec 7000, 248.72 °C for FC-72, 223 °C for Opteon SF-33, and 189.94 °C for R134a: in all cases, these values exceed — and, for the first three fluids, significantly exceed — the thermal class limit of the insulation (see Section 2.2.2). At the same time, the outlet vapor fraction reaches extremely high values for three of the four fluids (90.81% for Novec 7000, 93.61% for FC-72, and 87.74% for Opteon SF-33), approaching complete evaporation of the refrigerant, whereas R134a, owing to its higher heat capacity in the liquid phase, reaches a significantly lower value (49.37%), consistently with the better thermal performance exhibited by this fluid under these conditions.

This result, although apparently in contradiction with the objective of maximizing the exploitation of latent heat (see Section 6.2.3), is in fact fully consistent with the physics of convective boiling discussed in Section 3.3.3: an excessively low mass flow rate, while in principle allowing complete evaporation of the fluid, leads to the development of a high vapor quality already in the intermediate portions of the channel, resulting in the thinning and eventual breakdown of the liquid film at the wall (dryout) and a drastic deterioration of the local heat transfer coefficient in the affected regions, which is responsible for the anomalous increase in wall temperature observed. This interpretation is supported by the magnitude of the maximum temperatures reached within the channel, ranging from 161.13 °C (R134a) to 227.73 °C (FC-72), and by the fact that the fluid reaches outlet temperatures well above the saturation temperature — up to 102.56 °C for FC-72, compared with a T_{sat} of 60 °C — providing clear evidence of the presence of strongly superheated vapor in contact with walls under dryout conditions.

The initial trial estimate, based exclusively on an overall energy balance, therefore proved to be unrepresentative of a safe and effective operating condition for any of the four fluids analyzed, confirming the need for an exploratory parametric approach rather than a direct analytical determination of the mass flow rate.

Table 16 – Simulation Results at a Mass Flow Rate of 0.000568 kg/s per Channel

Parameter	Novec 7000	FC-72	Opteon SF-33	R134a
T_{out} [°C]	84.37	102.56	81.06	66.75
$T_{av\ duct}$ [°C]	106.62	117.67	107.13	92.35
$T_{max\ duct}$ [°C]	209.96	227.73	199.94	161.13
$T_{av\ w}$ [°C]	208.28	229.29	206.11	175.4
$T_{max\ w}$ [°C]	230.96	248.72	223	189.94
$T_{av\ lamination}$ [°C]	199.13	215.44	196.14	168.62
$T_{max\ lamination}$ [°C]	210.33	223.83	204	175.11
p_{drop} [Pa]	8.108	6.446	6.1	9.4
Liquid fraction [%]	12.53	8.89	12.42	26.99
Vapor fraction [%]	87.46	91.11	87.57	73.01
Vapor fraction outlet [%]	90.81	93.61	87.74	49.37

6.6.3 Progressive Increase in Mass Flow Rate: Results for the Four Fluids

Starting from the initial trial condition, a parametric analysis was carried out by progressively increasing the mass flow rate supplied to each cooling channel, according to the sequence 1.25 – 2.50 – 3.00 – 3.50 – 4.00 g/s, applied identically to all four fluids. The mass flow rate increments were defined by taking, as an order-of-magnitude reference, the typical flow rates employed in refrigeration circuits for automotive cabin air-conditioning systems, an application field characterized by dimensional similarities with the cooling system investigated in the present study. The objective is to identify the operating condition capable of ensuring an adequate thermal margin for the windings for all candidate fluids, while simultaneously maintaining effective utilization of the latent heat of vaporization.

In the following subsections, the results are presented grouped according to the different categories of monitored quantities in order to facilitate their comparative interpretation; all temperatures are expressed in degrees Celsius and all mass flow rates in grams per second per individual channel.

Duct and fluid temperature

Table 27 – Temperature outlet all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [°C]	84.37	67.78	62.78	62.05	60.99	60.18
FC-72 [°C]	102.56	74.88	68.3	64.83	62.4	62.69
Opteon SF-33 [°C]	81.06	68.66	62.3	61.27	60.25	60.43
R134a [°C]	66.75	63.66	62.3	61.27	60.25	60.43

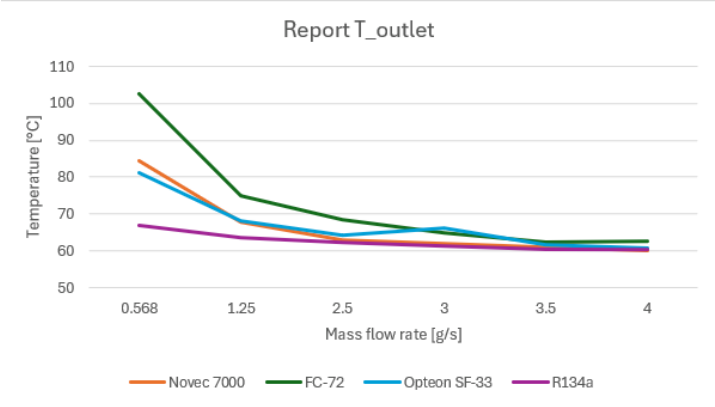


Figure 53 - Report temperature outlet all four fluids

Table 18 - Temperature average duct all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [°C]	106.62	80.6	67.73	65.67	63.41	61.78
FC-72 [°C]	117.67	88.49	73.92	69.5	67.36	65.72
Opteon SF-33 [°C]	107.13	81.29	68.2	67.9	63.72	61.69
R134a [°C]	92.35	75.27	65.7	61.79	60.33	58.43

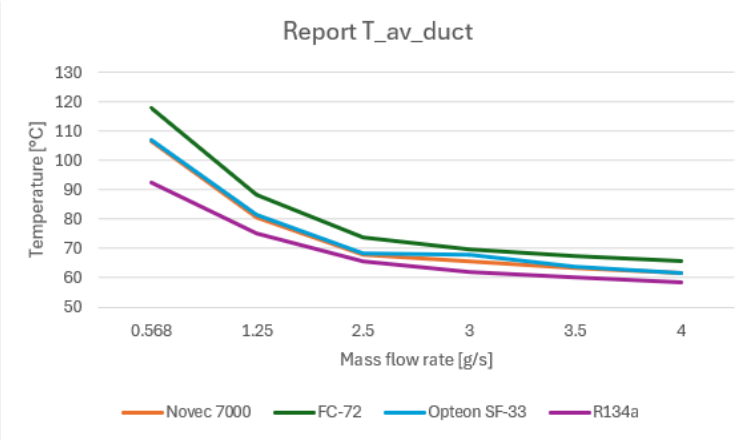


Figure 54 - Report temperature average duct all four fluids

Table 19 - Temperature max duct all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [°C]	209.96	159.22	128.67	125.81	119.9	114.23
FC-72 [°C]	227.73	182.24	149.53	138.69	140.04	131.31
Opteon SF-33 [°C]	199.94	158.29	128.4	123.68	118.66	114.02
R134a [°C]	161.13	128.62	107.68	99.9	99.94	92.86

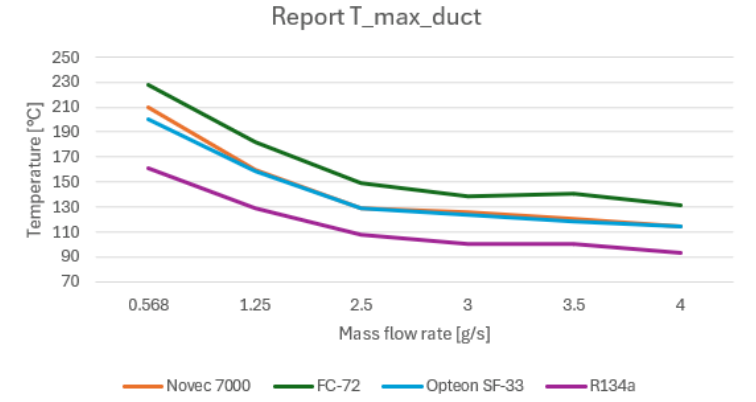


Figure 55 – Report temperature max duct all four fluids

The fluid temperature trends show that, for all refrigerants, increasing the mass flow rate leads to a progressive convergence of the outlet temperature toward the saturation temperature. Indeed, the fluid transitions from a highly superheated vapor condition, typical of the lower mass flow rates, to a saturated two-phase mixture condition, in which the temperature remains close to the saturation temperature due to the ongoing evaporation process. This convergence occurs more rapidly for fluids with a higher heat removal capability: R134a reaches an outlet temperature below 63 °C already at 2.50 g/s, whereas FC-72 reaches this condition only at 3.50 g/s. Nevertheless, some minor local reversals can be observed—such as the increase in the outlet temperature of Opteon SF-33 between

2.50 and 3.00 g/s, or those of FC-72 and R134a between 3.50 and 4.00 g/s—which can be attributed to the transient nature of the simulated phenomenon and to the fact that the reported values correspond to quantities evaluated at the end of a finite simulation time window and are therefore subject to residual numerical fluctuations.

Winding temperature

Table 20 - Temperature average winding all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [°C]	208.28	168.54	141.36	139.09	133.01	128.48
FC-72 [°C]	229.29	190.47	162.59	153.29	149.73	146
Opteon SF-33 [°C]	206.11	168.08	141.34	136.02	133.35	128.59
R134a [°C]	175.4	146.82	129.1	123.67	120.67	118.27

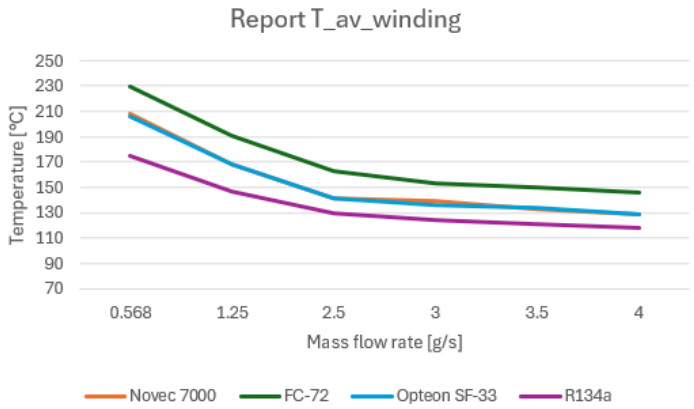


Figure 56 - Report temperature average winding all four fluids

Table 21 - Temperature max winding all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [°C]	230.96	185.71	157.87	156.13	149.91	145.12
FC-72 [°C]	248.72	207.62	179.04	169.5	166.52	162.18
Opteon SF-33 [°C]	223	184.13	157.79	152.26	150.21	145.04
R134a [°C]	189.94	162.11	144.98	139.51	136.65	134.34

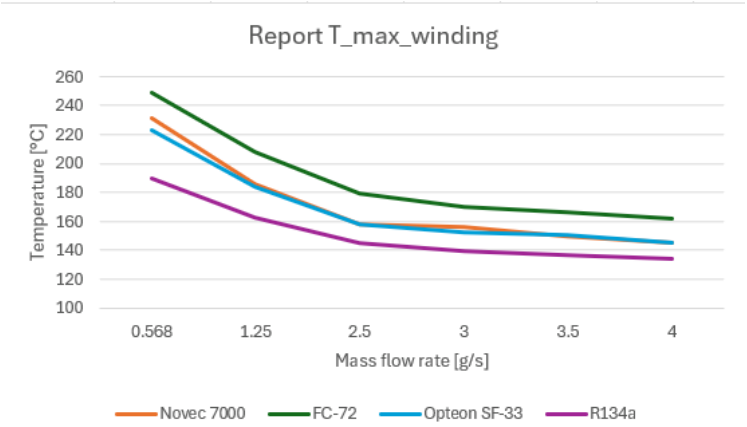


Figure 57 - Report temperature max winding all four fluids

The maximum winding temperature represents the most relevant parameter from a design perspective, as it is directly constrained by the thermal class limit of the electrical insulation. For all four fluids, it decreases monotonically as the mass flow rate increases, with an overall reduction ranging from 55.6 °C (R134a) to 86.5 °C (FC-72) over the investigated range. Of particular interest is the identification of the mass flow rate at which each fluid falls below the reference threshold of 180 °C: R134a reaches this condition already at a mass flow rate of 1.25 g/s, Novec 7000 and Opteon SF-33 at 2.50 g/s, whereas FC-72 requires a mass flow rate of 3.00 g/s. Furthermore, the performance ranking among the four fluids remains strictly unchanged throughout the entire investigated mass flow rate range—with R134a consistently exhibiting the best performance, followed by Novec 7000 and Opteon SF-33 with nearly equivalent performance, and FC-72 consistently showing the lowest performance—thus confirming the robustness of the comparative conclusions presented in Section 6.4 and their substantial independence from the specific mass flow rate considered.

Lamination and plastic insulation temperature

Table 22 - Temperature average lamination all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [°C]	199.13	164.38	140.54	138.68	132.8	128.26
FC-72 [°C]	215.44	182.44	158.59	150.44	148.84	144.69
Opteon SF-33 [°C]	196.14	163.55	140.17	135.24	133.43	128.47
R134a [°C]	168.62	143.01	127.58	122.6	120.17	117.61

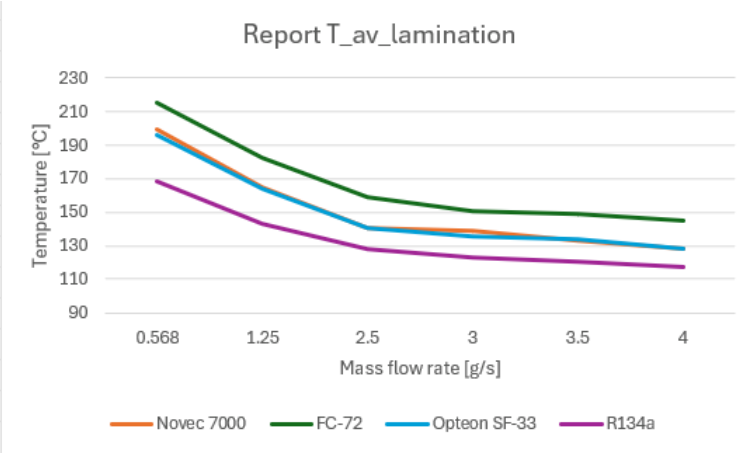


Figure 58 - Report temperature average lamination all four fluids

Table 23 - Temperature max lamination all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [°C]	210.33	171.25	146.66	144.58	138.61	134.07
FC-72 [°C]	223.83	189.72	165.08	156.54	155.28	150.88
Opteon SF-33 [°C]	204	170.31	146.2	140.8	139.2	134.13
R134a [°C]	175.11	148.58	132.74	127.83	125.56	122.91

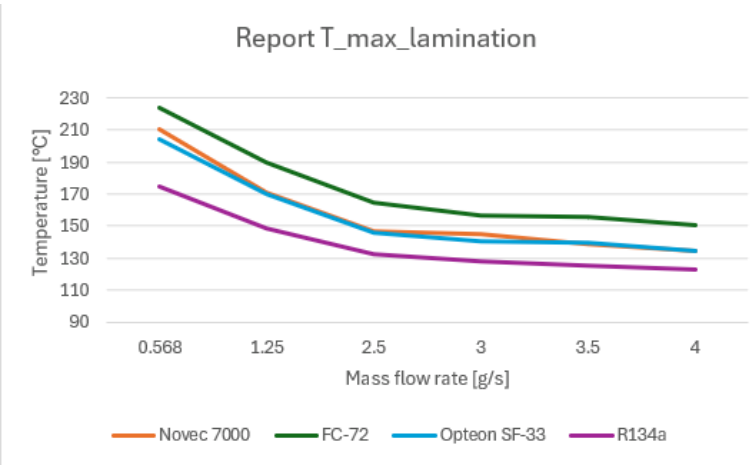


Figure 59 - Report temperature max lamination all four fluids

Table 24 - Temperature average plastic

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [°C]	-	105.04	91.38	89.11	84.69	82.56
FC-72 [°C]	-	118.56	101.39	96.26	92.96	91.27
Opteon SF-33 [°C]	-	104.86	91.44	89.21	87.03	82.47
R134a [°C]	-	93.57	83.61	79.01	76.92	74.87

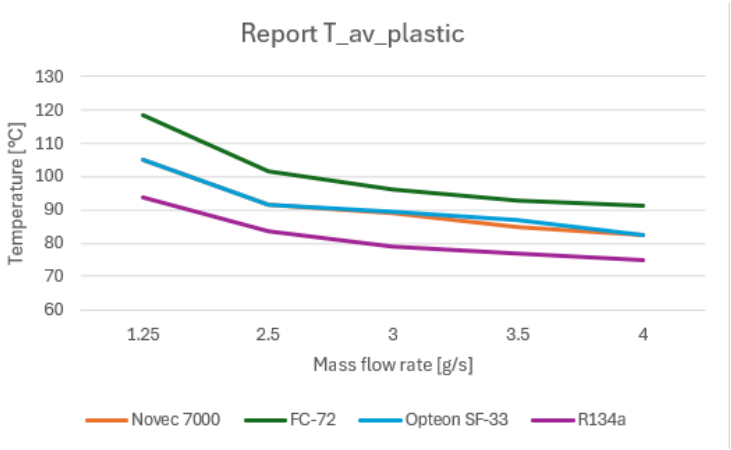


Figure 10 - Report temperature average plastic all four fluids

Table 25 - Temperature max plastic all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [°C]	-	157.02	134.11	132.06	126.33	121.77
FC-72 [°C]	-	174.82	151.42	143.05	141.6	137.89
Opteon SF-33 [°C]	-	155.75	132.96	128.81	127.66	121.45
R134a [°C]	-	134.95	119.34	114.02	112.29	109.42

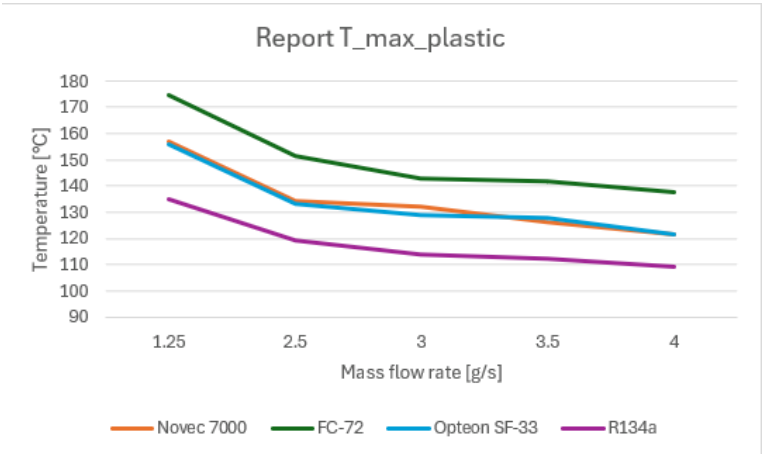


Figure 11 - Report temperature max plastic all four fluids

The thermal benefit resulting from the increase in mass flow rate is consistently transferred to all the solid components of the computational domain, although with different magnitudes depending on their position relative to the cooling channel and their respective thermal resistance. The plastic insulator, characterized by a thermal conductivity of only 0.2 W/m·K and located between the fluid and the windings, systematically exhibits lower maximum temperatures than the windings, while still maintaining a significant internal temperature gradient, as evidenced by the large difference between the mean and maximum values, exceeding 40 °C in all the configurations analyzed. The values relating to the plastic insulator are not available for the initial trial mass flow rate, as they were not monitored during that preliminary simulation.

Pressure drop

Table 26 – Pressure drop all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [Pa]	8.108	24.22	47.08	54.68	62.23	69.22
FC-72 [Pa]	6.446	17.091	35.231	39.911	62.717	60.606
Opteon SF-33 [Pa]	6.1	15.75	30.33	37.6	60.78	51.49
R134a [Pa]	9.4	15.1	34.8	40	64.3	57.2

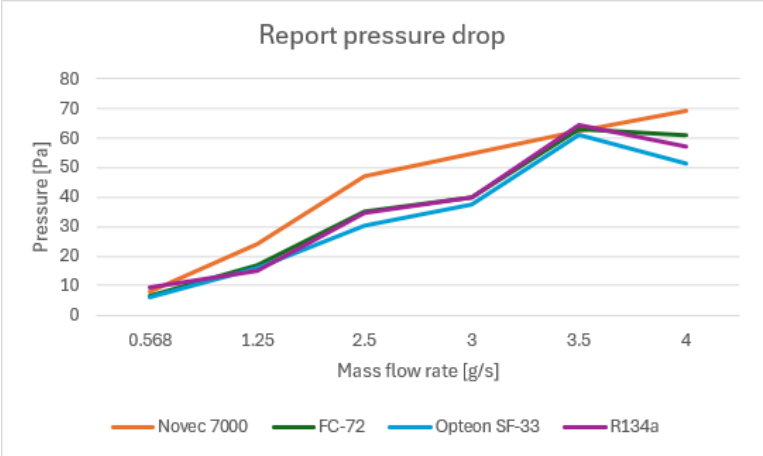


Figure 12 - Report pressure drop all four fluids

The pressure drop across the channel increases almost monotonically with mass flow rate for all the fluids, following a trend that is qualitatively consistent with the increase in mean flow velocity and with the acceleration contribution associated with the variation in mixture density along the channel. In all cases, the values remain extremely low in absolute terms, never exceeding 70 Pa under the most demanding operating condition, and can therefore be considered entirely negligible for the sizing of the refrigerant supply circuit and the associated pumping unit.

However, for FC-72, Opteon SF-33, and R134a, the pressure drop at a mass flow rate of 4.00 g/s is lower than that recorded at 3.50 g/s, in apparent contradiction with the expected trend. This behavior may be physically explained by the two-phase nature of the flow, since the overall pressure drop

depends not only on the mass flow rate but also on the vapor fraction and the resulting mean density of the mixture. Nevertheless, in the absence of experimental validation, a numerical contribution associated with the residual unsteadiness of the solution cannot be ruled out. This evidence was therefore interpreted with caution and was not used as a decisive criterion in the selection of the optimal operating point, particularly since the differences involved are negligible from a design perspective.

Phase fraction

Table 27 – Liquid fraction all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [%]	12.53	27.67	49.09	55.59	62.28	67.79
FC-72 [%]	8.89	19.62	37.12	43.51	49.72	53.64
Opteon SF-33 [%]	12.42	26.4	47.86	54.52	62.8	66.44
R134a [%]	26.99	50	73.2	81.64	89.21	93.69

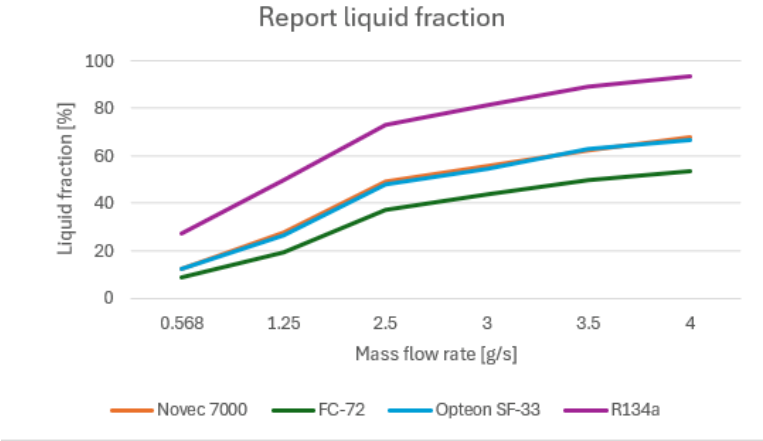


Figure 13 - Report liquid fraction all four fluids

Table 28 – Vapor fraction all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [%]	87.46	72.32	50.9	44.4	37.72	32.2
FC-72 [%]	91.11	80.37	62.87	56.49	50.27	46.35
Opteon SF-33 [%]	87.57	73.59	52.13	45.48	37.19	33.55
R134a [%]	73.01	49.99	26.79	18.35	10.78	6.31

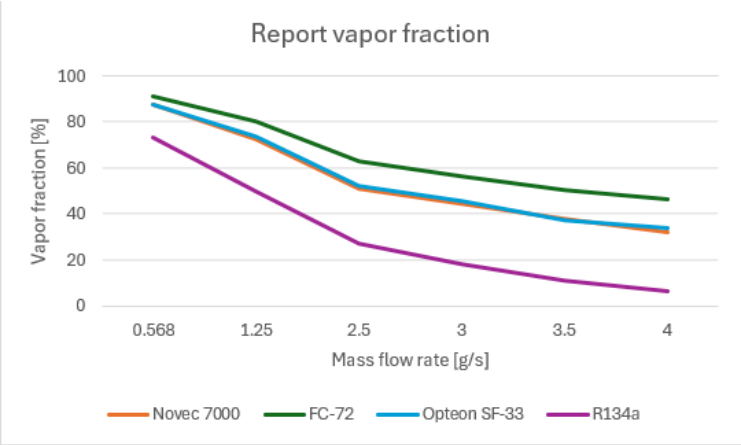


Figure 14 - Report vapor fraction all four fluids

Table 29 – Vapor fraction outlet all four fluids

Fluid	0.568 [g/s]	1.25 [g/s]	2.5 [g/s]	3 [g/s]	3.5 [g/s]	4 [g/s]
Novec 7000 [%]	90.81	75.6	56	51.18	45.2	40.6
FC-72 [%]	93.61	85.21	73.07	67.38	59.96	59.45
Opteon SF-33 [%]	87.74	77.07	61.13	56.87	45.33	46.84
R134a [%]	49.37	37.32	23.14	17.97	12.53	8.99

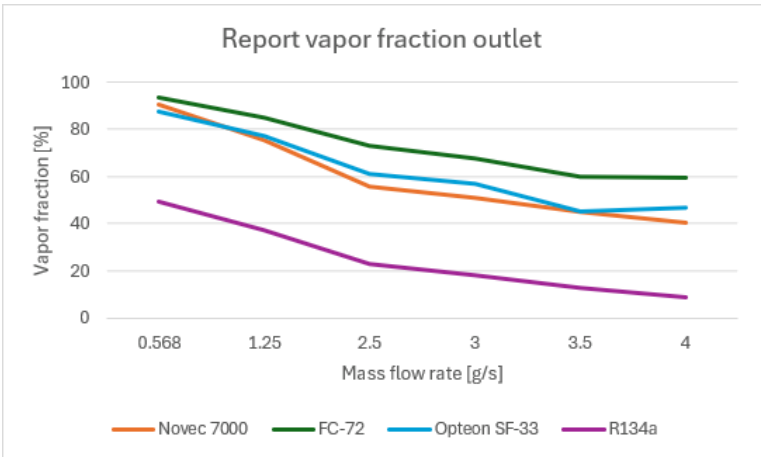


Figure 15 - Report vapor fraction outlet all four fluids

The analysis of the phase fractions represents the most relevant aspect for characterizing the two-phase operating regime. For all the fluids, as expected, a progressive decrease in vapor fraction is observed as the mass flow rate increases, since an increasing mass of refrigerant is available to remove the same thermal power. However, the behavior of the four fluids differs markedly: Novec 7000 and Opteon SF-33 exhibit nearly overlapping trends, with the outlet vapor fraction decreasing from values close to 90% to approximately 45% at a mass flow rate of 3.50 g/s; FC-72, penalized by the lowest latent heat among the four candidates, consistently maintains higher vapor fractions, never falling below 59% even at the highest mass flow rate investigated; R134a, by contrast, exhibits a much more

pronounced reduction, with the outlet vapor fraction decreasing to 12.53% at 3.50 g/s and to only 8.99% at 4.00 g/s.

The latter result deserves specific consideration: at the higher mass flow rates, R134a operates under a regime in which the contribution of latent heat to the overall heat transfer becomes marginal, with the refrigerant remaining in the liquid phase for almost the entire flow path. Although R134a remains the fluid providing the best thermal performance, this condition effectively results in an operating regime close to that of a single-phase system, moving away from the principle of latent heat utilization on which the cooling strategy investigated in the present work is based. While this aspect does not compromise the absolute thermal performance, it reduces the distinctiveness of the conceptual advantage associated with two-phase cooling and will be discussed further in Chapter 7.

6.6.4 Identification of the Optimal Operating Point

The selection of the optimal mass flow rate cannot be based solely on minimizing the maximum winding temperature, since such an approach would systematically favor the highest mass flow rates, resulting in an increase in the refrigerant flow rate required by the system, the associated pumping power, and, as highlighted in the previous section, a progressive reduction in the effectiveness of the two-phase heat transfer mechanism, which represents the fundamental principle underlying the adopted cooling strategy. It is therefore necessary to identify the operating point that provides the best compromise between thermal margin, limitation of the overall mass flow rate, and maintenance of an adequate vapor fraction to effectively exploit the latent heat of the refrigerant.

Since the operating point identified in this manner must provide a common basis for the subsequent comparison of the four fluids under identical supply conditions, the selection criterion was applied to the average values of the monitored quantities across the four candidate fluids, rather than to the results obtained for a single refrigerant. This approach ensures that the selected mass flow rate is not optimized according to the characteristics of a particular fluid, and therefore implicitly favorable to it, but instead represents a balanced operating condition for the entire set of candidates. Table 6.16 reports the average values, calculated across the four fluids, of the maximum winding temperature and outlet vapor fraction, together with the marginal thermal benefit associated with each increase in mass flow rate.

Table 30 - Average values and thermal benefit all four fluids

Mass flow rate [g/s]	$T_{max\,winding}$ media [°C]	Thermal benefit [°C]	Vapor fraction outlet average [%]
0.568	223.26	-	80.38
1.25	184.89	-38.27	68.8
2.5	159.92	-24.97	53.34
3	154.85	-5.57	48.35
3.5	150.82	-3.53	40.76
4	146.67	-4.15	38.97

The analysis of the marginal thermal benefit clearly highlights the existence of two distinct operating regimes. In the first part of the investigated range, the increase from 0.568 to 1.25 g/s and that from 1.25 to 2.50 g/s produce reductions in the average maximum winding temperature of 38.3 °C and 25.0 °C, respectively. These particularly significant improvements can be attributed to the progressive departure from the critical dryout conditions discussed in Section 6.3.3. Beyond a mass flow rate of 2.50 g/s, the trend changes substantially: each further increase of 0.50 g/s produces a thermal benefit ranging from 3.5 °C to 5.6 °C, indicating that the system has entered a region of markedly diminishing thermal returns, in which its performance is already approaching an asymptotic limit. It should also be noted that, within this region, the marginal benefit does not decrease in a strictly monotonic manner, since the increase from 3.50 to 4.00 g/s is slightly more effective than that from 3.00 to 3.50 g/s (−4.15 °C compared with −3.53 °C). This difference is marginal in magnitude and comparable with the residual numerical uncertainty of the solution, confirming that, within this portion of the investigated range, the thermal criterion alone is insufficient to discriminate among the different options and must therefore be complemented by the additional criteria discussed below.

Based on a combined assessment of all the criteria considered, the operating condition identified as the optimal operating point corresponds to a mass flow rate of 3.50 g/s per individual channel, equivalent to an overall mass flow rate of 56 g/s for the entire rotor. The reasons underlying this choice are summarized below.

- **Adequate thermal margin for all candidate fluids:** at a mass flow rate of 3.50 g/s, the maximum winding temperature remains below the thermal class limit of the insulation for all four refrigerants, with a margin ranging from 13.5 °C in the worst-performing case (FC-72, 166.52 °C) to 43.4 °C in the best-performing case (R134a, 136.65 °C). At the immediately lower mass flow rate of 3.00 g/s, the thermal margin for FC-72 would decrease to only 10.5 °C, a value considered insufficient to ensure a robust operating condition with respect to the uncertainties of the numerical model discussed in Chapter 7. This margin is particularly relevant from the perspective of the long-term reliability of the electrical insulation, whose service life is strongly influenced by operating temperature according to Arrhenius-type degradation laws.
- **Diminishing thermal returns:** as highlighted by the analysis of the marginal thermal benefit, beyond 2.50 g/s the system enters a region in which the thermal gain per unit increase in mass flow rate decreases by approximately one order of magnitude compared with the initial portion of the investigated range. A further increase from 3.50 to 4.00 g/s would result in a 14.3% increase in the overall mass flow rate while providing a reduction in the average maximum temperature of only 4.2 °C, thus making the 4.00 g/s mass flow rate an oversized operating condition with respect to the actual cooling demand and unjustified in terms of the cost-benefit ratio of the refrigerant supply system.
- **Effective utilization of latent heat:** at a mass flow rate of 3.50 g/s, the outlet vapor fraction is 45.20% for Novec 7000 and 45.33% for Opteon SF-33, values that ensure a balanced contribution between latent and sensible heat transfer, while avoiding both the risk of local dryout associated with excessively high vapor fractions, as observed at mass flow rates of 0.568 and 1.25 g/s, and an excessive shift toward nearly single-phase operation. Maintaining an intermediate vapor fraction also preserves a safety margin against complete evaporation of the refrigerant along the channel, a condition that would compromise heat transfer effectiveness in the downstream portion of the channel.

- **Limitation of the overall refrigerant mass flow rate:** a mass flow rate of 3.50 g/s per channel corresponds to 56 g/s for the entire rotor, a value that, as will be discussed in Section 6.5, is approximately four times lower than the mass flow rate required by the reference single-phase cooling solution, thereby preserving the main system-level advantage of the two-phase strategy in terms of supply-pipe dimensions and required pumping power.
- **Negligible pressure drop:** at the selected mass flow rate, the pressure drop remains within a range of 60.78 Pa to 64.30 Pa for all the fluids, values that are entirely negligible for the sizing of the supply circuit and substantially equivalent among the four refrigerants. Consequently, this parameter does not introduce any discriminating factor among the candidates.

In conclusion, a mass flow rate of 3.50 g/s per individual channel represents the best compromise among the different design requirements: it ensures an adequate thermal margin for the windings for all four candidate fluids, lies close to the knee of the marginal thermal benefit curve, limits the overall refrigerant mass flow rate required by the system, and preserves, for most of the fluids, effective utilization of the latent heat of vaporization. This operating condition was therefore adopted as the reference operating point for the detailed comparison among the four two-phase fluids presented in the following section.

6.7 Comparison of the Four Fluids at the Optimal Mass Flow Rate

Once the common optimal mass flow rate had been identified (3.50 g/s per individual channel, corresponding to an overall mass flow rate of 56 g/s for the entire rotor), the results of the four simulations corresponding to this operating condition were reconsidered and examined in detail in order to fully characterize the behavior of each refrigerant at the selected operating point and to provide the basis for the comparative analysis developed in Section 6.5.

It should be emphasized that, since this mass flow rate was kept identical for all the fluids and all the remaining boundary conditions were maintained strictly unchanged—namely, a fluid inlet temperature of 50 °C, a saturation temperature of 60 °C, Lee model coefficients equal to unity, the same spatial discretization of the computational domain, and the same numerical schemes—the only substantial differences among the four simulations lie in the sets of thermophysical properties of the two phases and in the saturation pressure, which were updated on a case-by-case basis according to the refrigerant considered. Consequently, the observed differences in thermal performance can be entirely attributed to the physical properties of the fluids, in accordance with the methodological approach outlined in Section 6.3.1.

It should be specifically noted that, by imposing the same saturation temperature of 60 °C for all the candidates, the corresponding saturation pressure varies significantly from one fluid to another, ranging from 113.5 kPa for FC-72 to 1.68 MPa for R134a, reflecting the different vapor-pressure curves of the four refrigerants. Although this circumstance has no direct consequences for the thermal results of the simulations, it is of considerable relevance from a system-design perspective, since it affects the sealing and mechanical strength requirements of the refrigerant supply circuit. This aspect will be further addressed and discussed in Chapter 7 as part of the overall assessment of the design trade-offs.

The following subsections present, for each fluid, the thermofluid-dynamic results and the corresponding energy balance obtained at the end of the simulation, together with a technical discussion aimed at correlating the observed performance with the characteristic physical properties of the refrigerant under consideration.

6.7.1 Novec 7000

For Novec 7000 at a mass flow rate of 0.0035 kg/s per channel, the maximum winding temperature reaches 149.91 °C, with an outlet vapor fraction of 45.2%, indicating a partial evaporation regime with a well-balanced distribution between the liquid and vapor phases. This balance is consistent with its intermediate latent heat of vaporization compared with the other fluids considered (128 kJ/kg) and with a liquid-phase thermal conductivity of 0.075 W/(m·K), which together provide a favorable compromise between the amount of fluid evaporated and the thermal resistance of the liquid film at the wall.

The thermal balance exhibits a relative imbalance of 1.75%, a limited value that confirms the good numerical quality of the simulation. The pressure drop (62.23 Pa) is comparable to that obtained for the other two-phase fluids, confirming that, at the considered mass flow rate, fluid-dynamic behavior does not represent a discriminating factor in refrigerant selection, unlike the thermal properties.

Table 31 – Novec 7000 Two-Phase Simulation Results (0.0035 kg/s per Channel)

Parameter	**Value**
T_{out} **[°C]**	60.99
$T_{av\ duct}$ **[°C]**	63.41
$T_{max\ duct}$ **[°C]**	119.9
$T_{av\ w}$ **[°C]**	133.01
$T_{max\ w}$ **[°C]**	149.91
$T_{av\ lamination}$ **[°C]**	132.8
$T_{max\ lamination}$ **[°C]**	138.61
$T_{av\ plastic}$ **[°C]**	84.69
$T_{max\ plastic}$ **[°C]**	126.33
p_{drop} **[Pa]**	62.23
Liquid fraction [%]	62.28
Vapor fraction [%]	37.72
Vapor fraction outlet [%]	45.2

Table 32 – Energy Balance of the Novec 7000 Two-Phase Simulation (0.0035 kg/s per Channel)

Parameter	Value
User source	145.42
$Q_{inlet\ 1}$ [W]	113.74
$Q_{inlet\ 2}$ [W]	113.74
$Q_{outlet\ 1}$ [W]	-170.6
$Q_{outlet\ 2}$ [W]	-168.6
$Q_{conv\ plastic}$ [W]	-3.545
$Q_{conv\ lamination}$ [W]	-32.7
Net results [W]	-2.546

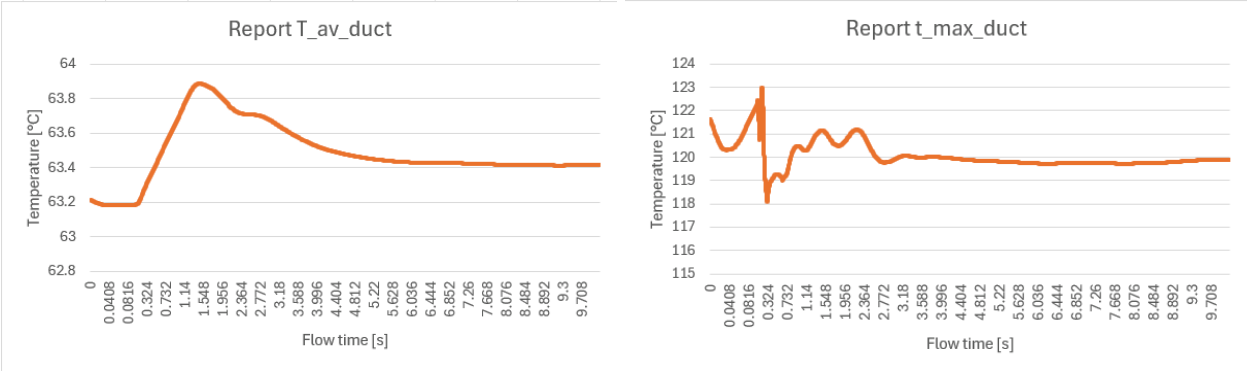


Figure 66 – Average Cooling-Channel Temperature Report, Novec 7000
Figure 67 – Maximum Cooling-Channel Temperature Report, Novec 7000

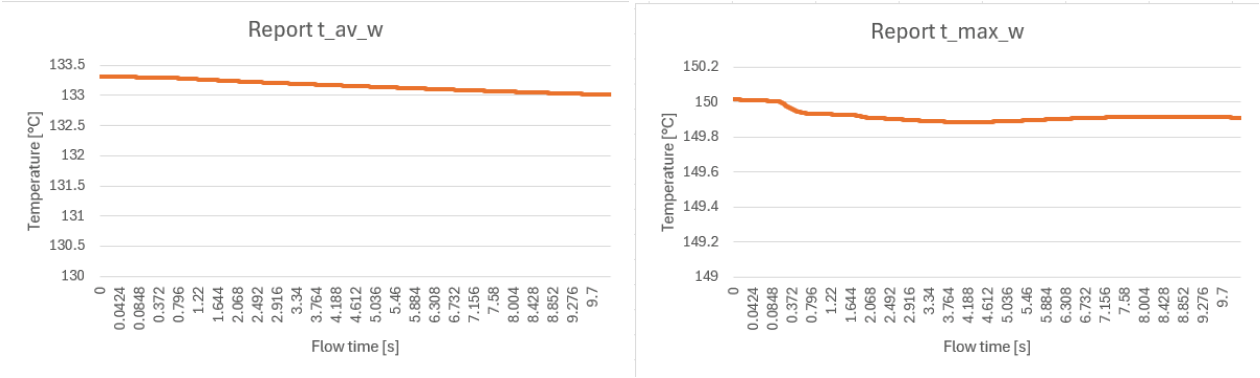
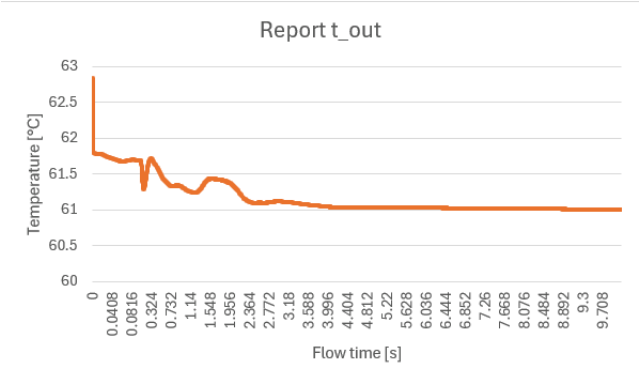


Figure 68 – Average Winding Temperature Report, Novec 7000
Figure 69 – Maximum Winding Temperature Report, Novec 7000



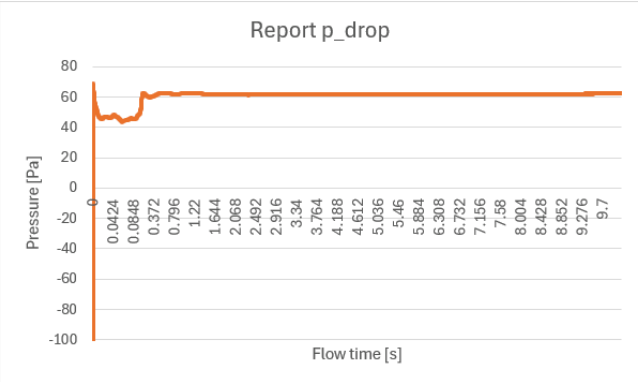
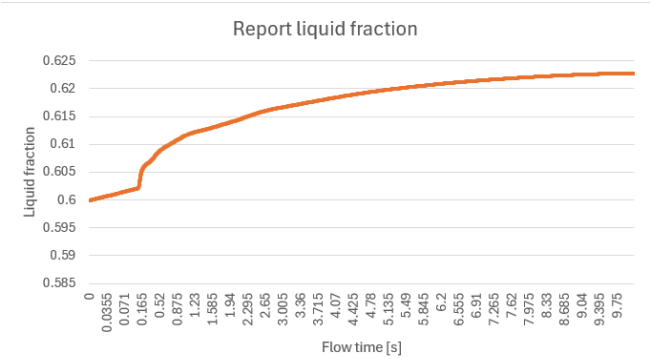


Figure 70 – Novec 7000 Outlet Temperature Report
Figure 71 – Pressure Drop Report, Novec 7000



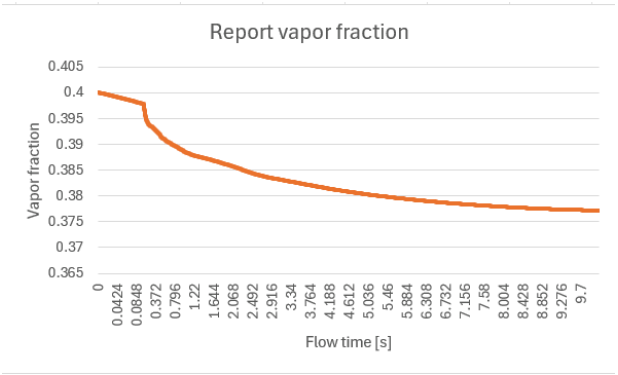


Figure 72 – Liquid Fraction Report, Novec 7000
Figure 73 – Vapor Fraction Report, Novec 7000

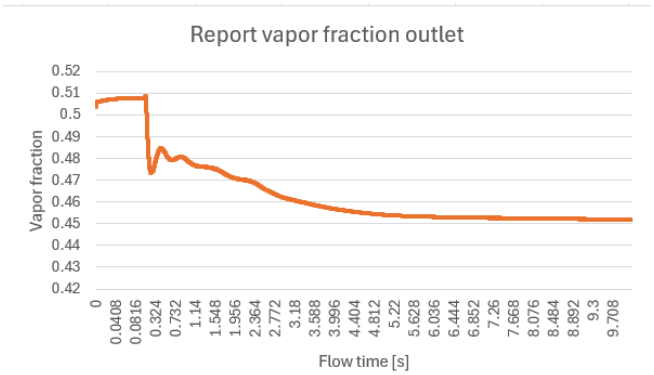


Figure 74 – Outlet Vapor Fraction Report, Novec 7000

Contour case Novec 7000

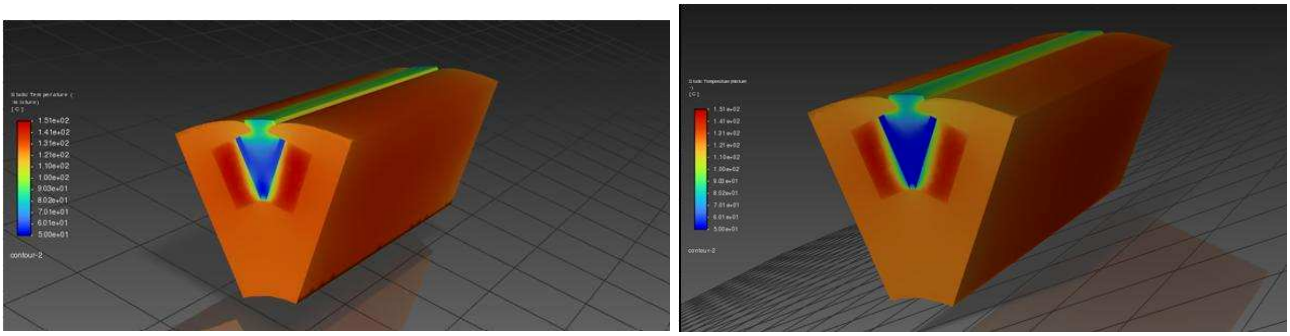


Figure 75 – Electric Rotor Computational-Domain Contour, Outlet Side, Isometric View
Figure 76 – Electric Rotor Computational-Domain Contour, Inlet Side, Isometric View

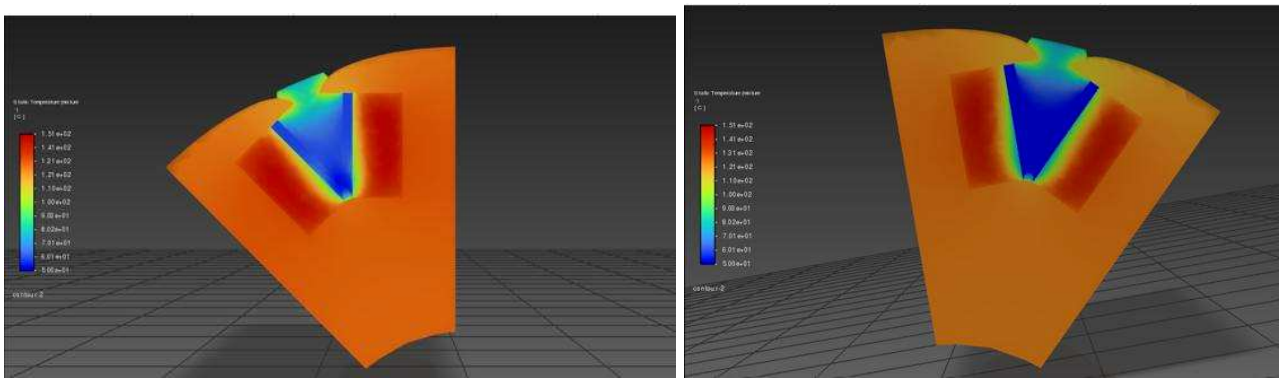


Figure 77 – Electric Rotor Computational-Domain Contour, Outlet Side, Front View
Figure 78 – Electric Rotor Computational-Domain Contour, Inlet Side, Front View

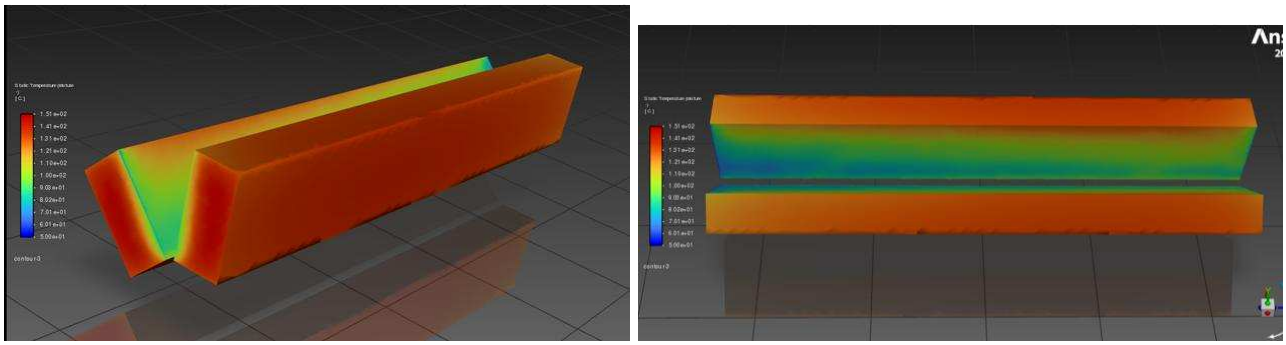


Figure 79 – Copper-Winding Contour, Isometric View
Figure 80 – Copper-Winding Contour, Top View

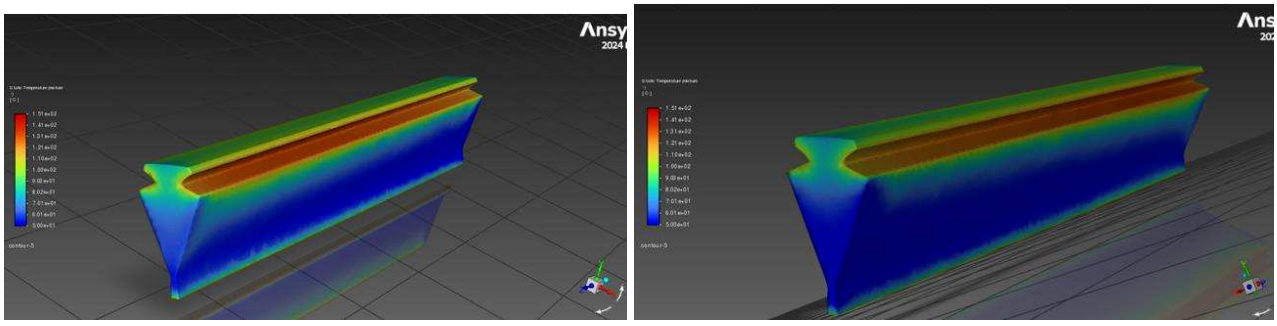


Figure 81 – Plastic-Partition Contour, Outlet Side, Isometric View
Figure 82 – Plastic-Partition Contour, Inlet Side, Isometric View

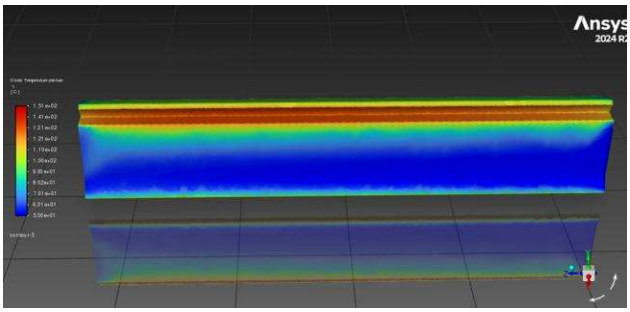
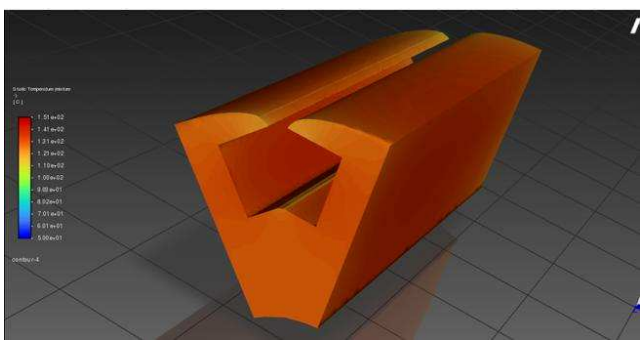


Figure 83 – Plastic-Partition Contour, Side View



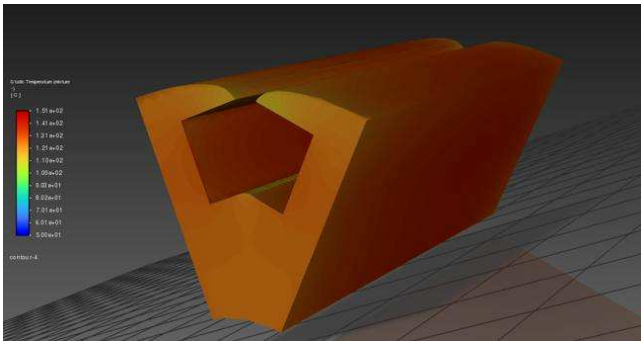
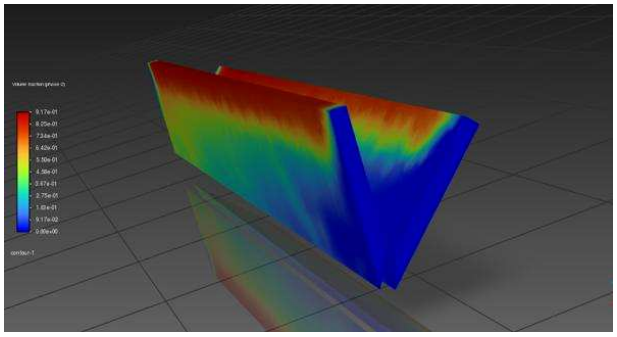


Figure 84 – Rotor Ferromagnetic-Structure Contour, Outlet Side
Figure 85 – Rotor Ferromagnetic-Structure Contour, Inlet Side



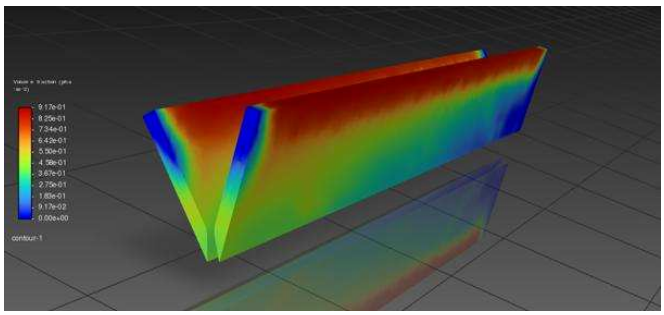


Figure 86 – Vapor-Phase Contour in the Cooling Channels, Inlet Side
Figure 87 – Vapor-Phase Contour in the Cooling Channels, Outlet Side

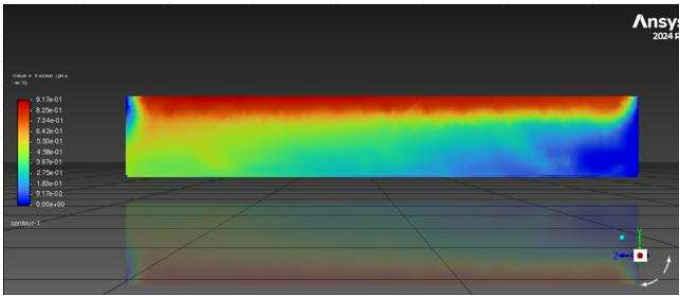


Figure 88 – Vapor-Phase Contour in the Cooling Channels, Side View

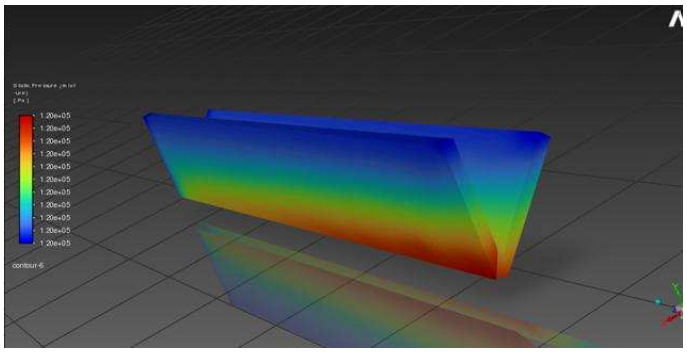


Figure 89 – Pressure-Drop Contour in the Cooling Channels, Inlet Side

6.7.2 FC-72

FC-72, which belongs to the perfluorocarbon family (Fluorinert, 3M), is characterized by the lowest latent heat of vaporization among the four candidate fluids (83.54 kJ/kg, corresponding to approximately 65% of that of Novec 7000) and by the lowest thermal conductivity in both the liquid phase (0.057 W/(m·K)) and the vapor phase (0.00909 W/(m·K)). These properties significantly penalize two-phase heat transfer: for the same thermal power to be removed, a lower latent heat requires the evaporation of a proportionally greater mass of fluid, while a lower thermal conductivity increases the thermal resistance of the liquid film adjacent to the heated wall, thereby increasing the wall superheat required to sustain the same heat flux.

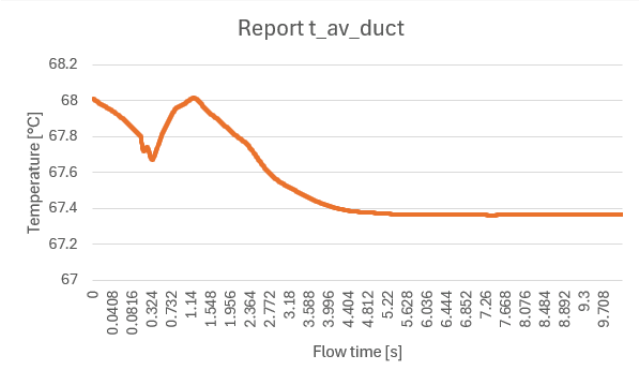
Consistently with these considerations, FC-72 exhibits the poorest thermal performance among all the fluids considered in the comparison: the maximum winding temperature reaches 166.52 °C, approximately 16.6 °C (+11.1%) higher than that obtained with Novec 7000 at the same mass flow rate. The outlet vapor fraction is the highest among the four fluids (59.96%), confirming that a larger fraction of the fluid evaporates in order to remove the same thermal power. However, this does not translate into an improvement in temperature control, owing to the adverse effect of the lower thermal conductivity. The energy balance also exhibits the highest relative imbalance among the cases analyzed (3.74%), likely associated with the greater unsteadiness of the evaporation front along the channel, which makes convergence of the Lee model more challenging for this specific fluid.

Table 33 – FC-72 Two-Phase Simulation Results (0.0035 kg/s per Channel)

Parameter	**Value**
T_{out} **[°C]**	62.4
$T_{av\ duct}$ **[°C]**	67.36
$T_{max\ duct}$ **[°C]**	140.04
$T_{av\ w}$ **[°C]**	149.73
$T_{max\ w}$ **[°C]**	166.52
$T_{av\ lamination}$ **[°C]**	148.84
$T_{max\ lamination}$ **[°C]**	155.28
$T_{av\ plastic}$ **[°C]**	92.96
$T_{max\ plastic}$ **[°C]**	141.6
p_{drop} **[Pa]**	62.717
Liquid fraction [%]	49.72
Vapor fraction [%]	50.27
Vapor fraction outlet [%]	59.96

Table 34 – Energy Balance of the FC-72 Two-Phase Simulation (0.0035 kg/s per Channel)

Parameter	Value
User source	145.42
$Q_{inlet\ 1}$ [W]	96.24
$Q_{inlet\ 2}$ [W]	96.24
$Q_{outlet\ 1}$ [W]	-150.78
$Q_{outlet\ 2}$ [W]	-149.05
$Q_{conv\ plastic}$ [W]	-4.31
$Q_{conv\ lamination}$ [W]	-39.2
Net results [W]	-5.44



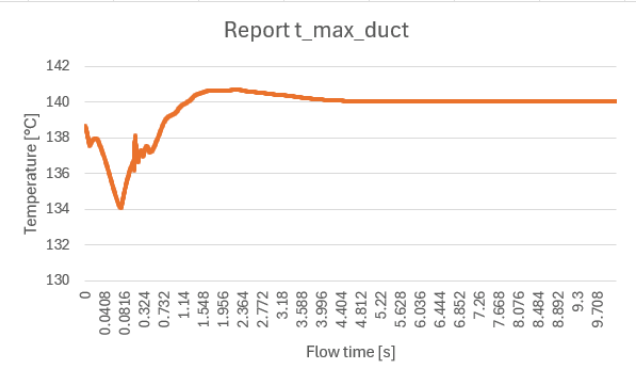
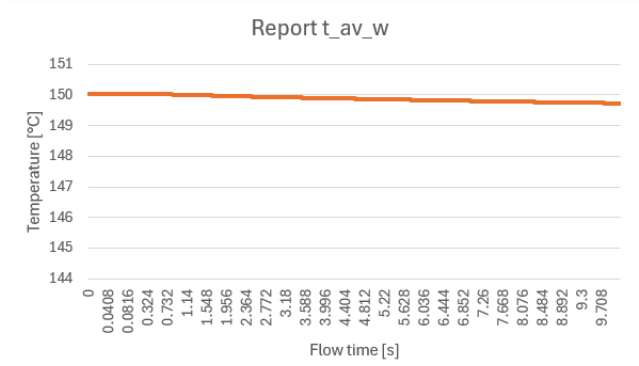


Figure 90 – Average Cooling-Channel Temperature Report, FC-72
Figure 91 – Maximum Cooling-Channel Temperature Report, FC-72



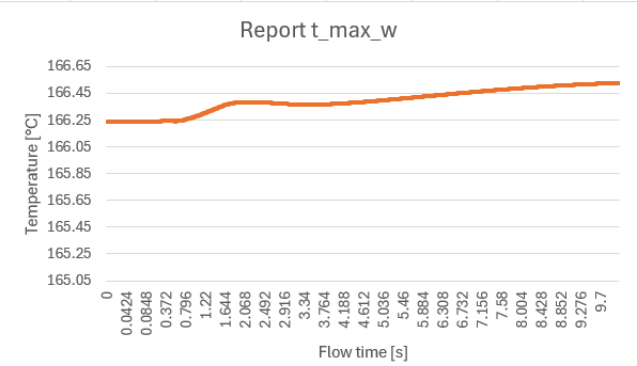
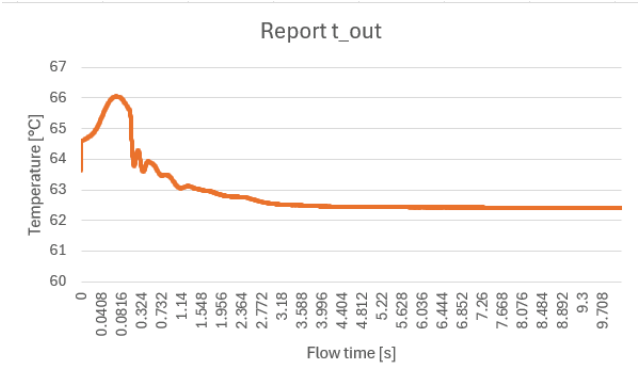


Figure 92 – Average Winding Temperature Report, FC-72
Figure 93 – Maximum Winding Temperature Report, FC-72



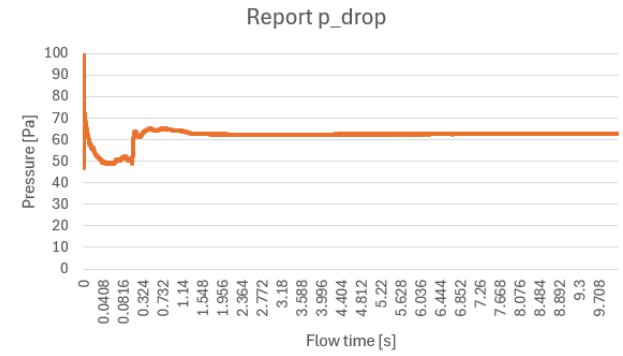
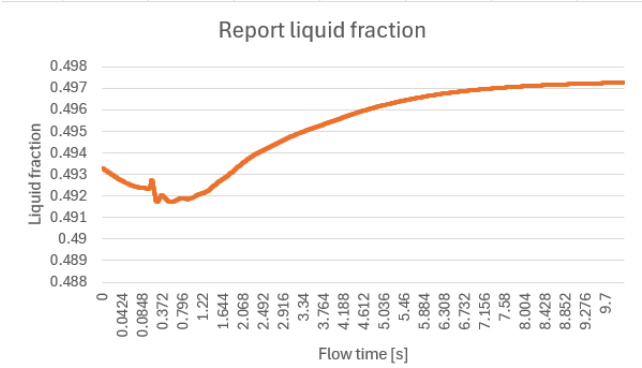


Figure 94 – FC-72 Outlet Temperature Report
Figure 95 – Pressure Drop Report, FC-72



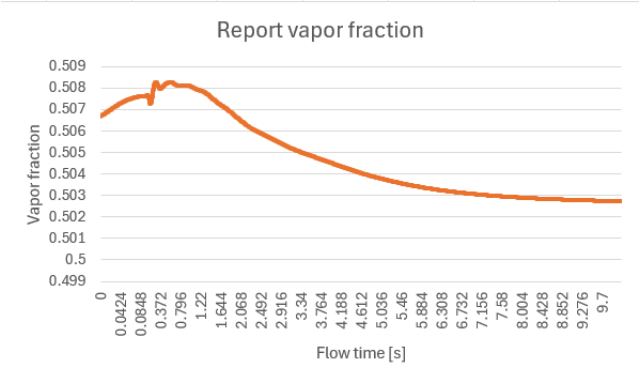


Figure 96 – Liquid Fraction Report, FC-72
Figure 97 – Vapor Fraction Report, FC-72

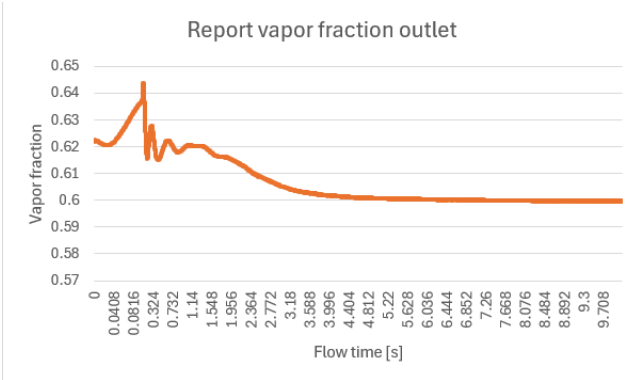


Figure 98 – Outlet Vapor Fraction Report, FC-72

Contour case FC-72

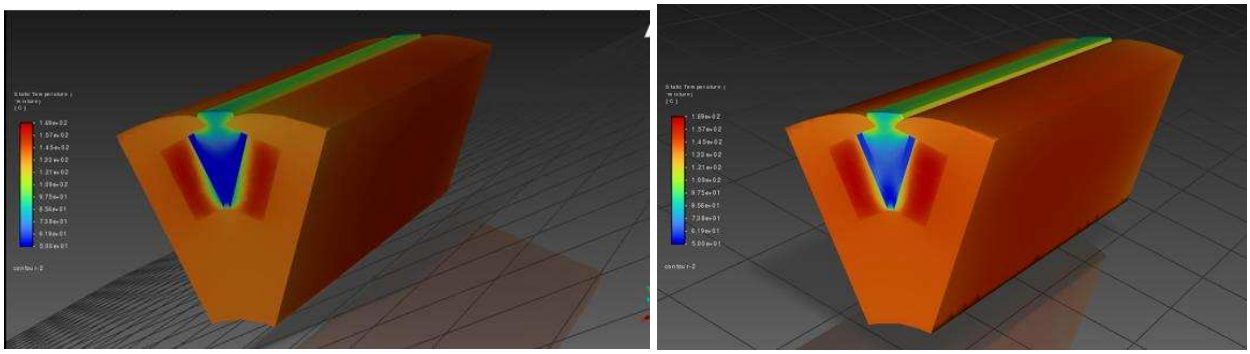


Figure 99 – Electric Rotor Computational-Domain Contour, Inlet Side, Isometric View
Figure 100 – Electric Rotor Computational-Domain Contour, Outlet Side, Isometric View

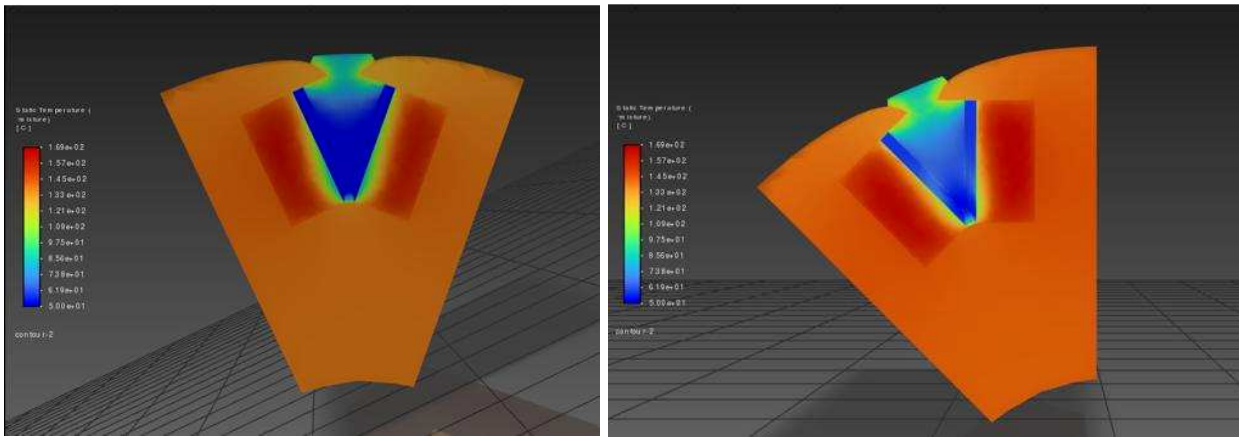


Figure 101 – Electric Rotor Computational-Domain Contour, Inlet Side, Front View
Figure 102 – Electric Rotor Computational-Domain Contour, Outlet Side, Front View

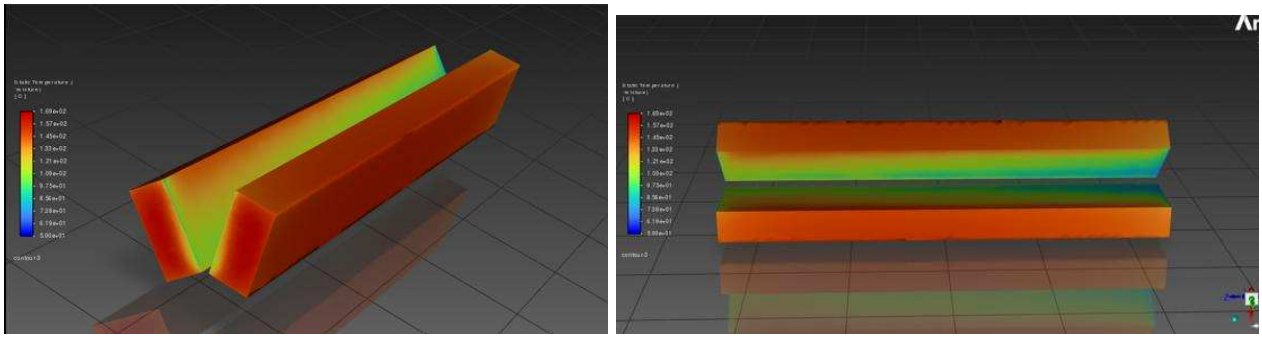
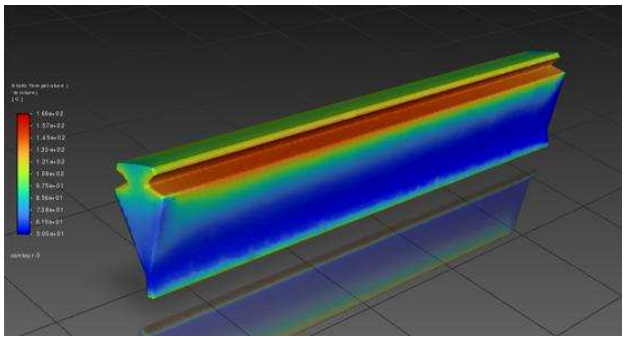


Figure 103 – Copper-Winding Contour, Isometric View
Figure 104 – Copper-Winding Contour, Top View



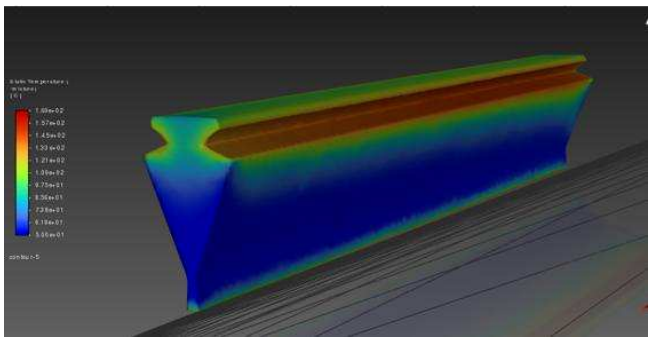


Figure 105 – Plastic-Partition Contour, Outlet Side, Isometric View
Figure 106 – Plastic-Partition Contour, Inlet Side, Isometric View

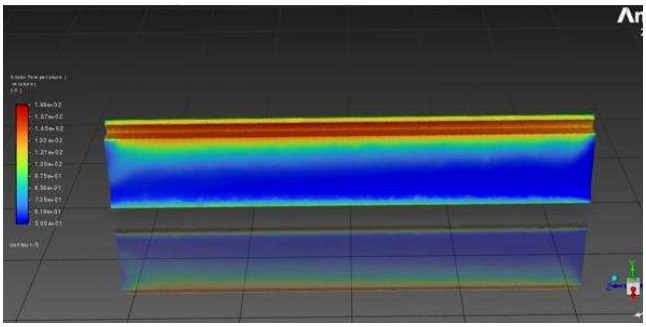
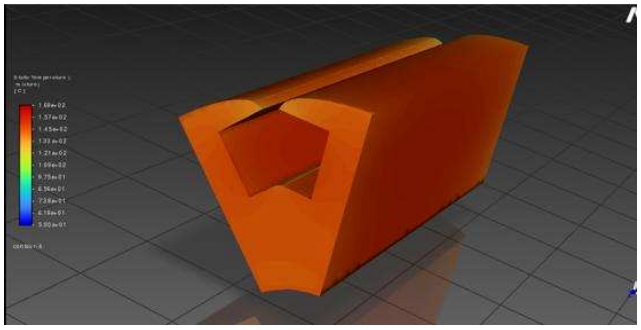


Figure 107 – Plastic-Partition Contour, Side View



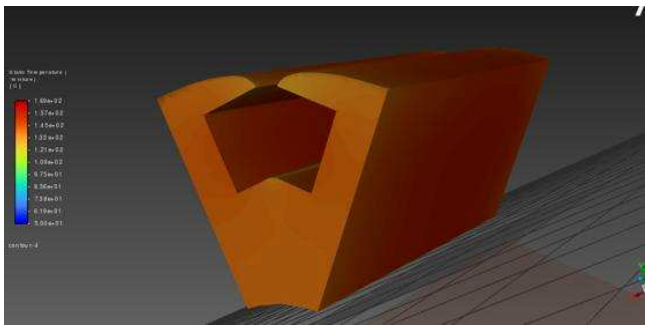
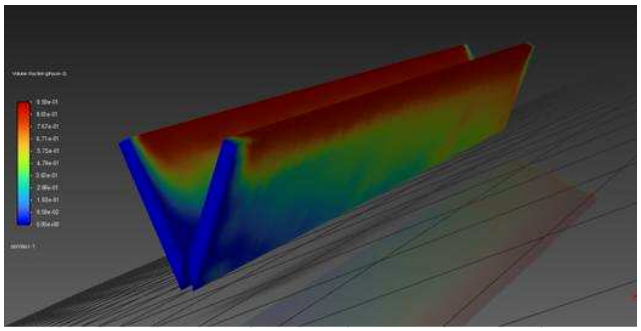


Figure 108 – Rotor Ferromagnetic-Structure Contour, Outlet Side
Figure 109 – Rotor Ferromagnetic-Structure Contour, Inlet Side



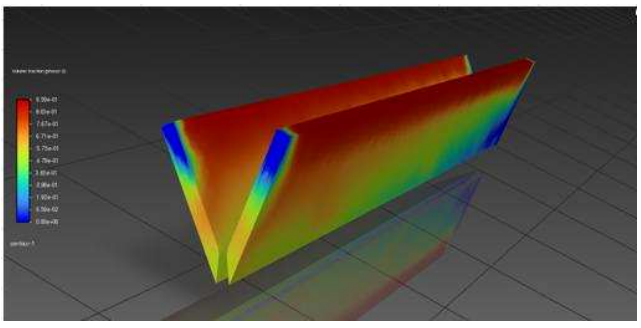


Figure 110 – Vapor-Phase Contour in the Cooling Channels, Inlet Side
Figure 111 – Vapor-Phase Contour in the Cooling Channels, Outlet Side

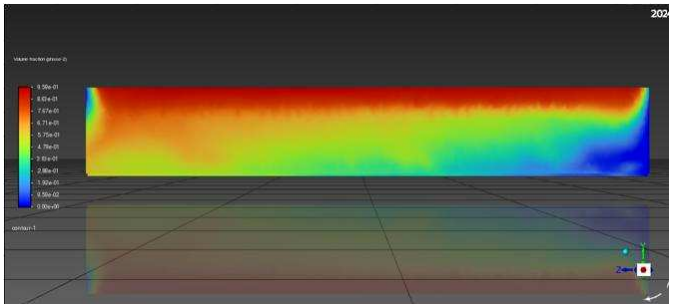


Figure 112 – Vapor-Phase Contour in the Cooling Channels, Side View

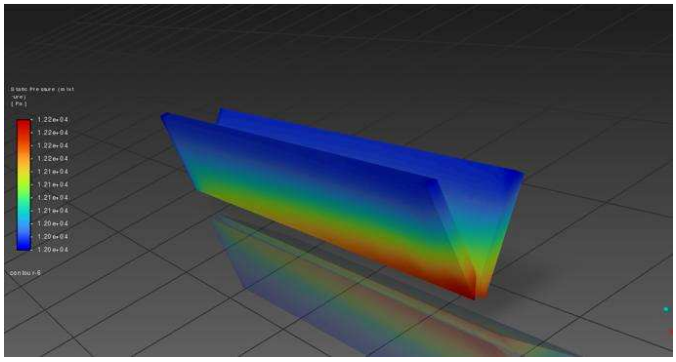


Figure 113 – Pressure-Drop Contour in the Cooling Channels, Inlet Side

6.7.3 Opteon SF-33

Opteon SF-33 (Chemours) exhibits thermophysical properties very similar to those of Novec 7000, both in terms of latent heat of vaporization (153 kJ/kg, the highest among the four fluids) and liquid-phase thermal conductivity (0.077 W/(m·K), a value close to the 0.075 W/(m·K) of Novec 7000). This similarity in thermophysical properties is reflected in simulation results that are virtually identical to those obtained with Novec 7000: the maximum winding temperature is 150.21 °C, corresponding to a difference of only 0.3 °C (+0.2%) compared with Novec 7000, while the outlet vapor fraction is 45.33%, a value almost identical to that obtained with Novec 7000 (45.2%).

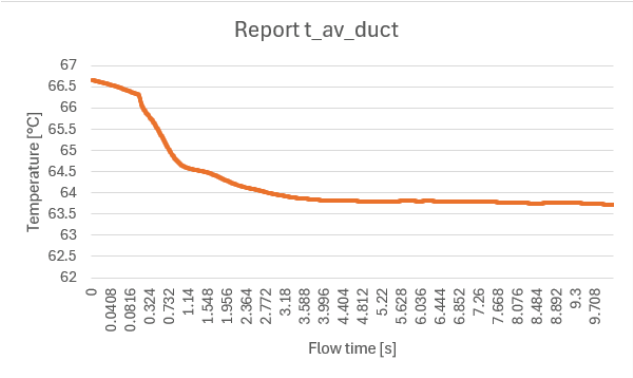
The temperature distribution in the other components (lamination and plastic insulation) also follows the same trend observed for Novec 7000, with deviations consistently below 1%, further confirming the close similarity between the two fluids in terms of thermofluid-dynamic behavior. The thermal balance exhibits a relative imbalance of 3.40%, higher than that obtained for Novec 7000 but still within a range considered acceptable for this type of transient multiphase simulation.

Table 35 – Opteon SF-33 Two-Phase Simulation Results (0.0035 kg/s per Channel)

Parameter	**Value**
T_{out} **[°C]**	61.54
$T_{av\ duct}$ **[°C]**	63.72
$T_{max\ duct}$ **[°C]**	118.66
$T_{av\ w}$ **[°C]**	133.35
$T_{max\ w}$ **[°C]**	150.21
$T_{av\ lamination}$ **[°C]**	133.43
$T_{max\ lamination}$ **[°C]**	139.2
$T_{av\ plastic}$ **[°C]**	87.03
$T_{max\ plastic}$ **[°C]**	127.66
p_{drop} **[Pa]**	60.78
Liquid fraction [%]	62.8
Vapor fraction [%]	37.19
Vapor fraction outlet [%]	45.33

Table 36 – Energy Balance of the Opteon SF-33 Two-Phase Simulation (0.0035 kg/s per Channel)

Parameter	**Value**
User source	145.42
$Q_{inlet\ 1}$ **[W]**	104.99
$Q_{inlet\ 2}$ **[W]**	104.99
$Q_{outlet\ 1}$ **[W]**	-162.6
$Q_{outlet\ 2}$ **[W]**	-161.2
$Q_{conv\ plastic}$ **[W]**	-3.594
$Q_{conv\ lamination}$ **[W]**	-32.957
Net results [W]	-4.946



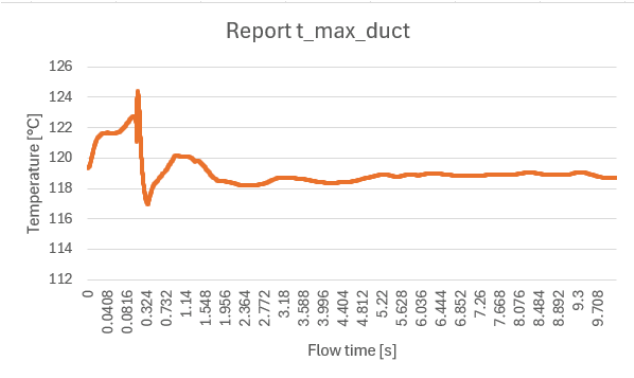
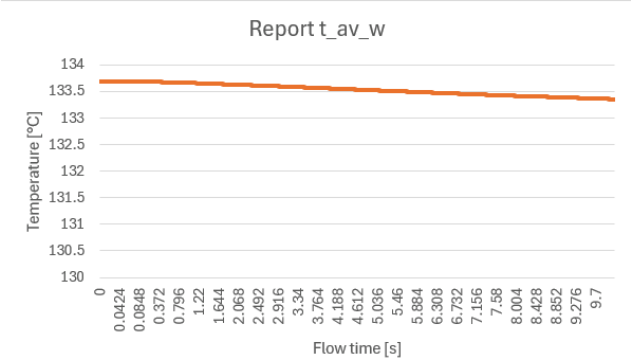


Figure 114 – Average Cooling-Channel Temperature Report, Opteon SF-33
Figure 115 – Maximum Cooling-Channel Temperature Report, Opteon SF-33



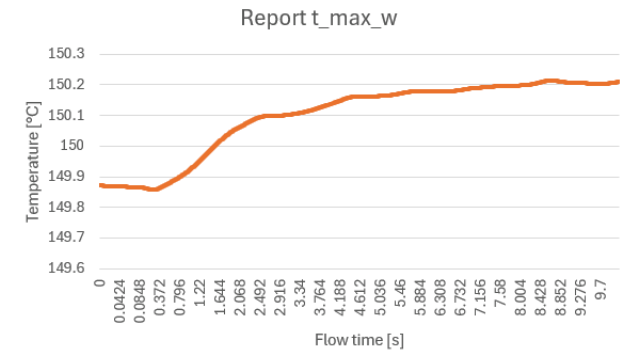
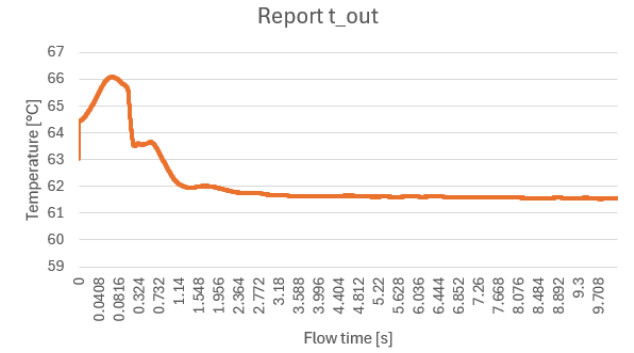


Figure 116 – Average Winding Temperature Report, Opteon SF-33
Figure 117 – Maximum Winding Temperature Report, Opteon SF-33



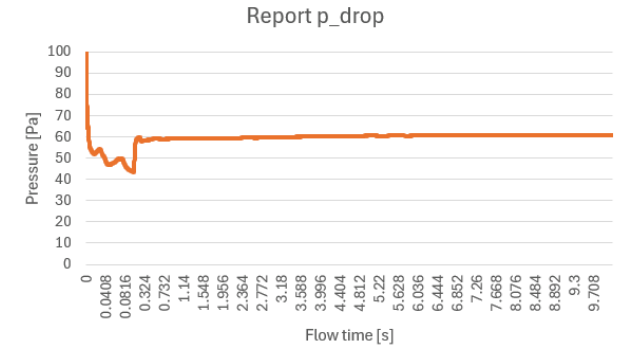
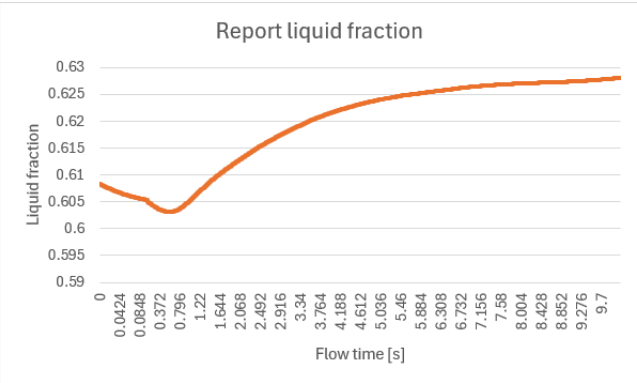


Figure 118 – Opteon SF-33 Outlet Temperature Report
Figure 119 – Pressure Drop Report, Opteon SF-33



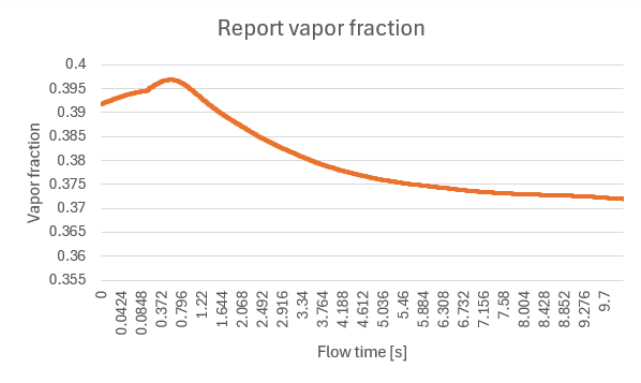


Figure 120 – Liquid Fraction Report, Opteon SF-33
Figure 121 – Vapor Fraction Report, Opteon SF-33

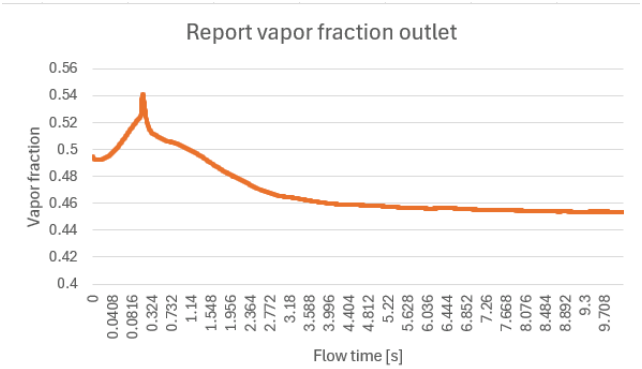
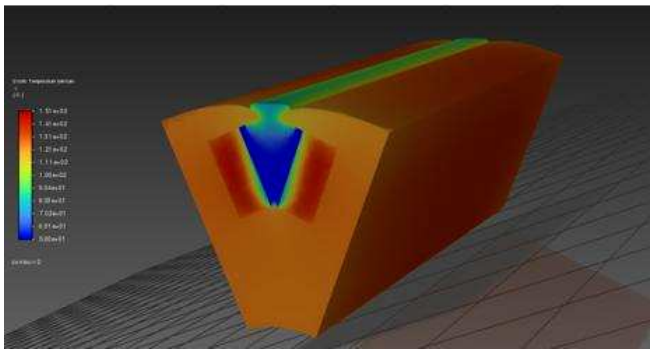


Figure 122 – Outlet Vapor Fraction Report, Opteon SF-33

Contour case Opteon SF-33



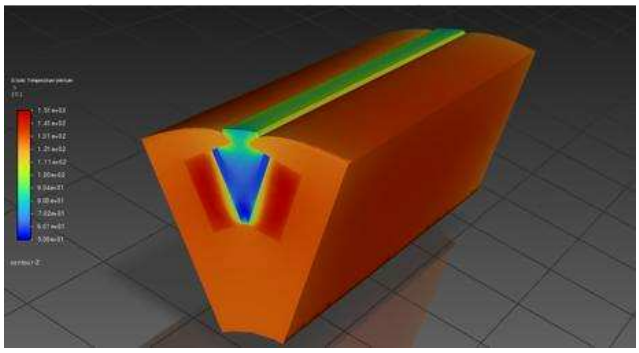


Figure 123 – Electric Rotor Computational-Domain Contour, Inlet Side, Isometric View
Figure 124 – Electric Rotor Computational-Domain Contour, Outlet Side, Isometric View

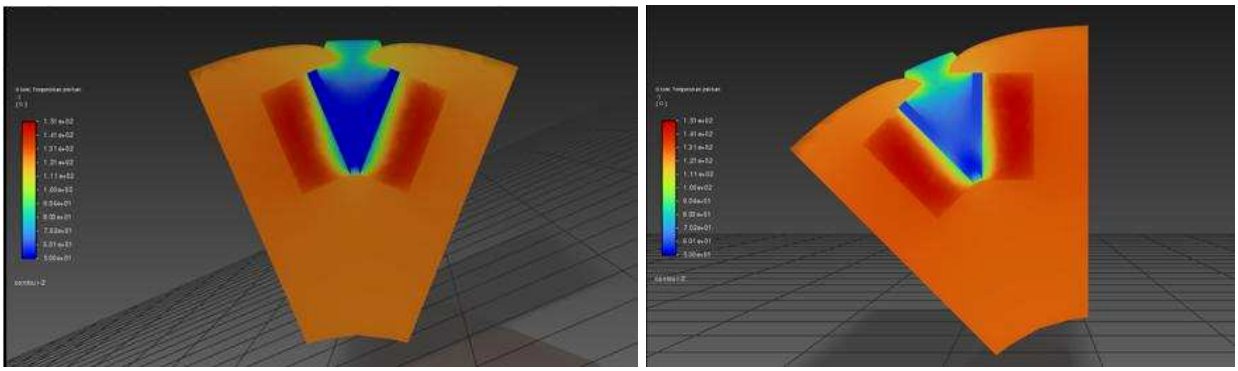


Figure 125 – Electric Rotor Computational-Domain Contour, Inlet Side, Front View
Figure 126 – Electric Rotor Computational-Domain Contour, Outlet Side, Front View

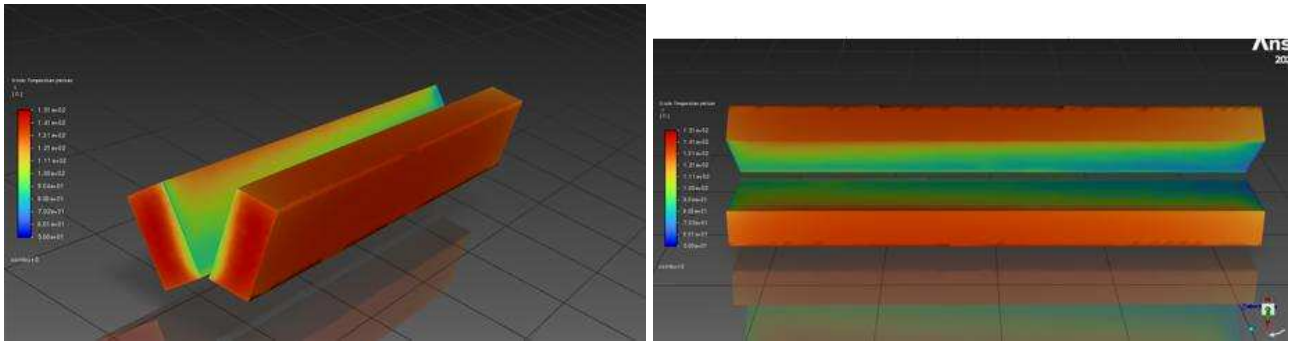


Figure 127 – Copper-Winding Contour, Isometric View
Figure 128 – Copper-Winding Contour, Top View

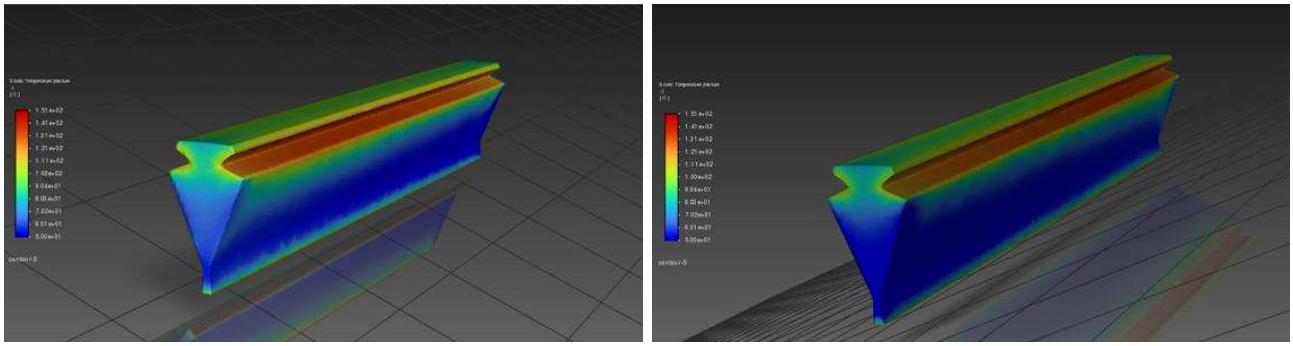


Figure 129 – Plastic-Partition Contour, Outlet Side, Isometric View
Figure 130 – Plastic-Partition Contour, Inlet Side, Isometric View

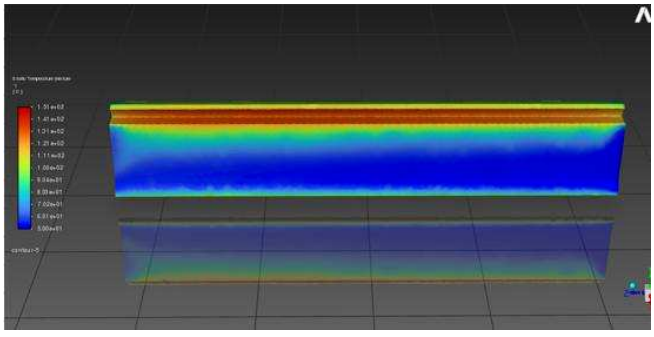
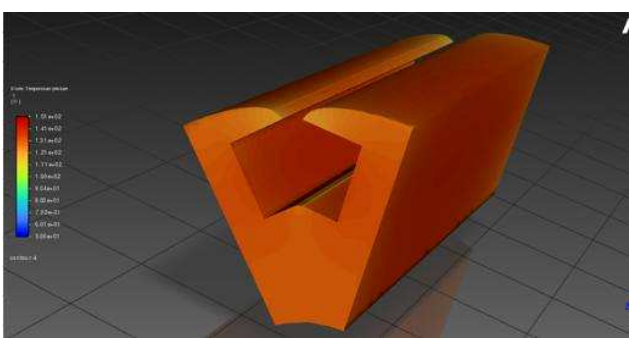


Figure 131 – Plastic-Partition Contour, Side View



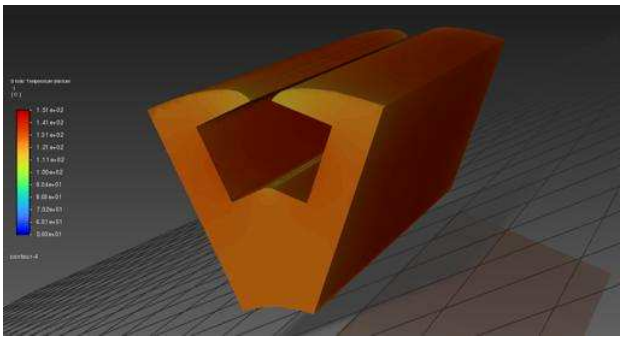
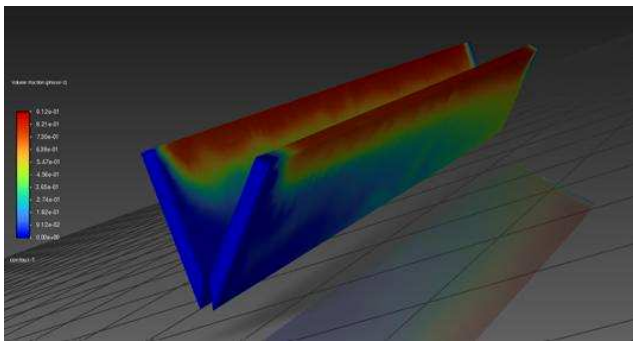


Figure 132 – Rotor Ferromagnetic-Structure Contour, Outlet Side
Figure 133 – Rotor Ferromagnetic-Structure Contour, Inlet Side



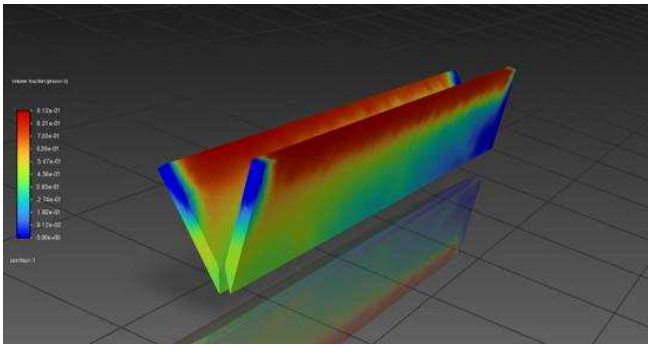


Figure 134 – Vapor-Phase Contour in the Cooling Channels, Inlet Side
Figure 135 – Vapor-Phase Contour in the Cooling Channels, Outlet Side

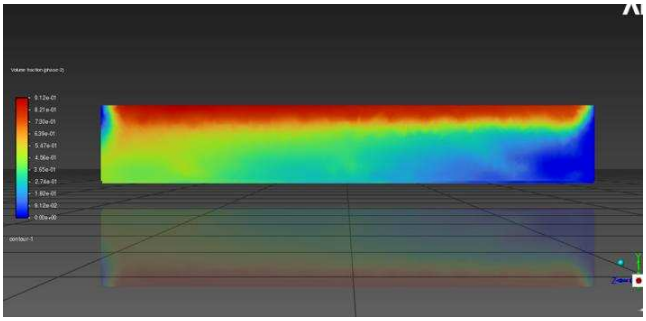


Figure 136 – Vapor-Phase Contour in the Cooling Channels, Side View

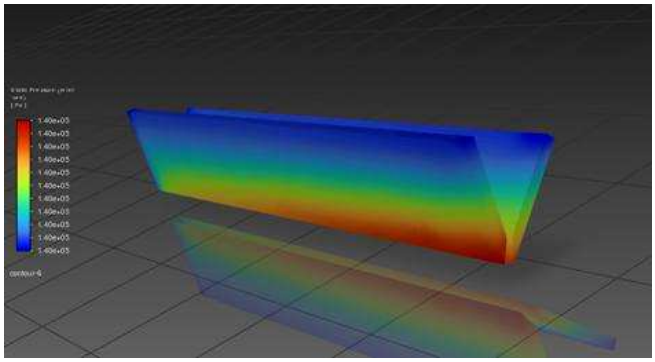


Figure 137 – Pressure-Drop Contour in the Cooling Channels, Inlet Side

6.7.4 R134a

R134a differs markedly from the other three fluids due to the high saturation pressure required to ensure T_{sat} = 60 °C (1.68 MPa, compared with a range of 113.5–240 kPa for the other three refrigerants), which is a direct consequence of the particularly steep vapor-pressure curve typical of low-molecular-weight hydrofluorocarbons (102.03 kg/kmol, the lowest among the four fluids). At this operating pressure, the outlet vapor fraction is the lowest among the four cases (12.53%), indicating an operating regime in which the contribution of latent heat transfer is more limited than for the other fluids and in which the refrigerant remains, for almost the entire flow path, in the subcooled-liquid state or only slightly superheated relative to saturation conditions.

Nevertheless, R134a provides the best thermal performance among all the fluids considered in the comparison: the maximum winding temperature is 136.65 °C, the lowest value recorded among the four fluids (−8.8% compared with Novec 7000), while the relative imbalance of the thermal balance is also the lowest (0.12%), confirming the excellent convergence quality of the simulation. This is likely facilitated by the more limited extent of the two-phase region and the resulting lower complexity of the phase-change front to be resolved numerically. This result can be attributed to the liquid-phase properties of R134a, particularly its high thermal conductivity (0.081 W/(m·K), the highest among the four fluids) and high specific heat capacity (1425 J/(kg·K), also the highest among the four fluids), which promote effective heat removal through both sensible and latent heat transfer, more than compensating for the lower evaporated fraction compared with the other refrigerants.

Table 37 – R134a Two-Phase Simulation Results (0.0035 kg/s per Channel)

Parameter	Value
T_{out} [°C]	60.25
$T_{av\ duct}$ [°C]	60.33
$T_{max\ duct}$ [°C]	99.94
$T_{av\ w}$ [°C]	120.67
$T_{max\ w}$ [°C]	136.65
$T_{av\ lamination}$ [°C]	120.17
$T_{max\ lamination}$ [°C]	125.56
$T_{av\ plastic}$ [°C]	76.92
$T_{max\ plastic}$ [°C]	112.29
p_{drop} [Pa]	64.3
Liquid fraction [%]	89.21
Vapor fraction [%]	10.78
Vapor fraction outlet [%]	12.53

Table 38 – Energy Balance of the R134a Two-Phase Simulation (0.0035 kg/s per Channel)

Parameter	Value
User source	145.42
$Q_{inlet\ 1}$ **[W]**	124.68
$Q_{inlet\ 2}$ **[W]**	124.68
$Q_{outlet\ 1}$ **[W]**	-183.84
$Q_{outlet\ 2}$ **[W]**	-181.1
$Q_{conv\ plastic}$ **[W]**	-2.84
$Q_{conv\ lamination}$ **[W]**	-27.17
Net results [W]	-0.181

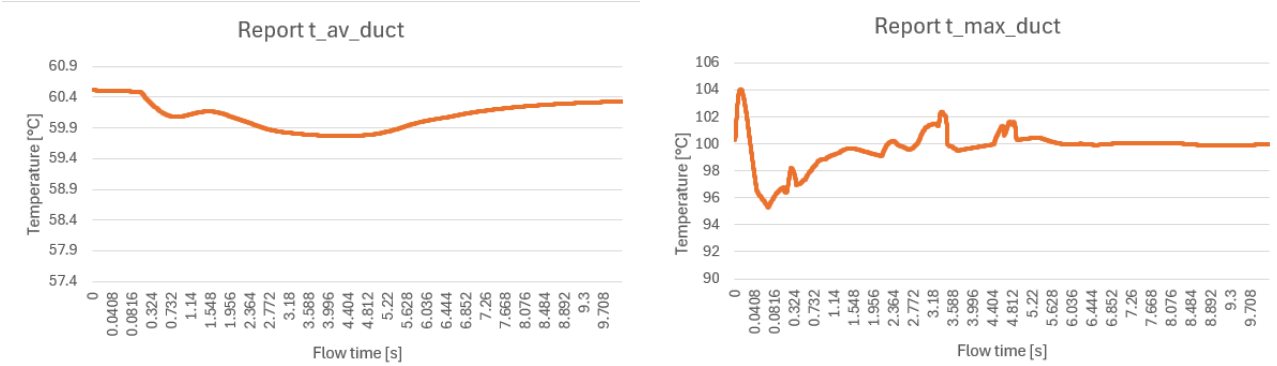


Figure 138 – Average Cooling-Channel Temperature Report, R134a
Figure 139 – Maximum Cooling-Channel Temperature Report, R134a

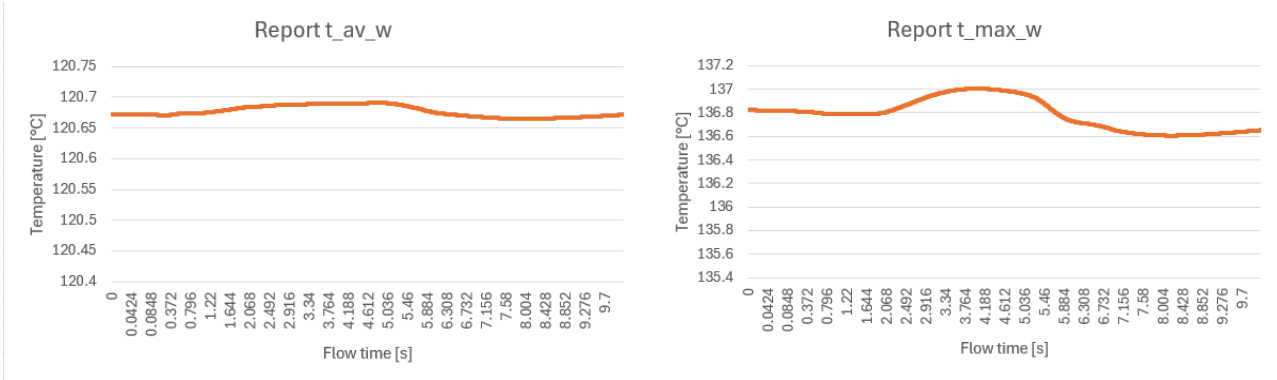
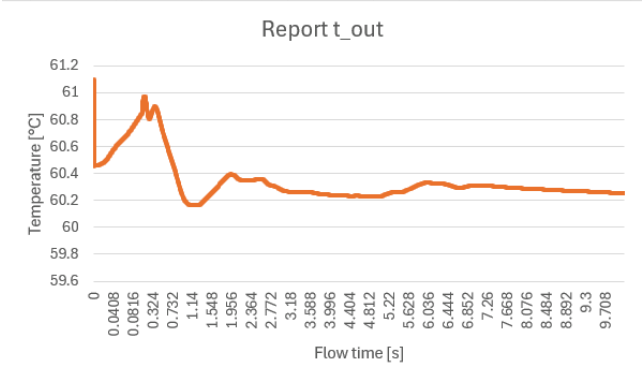


Figure 140 – Average Winding Temperature Report, R134a
Figure 141 – Maximum Winding Temperature Report, R134a



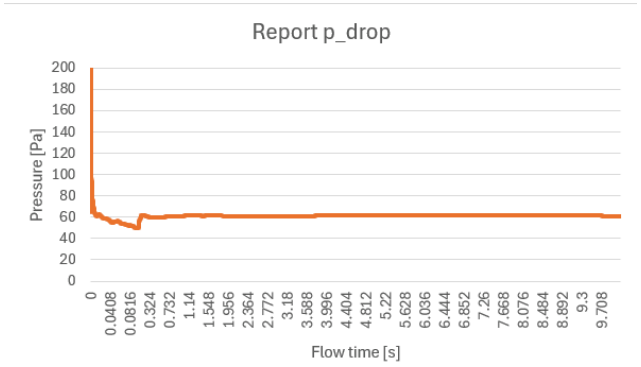
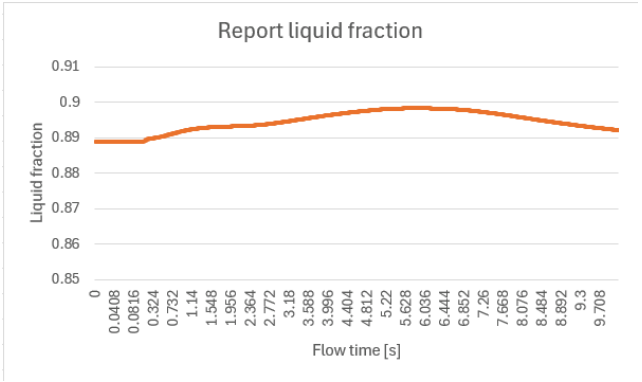


Figure 142 – R134a Outlet Temperature Report
Figure 143 – Pressure Drop Report, R134a



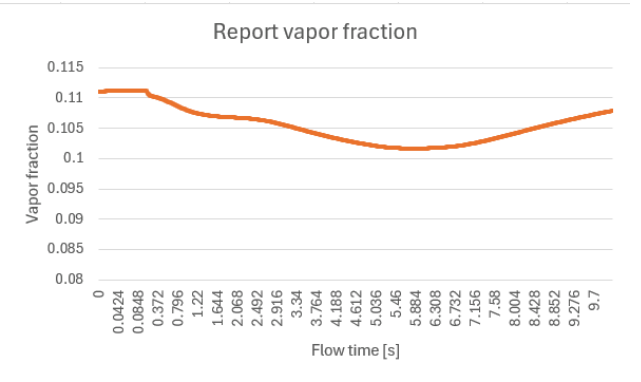


Figure 144 – Liquid Fraction Report, R134a
Figure 145 – Vapor Fraction Report, R134a

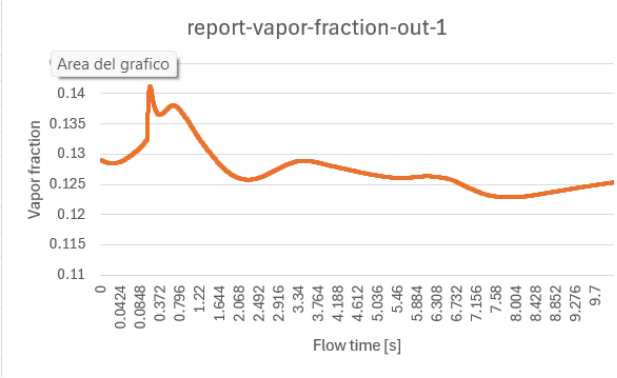
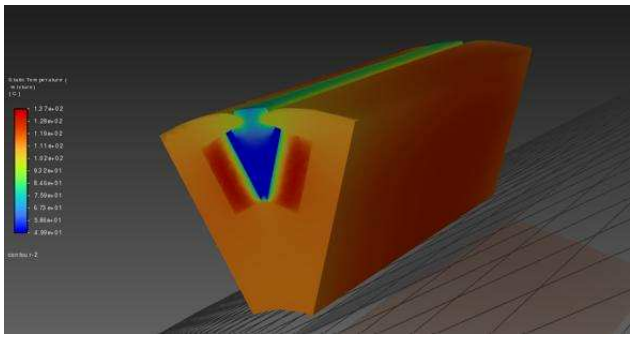


Figure 146 – Outlet Vapor Fraction Report, R134a

Contour case R134a



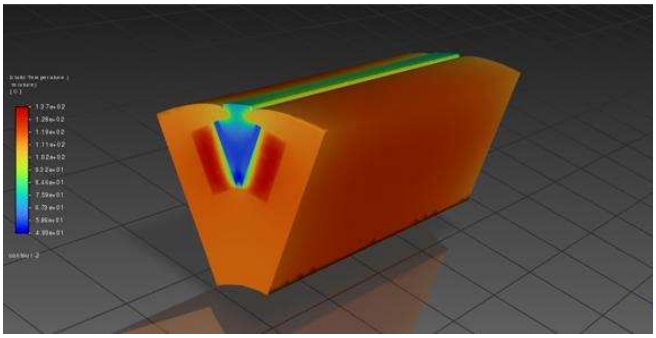
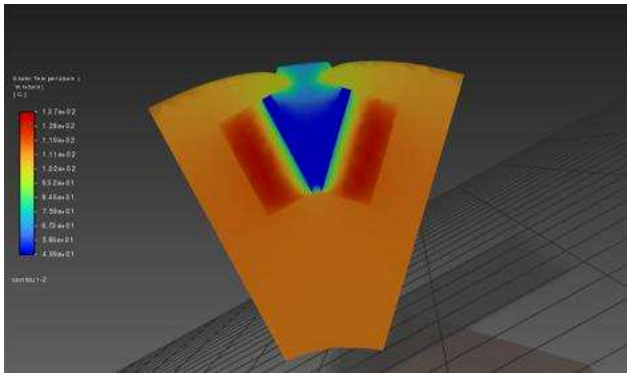


Figure 147 – Electric Rotor Computational-Domain Contour, Inlet Side, Isometric View
Figure 148 – Electric Rotor Computational-Domain Contour, Outlet Side, Isometric View



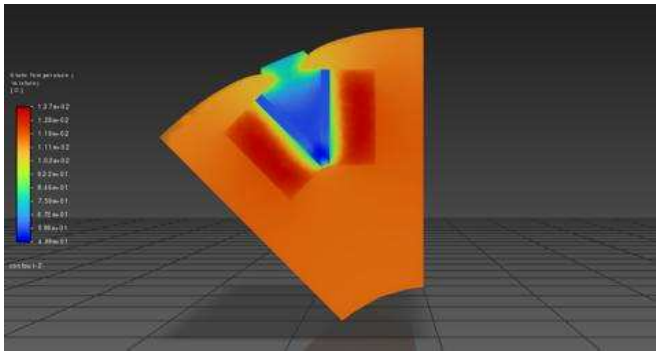
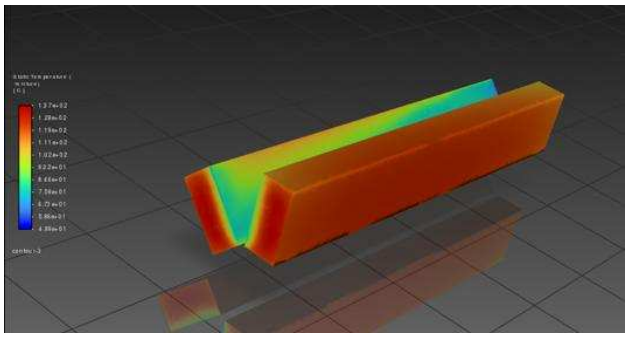


Figure 149 – Electric Rotor Computational-Domain Contour, Inlet Side, Front View
Figure 150 – Electric Rotor Computational-Domain Contour, Outlet Side, Front View



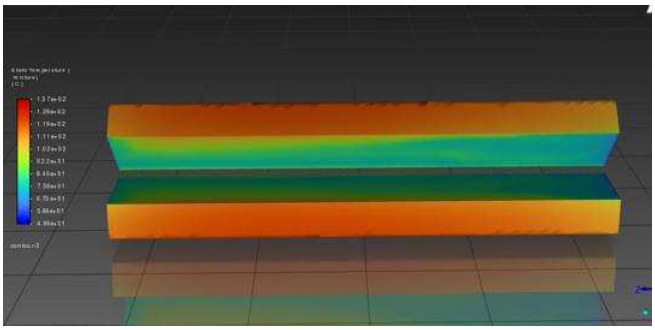


Figure 151 – Copper-Winding Contour, Isometric View
Figure 152 – Copper-Winding Contour, Top View

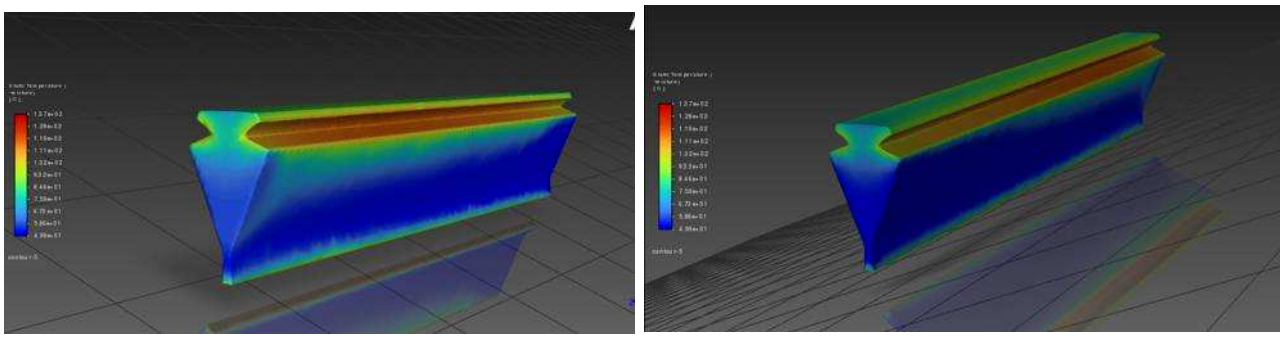


Figure 153 – Plastic-Partition Contour, Outlet Side, Isometric View
Figure 154 – Plastic-Partition Contour, Inlet Side, Isometric View

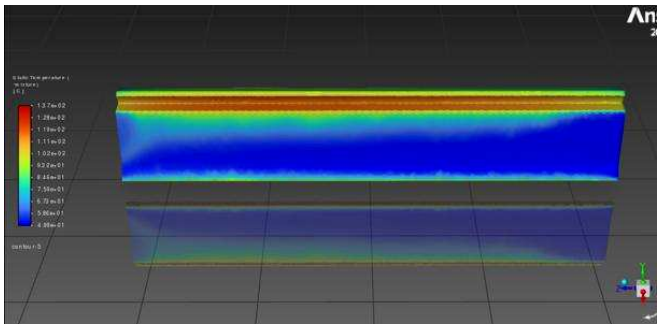


Figure 155 – Plastic-Partition Contour, Side View

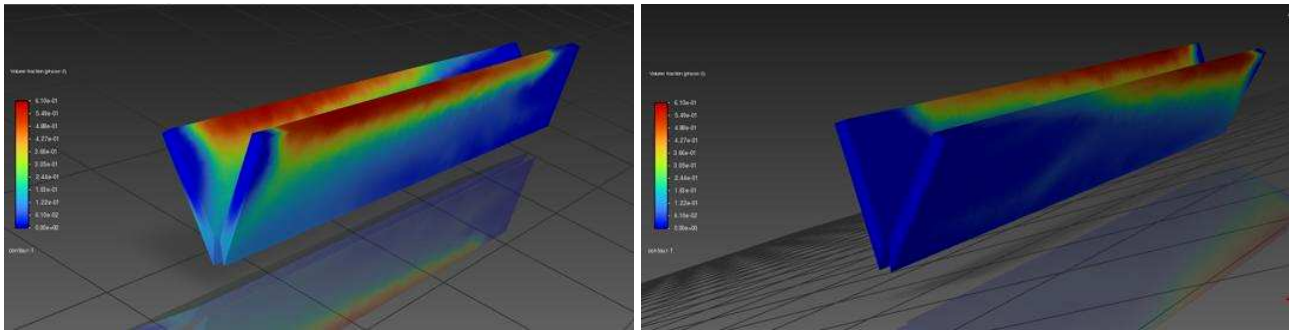


Figure 156 – Vapor-Phase Contour in the Cooling Channels, Outlet Side
Figure 157 – Vapor-Phase Contour in the Cooling Channels, Inlet Side

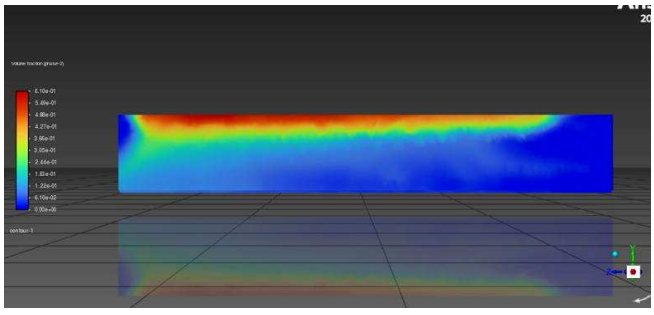


Figure 158 – Vapor-Phase Contour in the Cooling Channels, Side View

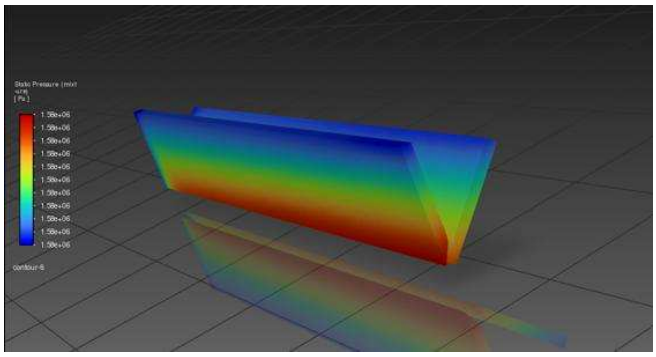


Figure 159 – Pressure-Drop Contour in the Cooling Channels, Inlet Side

6.8 Comparative Assessment of the Four Fluids and the Single-Phase Case

Table 39 – Comparison of the Properties and Two-Phase Simulation Results of the Four Fluids

Parameter	Novec 7000	FC-72	Opteon SF-33	R134a
Latent heat [kj/kg]	128	83.54	153	139.1
Liquid conductivity [W/m/K]	0.075	0.057	0.077	0.081
Liquid specific heat [J/kg/K]	1300	1100	1200	1425
Liquid density [kg/m^3]	1400	1680	1360	1207
Saturation pressure [Pa]	220000	113500	240000	1680000
T_{outlet} **[°C]**	60.99	62.4	61.54	60.25
$T_{max\ w}$ **[°C]**	149.91	166.52	150.21	136.65
$T_{max\ lamination}$ **[°C]**	138.61	155.28	139.2	125.56
$T_{max\ plastic}$ **[°C]**	126.33	141.6	127.66	112.29
Pressure losses [Pa]	62.23	62.72	60.78	64.3
Vapor fraction [%]	37.72	50.27	37.19	10.78
Vapor fraction outlet [%]	45.2	59.96	45.33	12.53
Imbalance [%]	1.75	3.74	3.4	0.12

It is interesting to observe that this ranking cannot be directly attributed solely to the vapor fraction produced, demonstrating the multifactorial nature of two-phase heat transfer. Indeed, FC-72

undergoes a greater degree of evaporation than Novec 7000 and Opteon SF-33 (outlet vapor fraction of 59.96% compared with approximately 45%), yet its heat-transfer performance remains significantly poorer due to its low latent heat of vaporization (83.54 kJ/kg, the lowest among the four fluids) and relatively low thermal conductivity in both the liquid and vapor phases. In particular, thermal conductivity governs the conductive thermal resistance of the thin liquid film separating the heated wall from the evaporating fluid core, and a reduction in this property directly results in greater wall superheat for the same amount of heat flux removed.

The case of R134a is even more significant from a physical perspective. Despite exhibiting the lowest vapor fraction among the four fluids (12.53%), indicating operation predominantly in the liquid phase, it achieves the lowest temperatures owing to the high thermal conductivity and high specific heat capacity of its liquid phase, which more than compensate for the lower contribution of latent heat. This result suggests that, for non-negligible refrigerant mass flow rates such as that considered here, the sensible contribution to heat transfer may be comparable to, or even greater than, the latent contribution, particularly for fluids with high heat capacity such as R134a (liquid-phase c_p equal to 1425 J/(kg·K), approximately 30% higher than that of Novec 7000). It should, however, be noted that this performance is achieved at the expense of a significantly higher saturation pressure (1.68 MPa compared with a range of 113.5–240 kPa for the other three fluids), resulting in more stringent sealing and mechanical-strength requirements for the components of the cooling circuit (pipes, joints, and the rotating seals required to supply a rotating component such as the rotor investigated in the present study), with a non-negligible impact on cost, system complexity, and potentially on system safety.

In terms of pressure drop, the four fluids exhibit substantially equivalent behavior, with values ranging from 60.78 Pa (Opteon SF-33) to 64.30 Pa (R134a), corresponding to a variation of approximately 5.8% between the minimum and maximum values. Refrigerant selection therefore does not result in significant fluid-dynamic penalties, and the sizing of the pumping unit of the cooling circuit is essentially independent of the selected refrigerant. Regarding the quality of the energy balance, R134a exhibits the lowest relative imbalance (0.12%), followed by Novec 7000 (1.75%), whereas Opteon SF-33 and FC-72 exhibit higher imbalances (3.40% and 3.74%, respectively), while still remaining within limits considered acceptable for this type of multiphase simulation. This ranking in terms of convergence quality appears to be correlated, although not in a strictly monotonic manner, with the extent of the two-phase region within the channel, which is more limited for R134a and more extensive for FC-72.

In addition to purely thermofluid-dynamic considerations, the selection of the most suitable refrigerant for the application should also take into account environmental and safety considerations, although these aspects fall outside the specific objectives of the present numerical analysis. In general terms, Novec 7000 is marketed by 3M as a non-flammable fluid with a lower global warming potential (GWP) than conventional hydrofluorocarbons, while also possessing dielectric properties that allow it to be used in direct contact with energized electrical components. R134a, although widely used and representing a mature solution from a system-design perspective, typically exhibits a higher GWP, characteristic of conventional hydrofluorocarbon refrigerants. FC-72, as a perfluorocarbon, is characterized by high environmental persistence. Although these aspects do not directly affect the thermal performance discussed in this section, they represent an additional criterion to be considered in an overall assessment and could be investigated in greater detail as part of future developments of the present thesis work.

Table 40 – Single-Phase Simulation Results (Refrigerant Inlet Temperature of 60 °C)

Parameter	Value
Mass flow rate [g/s] (per condotto)	12.81
T_{inlet} **[°C]**	60
T_{outlet} **[°C]**	61.92
$T_{max\,w}$ **[°C]**	138.42
$T_{av\,w}$ **[°C]**	122.2
$T_{max\,lamination}$ **[°C]**	126.31
$T_{max\,plastic}$ **[°C]**	113.15
Pressure losses [Pa]	898.26
Net results [W]	1.559

The comparison shows that the single-phase case, despite reaching a maximum winding temperature of 138.42 °C, which is comparable to that achieved by the best-performing two-phase fluids (Novec 7000, Opteon SF-33, and, in particular, R134a), requires a total mass flow rate of 205 g/s, approximately 3.7 times higher than that employed in the two-phase cases (56 g/s), together with a pressure drop of 898.26 Pa, more than one order of magnitude higher than the 60–65 Pa observed in the two-phase simulations. This comparison quantitatively highlights the main advantage of the adopted two-phase cooling strategy: by exploiting the latent heat of vaporization rather than relying solely on sensible heat, it is possible to achieve thermal performance comparable to, or in some cases better than (R134a), that of a single-phase system, while drastically reducing both the required refrigerant mass flow rate and the pressure drop, and consequently the pumping power required to operate the cooling circuit. This also provides clear benefits in terms of the size of the refrigerant supply lines and the energy consumption of the system auxiliaries.

In conclusion, considering jointly the thermal performance, the quality of numerical convergence, the system complexity associated with the operating pressure, and, on a preliminary basis, the environmental aspects discussed above, Novec 7000 and Opteon SF-33 represent the most balanced candidates among the four fluids analyzed, providing robust thermal performance at moderate saturation pressures while offering favorable properties in terms of electrical safety and environmental impact. R134a represents the option providing the best overall thermal performance, although at the expense of a significantly higher operating pressure, which complicates its integration into a cooling system for rotating components. Finally, FC-72 appears to be the least suitable candidate for the application considered, owing to its poorer thermal performance compared with the other three fluids at the same mass flow rate, further compounded by its greater environmental persistence.

Chapter 7

Discussion of the Results

This chapter provides an overall interpretation of the results obtained in Chapter 6, with the aim of moving beyond the direct comparison of individual simulations and integrating the numerical findings into a unified interpretative framework. This approach is intended both to support the selection of the refrigerant most suitable for the application under consideration and to highlight the inherent limitations of the adopted numerical methodology. The discussion is divided into three sections: the first (section 7.1) systematically compares the overall performance of the different cases analyzed, introducing summary metrics that enable the comparison of fluids and operating regimes characterized by different conditions; the second (section 7.2) translates these findings into design implications, discussing the main trade-offs to be considered in refrigerant selection and cooling-system sizing; and the third (section 7.3) critically examines the limitations of the numerical model employed and the main sources of uncertainty affecting its quantitative reliability.

7.1 Overall Performance Comparison

Although informative, a direct comparison of the maximum temperatures obtained in the different cases does not, by itself, provide a complete assessment of the relative effectiveness of the different cooling solutions, since each case operates under partially different boundary conditions (fluid inlet temperature of 50 °C for the two-phase cases and 60 °C for the reference single-phase case) and relies on heat-transfer mechanisms of different nature (latent heat transfer for the two-phase fluids and predominantly sensible heat transfer for the single-phase case). In order to provide a consistent basis for comparison, an equivalent thermal resistance R_{th} is introduced, defined as the ratio between the winding temperature rise relative to the common saturation temperature T_{sat} = 60 °C (also adopted as the reference inlet temperature for the single-phase case) and the total thermal power generated within the computational domain:

$$R_{th} = \frac{\left(T_{\max winding} - 60°C\right)}{Q_{gen}} \qquad con\ Q_{gen} = 145.42\ W$$

This quantity, expressed in °C/W, conceptually represents the overall thermal resistance experienced by the heat source (the windings) relative to a common temperature reference and enables a direct comparison of cooling effectiveness independently of the dominant heat-transfer mechanism. The values obtained for the five cases analyzed are reported in the table below.

Table 41 – Comparison between Two-Phase Refrigerants and Single-Phase Oil

Paremeter	Novec 7000	FC-72	Opteon SF-33	R134a	Ampcol oil (inlet 60°C)
$T_{max\,w}$ [°C]	149.91	166.52	150.21	136.65	138.42
R_{th} [°C/W]	0.618	0.732	0.62	0.527	0.539
Mass flow rate [g/s] (single channel)	3.5	3.5	3.5	3.5	12.81
Pressure losses [Pa]	62.23	62.72	60.78	64.3	898.26
Absolute pressure [Pa]	220000	113500	240000	1680000	101325
Imbalance [%]	1.75	3.74	3.4	0.12	1.07

Parameters analyzed for the comparison:

- **Cooling effectiveness:** the equivalent thermal resistance confirms the ranking already discussed in Section 6.5, while providing a more rigorous quantitative assessment of the differences among the various cases. R134a exhibits the lowest thermal resistance (0.527 °C/W), followed closely by the reference single-phase case (0.539 °C/W, +2.3% compared with R134a), whereas Novec 7000 and Opteon SF-33 fall within an almost identical intermediate range (0.618 and 0.620 °C/W, respectively, +17.3% and +17.6% compared with R134a), while FC-72 ranks last, with a significantly higher thermal resistance (0.732 °C/W, +38.9% compared with R134a). The most significant result emerging from this comparative analysis is the substantial equivalence, in terms of pure cooling effectiveness, between the single-phase case and the best-performing two-phase candidate (R134a). Considered in isolation, this result could suggest the absence of a clear advantage of the two-phase cooling strategy over a conventional single-phase solution. However, such a conclusion is incomplete unless complemented by an analysis of the pressure penalty, discussed below, which represents the actual distinguishing factor between the two cooling strategies.
- **Pressure penalty:** it is appropriate to conceptually distinguish between two types of pressure penalty, which are different in physical nature and have distinct design implications: on the one hand, the distributed pressure drop along the channel, resulting from viscous friction and, in the two-phase case, also from the acceleration of the liquid-vapor mixture caused by the variation in average density along the flow path; on the other hand, the absolute operating pressure of the circuit, required to ensure the design saturation temperature (60 °C) for the fluid under consideration.
- **Regarding the distributed pressure drop**: the comparison among the four two-phase fluids shows substantially equivalent values, within a narrow range between 60.78 Pa (Opteon SF-33) and 64.30 Pa (R134a), with a maximum variation of less than 6%. Therefore, the choice of the two-phase fluid does not result in significant penalties in terms of distributed pressure drop. By contrast, the comparison with the single-phase case is clear: the pressure drop of 898.26 Pa is more than one order of magnitude higher (approximately 14 times) than the

values obtained for the two-phase cases, as a direct consequence of both the higher total mass flow rate required (205 g/s compared with 56 g/s, corresponding to a factor of 3.7) and the absence of the latent cooling contribution, which forces the single-phase system to operate at a significantly higher Reynolds number in order to ensure an adequate convective heat-transfer coefficient. For the same pump efficiency, this difference results in a significantly higher pumping-power requirement for the single-phase circuit than for the two-phase circuits. This aspect must be properly taken into account in the overall energy balance of the cooling system, particularly in applications such as the automotive sector, where the energy efficiency of auxiliary systems directly affects vehicle driving range.

- **Regarding the absolute operating pressure**: by contrast, the variability among the four two-phase fluids is very pronounced: it ranges from 113.5 kPa for FC-72, close to atmospheric pressure, to 220–240 kPa for Novec 7000 and Opteon SF-33, and up to 1.68 MPa for R134a, almost one order of magnitude higher. This second pressure penalty, which is static rather than distributed in nature, represents the actual distinguishing factor among the two-phase fluids from a system-design perspective, since it directly affects the required circuit wall thickness, the type of seals required (particularly for the rotating coupling needed to supply the rotor), and, more generally, the safety class of the entire cooling system.
- **Practicality of implementation:** from the standpoint of practical implementation, the single-phase case offers the advantage of operating at a pressure close to atmospheric pressure and of not requiring a dedicated refrigeration circuit (compressor, condenser, and expansion valve), since it can rely on a simple hydraulic circuit consisting of a pump and a heat exchanger. However, this system-level simplicity is counterbalanced by the need to handle a significantly higher coolant mass flow rate and pressure drop, resulting in larger pumping components and supply lines.
- The two-phase fluids Novec 7000 and Opteon SF-33, operating at moderate pressures of approximately 2–2.5 bar absolute, make it possible to limit the required circuit wall thickness and sealing complexity, although they still require the integration of a complete refrigeration circuit to manage the phase-change process (evaporation on the rotor side and condensation on the external side). R134a, despite providing the best thermal performance, requires an operating pressure of almost 17 bar absolute, a condition that necessitates the use of components with higher pressure ratings, dedicated rotating seals for supplying the rotor, and, more generally, a greater degree of attention from a system-safety perspective. This aspect is particularly critical in an automotive application, where the cooling system is integrated into a high-speed rotating component. FC-72, despite exhibiting the lowest operating pressure among the four two-phase fluids, is penalized by its poorer thermal performance, which would require an increase in mass flow rate to achieve thermal margins comparable to those of the other fluids, thereby partially offsetting the practical advantage associated with its low operating pressure.

7.2 Design Implications

The findings presented in section 7.1 make it possible to define a structured decision-making framework for selecting the refrigerant most suitable for the application under consideration, based on three main criteria that are often in conflict with one another: thermal performance (summarized by the equivalent thermal resistance R_{th}), system complexity associated with the required operating pressure, and the safety and environmental sustainability aspects of the fluid employed.

Based on this framework, Novec 7000 and Opteon SF-33 emerge as the most balanced candidates for the application investigated in the present study. Both fluids provide a sufficiently large thermal margin with respect to the thermal-class limits of the electrical insulation (maximum winding temperatures of 149.91 °C and 150.21 °C, respectively, compared with the Class H limit of 180 °C), operate at moderate pressures of approximately 2–2.5 bar absolute, which are compatible with relatively simple system configurations and standard rotating seals, and, according to the technical documentation provided by their respective manufacturers, exhibit favorable properties in terms of non-flammability and reduced environmental impact compared with conventional hydrofluorocarbon refrigerants. The trade-off associated with selecting these two fluids consists of an equivalent thermal resistance approximately 17% higher than that of the best-performing candidate overall (R134a). Considering that the thermal limits of the electrical insulation are still comfortably satisfied in both cases, this penalty appears fully acceptable in view of the benefits in terms of system simplicity and safety.

R134a represents the preferred option when the primary design criterion is the maximization of the thermal margin, for example in applications characterized by particularly demanding load cycles or by packaging constraints requiring a further reduction in the cross-sectional area of the cooling channels. In this case, however, the trade-off to be accepted is the adoption of a high-pressure circuit (approximately 17 bar absolute), which requires components with higher pressure ratings, dedicated rotating seals for supplying refrigerant to the rotating rotor, and a greater level of design attention in terms of safety, with a likely increase in system manufacturing and certification costs.

Finally, FC-72 does not appear competitive with the other three fluids for the application under consideration. Despite its relatively low operating pressure, its poorer thermal performance (equivalent thermal resistance 39% higher than that of R134a) would reduce the available thermal safety margin and, for the same geometry, would require an increase in refrigerant mass flow rate, thereby reducing the advantage in terms of refrigerant consumption typically associated with the two-phase cooling strategy. In addition, the greater environmental persistence typically associated with perfluorocarbons represents a further unfavorable factor in the overall assessment.

In addition to refrigerant selection, the analysis conducted suggests several broader design implications that should be taken into account during subsequent stages of cooling-system development. First, the optimal mass flow rate identified in section 6.3.5 (0.0035 kg/s per channel) was determined with reference to a single motor operating point (4200 rpm, half load, continuous operation). Extending the cooling system to the entire operating range of the electric motor would require the assessment of a mass-flow-rate control strategy, for example through the use of a variable-flow pump capable of adapting the refrigerant mass flow rate to the instantaneous thermal load while maintaining the outlet vapor fraction within the identified optimal range (approximately 40–50%), even under partial-load or transient-overload conditions. Second, the transition from the reduced model, representing one-eighth of the rotor and including only two cooling channels, to the complete geometry requires the design of a distribution manifold capable of ensuring a uniform mass-flow-rate distribution among all the rotor cooling channels. Indeed, any flow maldistribution, which is

particularly critical in a two-phase system, could generate localized hot spots not predicted by the present analysis, which was conducted under the assumption of perfect domain symmetry. Finally, the selection of the refrigerant should also be verified in terms of its chemical compatibility with the polymeric materials used in the rotor plastic partition, an aspect that was not investigated in the present numerical analysis but may be relevant to the long-term durability of the system.

7.3 Limitations of the Numerical Model and Potential Sources of Uncertainty

Although the results presented in the previous chapters provide consistent and physically justified design insights, they were obtained using a numerical model characterized by several simplifications, whose influence on the quantitative accuracy of the results deserves critical discussion. The main sources of uncertainty identified are presented below, grouped by category.

Simplifications of the Geometry and Boundary Conditions

- The computational domain was reduced to one-eighth of the complete rotor, including only two cooling channels, by exploiting the cyclic symmetry conditions of the component. Although this approach is justified by the geometric periodicity of the rotor, it assumes perfectly uniform coolant distribution and heat generation among all the actual channels, an assumption that may not be fully satisfied in the presence of manufacturing tolerances, machining non-uniformities, or edge effects that are not represented in the reduced model;
- Heat generation within the windings was imposed as uniform along the axial length, after appropriately scaling the overall Joule losses (from 88.75 W to 72.7 W for each half-winding) to account for the actual concentration of losses in the end-winding regions. Although necessary in the absence of a dedicated electromagnetic model, this approximation does not reproduce the actual non-uniform distribution of dissipated power density along the winding, which could generate localized hot spots that are not captured by the present model;
- The convective heat-transfer coefficient imposed on the external surfaces of the lamination and plastic insulation (h_{conv} = 40 W/(m²·K), with a $T_{free\ stream}$ of 50 °C) was assumed to be constant and uniform, in the absence of a dedicated analysis of the heat-transfer conditions between the rotor and its external surroundings (air gap, motor housing, and any possible ventilation). In reality, these conditions vary as a function of rotational speed, air-gap geometry, and motor operating conditions;
- The model does not account for the dynamic effects associated with rotor rotation (centrifugal and Coriolis forces), which, under actual operating conditions, could induce stratification of the liquid and vapor phases within the rotating channel, with the denser liquid phase being driven toward the outer wall of the channel by centrifugal forces. This phenomenon, which cannot be reproduced in a stationary domain such as that employed in the present study, could locally alter the heat-transfer coefficient and temperature distribution compared with the predictions of the present analysis, which was conducted under non-rotating conditions.

Phase-Change Modelling and Numerical Schemes

- The Lee phase-change model, employed in conjunction with the Mixture model for multiphase simulation, requires the definition of evaporation and condensation coefficients (the so-called "Lee coefficients"), which were set equal to unity for all the fluids analyzed in the present study. These coefficients are known in the literature to significantly affect the simulated phase-change rate, and no fluid-specific calibration was performed in the present work, for example through comparison with experimental data or literature correlations for convective boiling (such as the Chen correlation). This therefore represents a potential source of discrepancy between the numerical results and the actual behavior of the fluids;
- The spatial discretization schemes adopted are not uniform for all the equations solved. While the convective term of the momentum equation was discretized using a second-order scheme (Second Order Upwind) starting from the parametric mass-flow-rate study, the volume-fraction equation was maintained at first order (First Order Upwind) for reasons of numerical stability, in accordance with common practice in multiphase simulations using the Mixture model. Although necessary, this choice introduces greater numerical diffusion in the resolution of the liquid-vapor interface, potentially resulting in a smoothing of the vapor-fraction gradients near the heated wall;
- The saturation temperature was imposed as a constant and uniform value (60 °C) throughout the entire domain, consistently with the saturation pressure assigned as a fluid property. In reality, the pressure drop along the channel induces a local variation in static pressure and, consequently, a corresponding variation in the local saturation temperature. This effect was neglected in the present study because of the relatively small pressure drops observed (on the order of 60–70 Pa for the two-phase cases), although it could be taken into account in a more refined modelling approach.

Material and Fluid Properties

- The thermophysical properties of the four two-phase refrigerants were assumed to remain constant with temperature and were obtained from charts, saturation curves, technical datasheets, and reference databases through graphical reading or interpolation. Although this procedure represents common practice during preliminary investigations, it introduces a degree of uncertainty into the numerical values employed, which inevitably affects the simulation results, particularly for properties such as the latent heat of vaporization and surface tension, which have a more direct influence on phase-change dynamics;
- Similarly, the properties of the solid materials included in the computational domain (copper, M250-35A lamination, and plastic insulation) were assumed to be constant, whereas some of these properties, particularly the thermal conductivity of copper and the lamination, exhibit a non-negligible temperature dependence over the temperature range explored in the simulations (up to more than 180 °C in the poorest-performing cases).

Grid- and Time-Step-Independence Verification

- The grid-independence study described in this chapter was conducted exclusively for the single-phase case using the cooling oil, at a single inlet-temperature condition (40 °C). This study was not repeated for the two-phase configurations, for which grid resolution near the heated wall may play an even more critical role, since the vapor-fraction and temperature gradients within the thin phase-transition region must be adequately resolved. Extending the

grid validated for the single-phase case to the two-phase simulations therefore represents an assumption that was not directly verified, although it can be considered reasonable as a first approximation;

- The overall duration of the transient stage of the two-phase simulations (approximately 10 s of physical time, divided into an initial stage with a smaller time step followed by a stage with a larger time step) was not accompanied by an explicit and systematic verification that a quasi-steady-state condition had been reached for all the quantities of interest in all the cases analyzed. Although the reported thermal-balance values indicate a good level of convergence for most cases, the possibility remains that some local quantities, particularly the vapor-fraction distribution, had not yet reached a fully stabilized condition by the end of the simulated time window;
- Although relatively limited, the relative thermal-balance imbalances (ranging from 0.12% to 4.01% of the generated thermal power, depending on the case) nevertheless represent a non-zero residual in the closure of the overall energy balance and should therefore be regarded as an additional source of uncertainty to be considered when interpreting the presented comparisons quantitatively rather than qualitatively.

Lack of Experimental Validation

Finally, it should be emphasized that the entire analysis presented in this work is purely numerical, in the absence of direct experimental validation on the component under investigation or on a representative test configuration. Although the results obtained are consistent with the expected physical behavior of single-phase and two-phase heat-transfer phenomena, and although the observed qualitative trends (reduction in temperature with increasing mass flow rate and performance ranking among the fluids correlated with their respective thermophysical properties) appear robust and well supported from a physical standpoint, the absence of an experimental validation campaign requires caution when attributing to the absolute numerical values presented (temperatures, pressure drops, and vapor fractions) an accuracy greater than that typically associated with a model characterized by this level of simplification. Future experimental validation, even using a simplified test rig reproducing a single cooling channel, would represent a natural development of the present work and would help confirm, or potentially recalibrate, the design recommendations presented herein.

Chapter 8

Conclusions and Future Developments

This chapter concludes the thesis by summarizing, in section 8.1, the main results obtained throughout the modelling and numerical simulation process developed in the previous chapters, with particular attention to their internal consistency and overall engineering significance, and by outlining, in section 8.2, the possible extensions and future developments of the work, systematically identified on the basis of the limitations and simplifications critically discussed in Chapter 7. The purpose of this concluding chapter is not merely to repeat the numerical results already presented, but rather to integrate them into an overall engineering assessment of the feasibility and effectiveness of the two-phase cooling strategy investigated, while also outlining a credible pathway for the continuation of the research.

8.1 Summary of the Main Results

The present thesis investigated the numerical modelling of the cooling of the rotor of an Electrically Excited Synchronous Motor (EESM), with the aim of assessing the feasibility and performance of a two-phase cooling strategy based on exploiting the latent heat of vaporization of a refrigerant, as an alternative to a conventional single-phase forced-convection cooling solution. This objective arises from the increasingly stringent need in the automotive electric-traction sector to dissipate increasing thermal power densities within limited volumes, without excessively penalizing the vehicle auxiliary energy balance through oversized pumping systems. The study was conducted through CFD simulations in Ansys Fluent using a computational domain reduced to one-eighth of the complete rotor by exploiting the cyclic symmetry conditions of the component. The domain included two cooling channels, the corresponding copper windings, the ferromagnetic lamination, and the plastic insulation separating the fluid from the windings.

The adopted methodological approach can be summarized into six main stages, logically interconnected and each serving as a prerequisite for the subsequent stage, as summarized in the table below.

Table 42 – Summary of the Key Findings of the Thesis Study

Study Phase	Reference Configuration	Main Outcome
Grid-Independence Study	Single-phase case, AC-110 oil, $T_{in} = 40°C$, 3 meshes	Deviations < 2% among the meshes for $T_{\max w}$
Single-Phase Case	0.205 kg/s total, T_{in} 20-80°C	$T_{\max w}$ 104-154 °C, Δp 617-2937 Pa
Two-Phase Case, First Attempt	All four fluids, 0.568 g/s per channel	$T_{\max w}$ 189.9-248.7 °C, vapor fraction outlet 49.9-93.6 %
Mass-Flow-Rate Parametric Study	All four fluids, 1.25-4 g/s per channel	Mean $T_{\max w}$ 223.2→146.7 °C, $\dot{m}$ optimal 3.5 g/s
Comparison of Four Two-Phase Fluids	3.5 g/s per channel, T_{sat} 60°C	$T_{\max w}$ 136.7-166.5°C depending on the fluid
Comparative Discussion	Summary R_{th} , pressure penalty	Novec 7000 / Opteon SF-33 as the most balanced choice

Preliminary Stage: Grid-Independence Assessment and Characterization of the Single-Phase Case

Prior to the analysis of the two-phase system, it was necessary to establish an appropriate level of spatial discretization through a grid-independence study, conducted on the single-phase case using AmpCool AC-110 dielectric oil at an inlet temperature of 40 °C, by comparing three meshes with different resolutions (41 938, 38 833, and 129 488 cells, respectively). The comparison showed limited deviations among the three grids for the quantities of greatest design interest (maximum winding temperature ranging from 122.62 °C to 124.35 °C, corresponding to a relative deviation of less than 2%), making it possible to select a computational resolution capable of providing an appropriate compromise between solution accuracy and computational cost, which was subsequently employed for all the single-phase simulations presented in this work.

The single-phase case, analyzed by varying the fluid inlet temperature (20, 40, 60, and 80 °C, at a constant total mass flow rate of 0.205 kg/s), made it possible to characterize the thermofluid-dynamic behavior of the system under a purely convective sensible heat-transfer regime. The results highlighted an approximately linear relationship between inlet temperature and maximum winding temperature (from 104.08 °C to 154.18 °C over the investigated range) and, conversely, a strongly nonlinear decreasing relationship between inlet temperature and pressure drop (from 2936.91 Pa to 616.89 Pa), the latter being attributable to the pronounced temperature dependence of the oil viscosity. This characterization provided a physically well-established and internally consistent reference case against which the actual advantages of the two-phase cooling strategy discussed in Chapter 7 could subsequently be quantitatively assessed.

Two-Phase Refrigerant Mass-Flow-Rate Sizing

The transition from the single-phase to the two-phase case was initially addressed through a first-attempt estimate of the required mass flow rate, obtained from a simplified energy balance in which the entire generated thermal power (145.42 W) was assumed to be removed solely as latent heat of

vaporization. The latent heat of Novec 7000 (128 kJ/kg) was adopted as the reference value for this preliminary estimate, since it lies close to the mean latent heat of the four candidate fluids (125.9 kJ/kg) and therefore provides a balanced basis that is neither systematically favorable nor systematically penalizing for any of them. The resulting value of 0.568 g/s per channel was applied unchanged to all four refrigerants, consistently with the methodological requirement — discussed below — of exploring an identical mass-flow-rate grid for every candidate fluid.

The simulations performed at this first-attempt mass flow rate proved it to be excessively low for all four refrigerants, yielding clearly unacceptable maximum winding temperatures (230.96 °C for Novec 7000, 248.72 °C for FC-72, 223.00 °C for Opteon SF-33, and 189.94 °C for R134a) together with very high outlet vapor fractions for three of the four fluids (90.81%, 93.61%, and 87.74% for Novec 7000, FC-72, and Opteon SF-33 respectively, against 49.37% for R134a). These results indicate a regime in which almost the entire fluid mass evaporates well before reaching the channel outlet section, consequently leading to the development of a near-dry-out condition in the downstream portion of the channel, where the local heat-transfer coefficient collapses due to the absence of a liquid film at the wall. Far from representing a failure of the adopted approach, this outcome provided an important methodological insight: the sizing of the mass flow rate of a two-phase cooling system cannot be based solely on the overall energy balance, but must necessarily account for the local distribution of the vapor fraction along the channel and therefore requires iterative refinement supported by numerical simulation.

Based on this evidence, a systematic parametric study of the refrigerant mass flow rate was carried out on all four candidate fluids, exploring for each of them the same sequence of values per channel — 0.568, 1.25, 2.50, 3.00, 3.50, and 4.00 g/s — for a total of twenty-four two-phase simulations. The adoption of a mass-flow-rate grid common to all fluids responds to a specific methodological criterion: by imposing identical supply conditions, the only genuinely differentiating variable among the simulations remains the set of thermophysical properties of the refrigerant, so that the observed differences in thermal performance become directly attributable to the fluid characteristics alone. To ensure full comparability of the results, all twenty-four simulations were performed using the same multi-stage solution strategy, consisting of a steady-state initialization of 6000 iterations followed by a first transient stage with a time step of 0.0001 s for 1000 time steps and a second transient stage with a time step of 0.001 s for 10000 time steps, corresponding to an overall simulated physical time of approximately 10.1 s. The boundary conditions, the spatial discretization of the domain, and the numerical schemes were likewise kept unchanged throughout the entire study.

The optimal operating point was identified by applying the selection criteria to the mean values of the monitored quantities across the four candidate fluids, rather than to the results of any individual refrigerant. This ensures that the selected mass flow rate is not implicitly optimized for the characteristics of one particular fluid, but instead represents a balanced operating condition for the entire set of candidates, thereby providing a legitimate common basis for the subsequent comparison. The analysis of the marginal thermal benefit associated with each mass-flow-rate increment revealed the existence of two distinct regimes: below 2.50 g/s per channel each increment produced substantial reductions in the mean maximum winding temperature (38.3 °C and 25.0 °C for the first two steps respectively), reflecting the progressive departure from the near-dry-out conditions described above, whereas beyond this threshold each further increment of 0.50 g/s yielded a benefit of only 3.5 °C to 5.6 °C, indicating the onset of a region of markedly diminishing thermal returns.

Within this region, the operating point of 3.50 g/s per channel (56 g/s for the entire rotor) was selected as the optimum. The governing criterion proved to be the thermal margin of the least performing fluid: at 3.50 g/s all four refrigerants remain below the Class H insulation limit, with a margin ranging from 13.5 °C in the worst case (FC-72, 166.52 °C) to 43.4 °C in the best case (R134a, 136.65 °C), whereas at the immediately lower value of 3.00 g/s the margin of FC-72 would be reduced to only 10.5 °C, a value considered insufficient to guarantee robust operation in view of the model uncertainties discussed in Chapter 7. This choice is further supported by the outlet vapor fraction, which at 3.50 g/s settles at approximately 45% for both Novec 7000 and Opteon SF-33, thus preserving a balanced contribution between latent and sensible heat transfer, and by the fact that a further increase to 4.00 g/s would entail a 14.3% increase in the total refrigerant mass flow rate in exchange for a mean temperature reduction of only 4.2 °C.

Multi-Fluid Comparison and Decision-Making Summary

The results obtained at the optimal operating point were subsequently examined in detail in order to characterize the behavior of each refrigerant at the selected working condition. The comparison revealed a clear performance ranking, with R134a providing the lowest maximum winding temperature (136.65 °C), followed by Novec 7000 and Opteon SF-33, which exhibited virtually equivalent performance (149.91 °C and 150.21 °C respectively, with a mutual deviation of less than 1%), and finally by FC-72, which showed significantly poorer performance (166.52 °C) due to its lower latent heat of vaporization (83.54 kJ/kg, the lowest among the four fluids) and its relatively low liquid-phase thermal conductivity (0.057 W/(m·K)).

A result of particular methodological value, made possible by the decision to conduct the parametric study on all four fluids rather than on a single reference refrigerant, concerns the invariance of this performance ranking across the entire range of mass flow rates investigated. R134a proved systematically superior, Novec 7000 and Opteon SF-33 remained consistently equivalent to one another, and FC-72 was invariably the least performing fluid, at every one of the six mass flow rates examined. This invariance indicates that the comparative conclusions drawn regarding the selection of the most suitable refrigerant are not contingent upon the specific operating point adopted, and would therefore remain valid even if the design mass flow rate were subsequently revised, thereby considerably strengthening the robustness of the design recommendations formulated in this work.

The parametric study also brought to light a behavior specific to R134a that deserves explicit mention: owing to its high liquid-phase specific heat capacity and thermal conductivity, the vapor fraction produced by this fluid decreases far more rapidly with increasing mass flow rate than for the other candidates, falling to 12.53% at the optimal operating point and to 8.99% at 4.00 g/s. At the higher mass flow rates, therefore, R134a operates in a regime in which the contribution of latent heat to the overall heat transfer becomes marginal and the refrigerant remains in the liquid phase throughout almost the entire channel. Although this does not compromise its absolute thermal performance — indeed R134a remains the best performing fluid — it does mean that the conceptual advantage specific to two-phase cooling is only partially exploited by this refrigerant at the selected operating point.

The critical discussion presented in Chapter 7 made it possible to translate these numerical findings into a structured decision-making framework by introducing an equivalent thermal resistance R_{th},

defined as the ratio between the winding temperature rise relative to a common reference temperature (T_{sat}, = 60 °C) and the generated thermal power, as a summary metric for consistently comparing the cooling effectiveness of the different cases analyzed, independently of the dominant heat-transfer mechanism (latent or sensible). This analysis made it possible to quantify the main advantage of the two-phase cooling strategy over an equivalent single-phase solution, appropriately selected as a reference case at the same fluid outlet temperature: while providing comparable or superior thermal performance (equivalent thermal resistance ranging from 0.527 °C/W to 0.732 °C/W for the two-phase fluids, compared with 0.539 °C/W for the reference single-phase case), the two-phase fluids require an approximately 3.7 times lower refrigerant mass flow rate and a pressure drop more than one order of magnitude lower (approximately 14 times lower). This results in benefits in terms of required pumping power, size of the refrigerant supply lines, and energy consumption of the system auxiliaries, which are particularly relevant in automotive applications, where overall vehicle energy efficiency represents a primary design requirement and every watt diverted from traction to power auxiliary systems directly translates into a reduction in driving range.

Table 43 – Summary of the Main Quantities from the Two-Phase and Single-Phase Simulations

Parameter	Value / Range
Total Thermal Power to Be Dissipated	145.42 W
Optimal Two-Phase Refrigerant Mass Flow Rate (per Channel)	3.5 g/s
Mass Flow Rate Reduction Compared with the Single-Phase Case	3 / 7 times
Pressure Drop Reduction Compared with the Equivalent Single-Phase Case	14 times
Maximum Winding Temperature, Best-Performing Two-Phase Fluid (R134a)	136.65 °C
Maximum Winding Temperature, Compromise Two-Phase Fluids (Novec 7000 – Opteon SF-33)	149.91 – 150.21 °C
Thermal Margin Relative to the Class H Limit (180 °C)	13.5 – 43 °C
Operating Pressure of the Candidate Two-Phase Fluids	113500 – 1680000 Pa

The decision-making framework outlined above led to the identification of Novec 7000 and Opteon SF-33 as the most balanced candidates for the application under consideration, owing to their sufficiently large thermal margin, moderate operating pressure (on the order of 2–2.5 bar absolute, compatible with relatively simple system configurations and standard rotating seals), and favorable properties in terms of non-flammability and reduced environmental impact compared with conventional hydrofluorocarbon refrigerants. R134a was identified as the preferred option when the primary criterion is the maximization of the thermal margin, for example in applications characterized by particularly demanding load cycles, while accepting the greater system complexity associated with the high operating pressure required (approximately 17 bar absolute) and the more limited exploitation of the latent-heat mechanism discussed above. Finally, FC-72 was found to be the least suitable candidate, being penalized by poorer thermal performance that is not sufficiently

compensated by an advantage in terms of operating pressure, and further disadvantaged by considerations related to environmental persistence.

Overall, the work carried out demonstrated the technical feasibility and performance advantages of the two-phase cooling strategy for the rotor of the EESM investigated in this study, while also providing a structured methodological framework — from the grid-independence study, through the determination of the optimal mass flow rate by means of a diminishing marginal-benefit criterion applied to the mean response of all four candidate fluids, to the multi-criteria comparison summarized by an equivalent thermal resistance metric — that is fully reproducible and applicable to future design developments of the component, as well as to similar rotating electrical components characterized by comparable cooling requirements in terms of power density and packaging constraints. Therefore, the contribution of the present work lies not so much in the individual numerical values of temperature or optimal mass flow rate identified, but rather in the developed comparison methodology, which can be transferred to future variants of the component, different motor operating conditions, or alternative refrigerants that may become available on the market.

8.2 Possible Extensions and Future Studies

The limitations of the numerical model critically discussed in section 7.3 naturally identify the main directions along which the present work could be extended and further investigated in future developments. Although presented separately for clarity, these research directions are in fact strongly interconnected: for example, a computational domain extended to multiple channels would most likely also require the introduction of the rotational effects discussed in the second research direction, while experimental validation represents a necessary prerequisite for assigning full quantitative reliability to any further refinement of the model. These possible extensions are presented below according to five main research directions, discussed in sufficient detail to potentially provide the basis for a future research plan.

Fluid-Specific Optimization of the Mass Flow Rate

The parametric study presented in Chapter 6 was deliberately conducted using a mass-flow-rate grid common to all four candidate fluids, in order to isolate the influence of the thermophysical properties of the refrigerant and thereby guarantee the fairness of the subsequent comparison. This methodological choice, while essential to the comparative objective of the present work, implies that the selected operating point of 3.50 g/s per channel represents a compromise optimized for the set of candidates as a whole rather than for any individual refrigerant. The results themselves indicate that the fluid-specific optima would in fact differ appreciably: R134a already achieves an adequate thermal margin at substantially lower mass flow rates, and would arguably benefit from operating below 3.50 g/s in order to preserve a more significant contribution of latent heat transfer, whereas FC-72, which exhibits the narrowest thermal margin at the selected operating point, would require a higher mass flow rate to attain a margin comparable to that of the other candidates.

A natural extension of the present work would therefore consist of performing a dedicated optimization for each refrigerant, aimed at identifying the mass flow rate that maximizes the thermal margin while simultaneously maintaining the outlet vapor fraction within a prescribed target range. Such an analysis would shift the objective from a comparison at equal supply conditions, as adopted here, to a comparison between individually optimized configurations, which is ultimately the more representative basis for a final design decision. It would also make it possible to quantify, for each fluid, the total refrigerant mass flow rate actually required by the system, a quantity of direct relevance to the sizing of the supply circuit and to the overall energy balance of the cooling system.

Extension of the Computational Domain to Multiple Cooling Channels

The model developed in the present work considers a computational domain reduced to one-eighth of the complete rotor, comprising only two cooling channels and exploiting the assumption of cyclic symmetry of the component. Although this assumption is reasonable under nominal operating conditions and in the presence of perfect manufacturing symmetry, it does not allow potential interactions between adjacent channels to be captured, nor does it make it possible to verify the actual uniformity of refrigerant distribution among the different channels of the complete rotor. A natural future development would therefore consist of extending the computational domain to the entire rotor, with all cooling channels supplied by a common distribution manifold, in order to verify the assumption of uniform flow distribution adopted in the present study and to assess potential flow-maldistribution phenomena among the different channels. This phenomenon is particularly critical in a two-phase cooling system for at least two concurrent reasons: on the one hand, two-phase pressure drop does not vary linearly with mass flow rate, meaning that small pressure asymmetries within the supply manifold may be amplified into significant differences in mass flow rate among parallel channels; on the other hand, even a modest reduction in mass flow rate within a single channel results in a proportional increase in the local vapor fraction and, potentially, in the onset of localized dry-out conditions not predicted by the present model, with the formation of hot spots that could become critical to the service life of the electrical insulation in that specific region of the rotor.

From an implementation perspective, such an extension would require the design of an appropriate distribution manifold, whose geometry (inlet cross-sectional area and any calibrated restrictions introduced to equalize the pressure drops among channels located at different distances from the supply point) would itself become an optimization objective. It would also involve a significant increase in the computational resources required, both in terms of the number of computational cells and simulation time, owing to the increased geometric complexity of the domain and the need to maintain sufficient grid resolution to accurately resolve the phase-change process within each channel. A possible intermediate solution, from the standpoint of computational cost, could consist of modelling a larger rotor sector (for example, one-quarter or one-half of the component), while still retaining a certain degree of symmetry but including a greater number of adjacent channels. This would make it possible to capture potential thermal interactions between neighboring channels without incurring the full computational cost associated with the complete geometry. Alternatively, a co-simulation approach could be adopted by coupling a detailed CFD model of a single representative channel with a simplified one-dimensional (1D) fluid-network model of the entire distribution manifold, capable of estimating the mass-flow-rate distribution among the channels with a significantly lower computational cost than a complete three-dimensional CFD simulation.

Modelling of Three-Dimensional and Rotational Effects

Although the present model is inherently three-dimensional in terms of computational-domain geometry, it does not account for several physical effects associated with the rotating nature of the component and the actual three-dimensional distribution of the heat sources. In particular, one possible development would consist of introducing a rotating reference frame into the computational domain (rotating frame or Multiple Reference Frame, potentially coupled with a dynamic-mesh approach for the transient stage), capable of capturing the effects of centrifugal and Coriolis forces on the dynamics of the two phases within the channel. These forces, which are non-negligible at rotational speeds typical of an automotive electric motor (4200 rpm in the case considered here), could induce radial stratification of the liquid and vapor phases, with the denser liquid phase being driven toward the outer wall of the channel (located at a larger radius from the axis of rotation). This could locally alter the heat-transfer coefficient compared with that predicted by an analysis conducted under non-rotating conditions and could potentially also modify the circumferential location within the channel at which the maximum temperatures occur, an aspect that is not captured by the current model.

A second possible development concerns the modelling of a non-uniform heat-generation distribution along the axial length of the winding, replacing the assumption of scaled uniform heat generation adopted in the present work (72.7 W for each half-winding, obtained by scaling the actual losses of 88.75 W to account for the concentration of losses in the end-winding regions). This could potentially be achieved by coupling the thermofluid-dynamic model with a dedicated electromagnetic model (for example, based on the Finite Element Method and implemented in Ansys Maxwell, Motor-CAD, or JMAG), capable of providing a realistic map of Joule losses, including those concentrated in the end-winding regions and any additional eddy-current losses in the lamination and pole expansions, which are not currently modelled explicitly. Although demanding from an implementation perspective, such an electromagnetic–thermal–fluid-dynamic co-simulation approach would represent a significant step toward the development of a truly predictive model of the thermal behavior of the motor under actual operating conditions.

Finally, extending the computational domain to include the air gap and, potentially, the motor stator would make it possible to replace the current simplified boundary conditions (constant and uniform convective heat-transfer coefficient, h_{conv} = 40 W/(m²·K), with a $T_{free\ stream}$ of 50 °C) with a conjugate heat-transfer analysis that is more representative of the actual motor operating conditions. In such a model, the heat-transfer coefficient within the air gap would itself depend on the rotor rotational speed and the local fluid-dynamic conditions, while the thermal interaction between the rotor and stator, both of which are sources of heat generation, could no longer be neglected in an overall thermal analysis of the electric motor.

Experimental Validation

As discussed in section 7.3, the entire analysis presented in this thesis is purely numerical, in the absence of direct experimental validation. A natural and high-priority future development would therefore consist of designing and constructing an experimental test rig, potentially using a configuration simplified with respect to the complete rotor geometry, reproducing a single cooling channel with an electrically heated winding (for example, using a calibrated electrical resistance

capable of reproducing Joule heat generation and instrumented with thermocouples at the locations of greatest interest). The channel would be supplied with the two-phase refrigerant under controlled mass-flow-rate, inlet-temperature, and saturation-pressure conditions, potentially adjustable through a dedicated conditioning circuit equipped with a variable-flow pump, a condensation heat exchanger, mass-flow-rate measurement instrumentation (for example, a Coriolis flow meter capable of operating correctly even in the presence of two-phase mixtures), and pressure-drop measurement instrumentation (differential pressure transducers).

Such an experimental campaign would enable quantitative validation of the numerical results obtained, particularly in terms of maximum temperatures, pressure drop, and outlet vapor fraction. The latter could be estimated indirectly through an enthalpy balance applied to the conditioning circuit or through optical visualization techniques for the two-phase flow using a transparent section of the cooling channel. The experimental data would also provide direct evidence regarding the adequacy of the Lee-model coefficients, which are currently set equal to unity, potentially allowing fluid-specific calibration through an inverse optimization procedure in which the evaporation and condensation coefficients are iteratively adjusted to minimize the discrepancy between simulated and measured temperatures. A further aspect that would benefit from experimental verification concerns the non-monotonic pressure-drop behavior observed for three of the four fluids between 3.50 and 4.00 g/s per channel, discussed in Chapter 6, for which a purely numerical origin could not be ruled out in the absence of measured data. More generally, the availability of experimental results would make it possible to assess the influence of the numerical simplifications discussed in section 7.3 (discretization schemes, grid resolution, and duration of the transient stage) on the quantitative accuracy of the results and, if necessary, to recalibrate selected parameters of the numerical model, while also providing specific experimental verification of the assumption of a constant saturation temperature along the channel, which has not been verified in the present study.

Extension to Other Refrigerants

The comparison conducted in the present work considered four two-phase refrigerants (Novec 7000, FC-72, Opteon SF-33, and R134a), selected as representative of different chemical families of refrigerants currently available on the market. A possible future development would consist of extending the comparison to other next-generation refrigerants, particularly hydrofluoroolefins (HFOs), such as R1234yf and R1234ze, which are already widely used in automotive air-conditioning systems, as well as to their corresponding blends. These fluids are currently attracting increasing interest in the refrigeration and automotive air-conditioning sectors owing to their lower global warming potential (GWP) compared with conventional hydrofluorocarbons, in accordance with European regulations on fluorinated greenhouse gases (F-Gas Regulation), which progressively restrict the use of fluids with higher GWP values and promote their replacement with alternatives characterized by a lower environmental impact. Such an extension would be particularly significant from the perspective of ensuring the long-term regulatory compliance of the proposed design solution, considering the typical development times of an automotive component, which are on the order of several years, during which the regulatory framework governing refrigerants may evolve further.

A further possible development concerns the evaluation of refrigerant mixtures (blends), which could, in principle, provide thermophysical properties intermediate between those of the pure fluids

considered in the present work, offering the possibility of jointly optimizing thermal performance and the required operating pressure. This could enable the development of a fluid approaching an ideal combination of the high thermal conductivity and high specific heat capacity of R134a with the moderate operating pressure of Novec 7000 and Opteon SF-33. Such optimization, however, would require the availability of accurate thermodynamic models for the mixture under consideration, which may not be readily available for non-standard compositions. An additional selection criterion to be investigated in future developments concerns the safety classification of the fluids in terms of toxicity and flammability, according to standards such as the ASHRAE 34 classification. This aspect was not quantitatively addressed in the present work but could potentially represent a decisive factor in the final selection of the refrigerant for a production automotive application.

Finally, extending the study to variable operating conditions (different motor operating points in terms of rotational speed and load level, up to and including an entire standardized driving cycle such as the WLTP) would make it possible to assess the robustness of the refrigerant selection and optimal mass flow rate identified in the present work, which was conducted with reference to a single operating point (4200 rpm, half load, continuous operation), over the entire operating range of the electric motor. Such an extension would also make it possible to develop, if necessary, a refrigerant mass-flow-rate control strategy capable of dynamically adapting to the instantaneous thermal load, for example through a control logic based on monitoring the fluid outlet temperature or its degree of superheat relative to saturation conditions. The objective would be to maintain the outlet vapor fraction within the identified optimal range (approximately 40–50%) even under partial-load or transient-overload conditions, while simultaneously avoiding excessive energy consumption by the circulation pump under reduced-load operating conditions.

In conclusion, although the future developments outlined in this section do not undermine the validity of the design recommendations provided by the present thesis, they represent their natural continuation toward the progressive refinement of the numerical model and the consolidation of the quantitative findings with a view to the future industrial application of two-phase cooling systems for high-performance electric motors. A potential continuation of this research could reasonably be structured according to a sequence of increasing implementation effort and increasing solution maturity: first, the fluid-specific optimization of the mass flow rate, which requires no modification of the existing numerical model and would immediately refine the design recommendations formulated here; subsequently, the extension to multiple cooling channels and the modelling of rotational effects, as these are directly preparatory to a realistic assessment of the performance of the complete component; then experimental validation, as an essential prerequisite for assigning full quantitative reliability to the numerical model; and finally, the extension to new refrigerants and variable operating conditions, as the concluding step toward the definition of a mature design solution that complies with the current regulatory framework and is ready for potential industrialization.

Bibliography

[1] Chengdi Xiao, Longchao Liu, Gongbing Chen, Haitao Zhang, Xixin Rao, Yi Zhou, Mechanisms of temperature distribution in large rotary compressor motors: experimental and simulation study on refrigerant-oil two-phase heat transfer, Applied Thermal Engineering 279 (2025) 127757.

[2] X. Wang, B. Li, D. Gerada, K. Huang, I. Stone, S. Worrall, Y. Yan, A critical review on thermal management technologies for motors in electric cars, Applied Thermal Engineering 201 (2022) 117758.

[3] Sudharsan Vasudevan, Sassan Etemad, Lars Davidson, Mirko Bovo, Comparative analysis of single and multiphase numerical frameworks for subcooled boiling flow in an internal combustion engine coolant jacket, Applied Thermal Engineering 219 (2023) 119435.

[4] Jong Seung Na, Choong-Mo Ryu, Seung Jae Moon, Thermal analysis methods for wound rotor synchronous motors with water-cooling jackets, Thermal Science and Engineering Progress 46 (2023) 102177.

[5] Amitav Tikadar, Joon W. Kim, Yogendra Joshi Fellow, Satish Kumar, Flow assisted evaporative Cooling for electric motor.

[6] Yew Chuan Chong, Melanie Michon, David Gerada, Saeed Jahangirian, Chuan Liu, Chris Gerada, Husain Adam, Zeyuan Xu, Analytical modelling technique of oil-cooled electric machine using empirical correlation and LPTN.

[7] Yaohui Gai, Mohammad Kimiabeigi, Yew Chuan Chong, James D. Widmer, Xu Deng, Mircea Popescu, James Goss, Dave A. Staton, Andrew Steven, Cooling of automotive traction motors: schemes, examples, and computation methods, IEEE transactions on industrial electronics, vol. 66, NO. 33, March 2019.

[8] Luca Boscaglia, Hari Sankan Natesan Sugumar, Nimanada Sharma, Yujing Liu, Design and verification of an electrically excited synchronous machine rotor with direct oil cooling for truck applications, IEEE transactions on industrial electronics, vol. 11, NO. 1, February 2025.

[9] Yue Tang, Lei Chen, Feng Chai, Tanci Chen, Thermal modeling and analysis of active and end windings of enclosed permanent-magnet synchronous in-wheel motor based on multi-block method, IEEE transactions on industrial electronics, vol. 35, NO. 1, March 2020.

[10] Yu Wang, Lin Ruan, Self-circulating evaporative cooling system of a rotor and its experimental verification, Processes MDPI.

[11] Chengyang Ye, Irving S. Aguilar-Zamorate, Andrea Tonoli, Renato Galluzzi, Edoardo Lagorio, Federica Graffeo, Simone Ferrari, Silvio Vaschetto, Gianmario Pellegrino, Vittorio Ravello, Raffaele Bonavolontà, Oil rotor cooling comparison for electrically excited synchronous machine.

[12] Li Zhenguo, Ruan Lin, Tang Longyao, Heat transfer characteristics of spray evaporative cooling system for a large electrical machine, 2015 18th international conference on electrical machines and system (ICEMS), Oct 25-28, 2015, Pattaya City, Thailand.

[13] Dmytro Konovalov, Ignat Tolstorebrov, Yuhiro Iwamoto, Halina Kobalava, Jacob Joseph Lamb, Trygve Magne Eikevik, Optimizing low-temperature three-circuit evaporative cooling system for an electric motor by using refrigerants, Energies MDPI.

[14] Dmytro Konovalov, Ignat Tolstorebrov, Trygve Magne Eikevik, Halina Kobalava, Mykola Radchenko, Armin Hafner, Andrii Radchenko, Recent developments in cooling systems and cooling management for electric motors, Energies MDPI.

[15] Jeremy B. Campbell, Leon M. Tolbert, Curt W. Ayers, Burak Ozpineci, Kirk T. Lowe, Two-phase cooling method using the R134a refrigerant to cool power electronic device, IEEE transactions on industry applications, vol. 43, NO. 3, May/June 2007.

[16] L. Cattani, M. Malavasi, F. Bozzoli, C. Sciancalepore, Two-phase cooling system for electric vehicles' battery, Journal of physics: conference series 2766 (2024) 012011.

[17] Aleksandar Ristic-Smith, Daniel J. Rogers, Compact two-phase immersion cooling with dielectric fluid for PCB-based power electronics, IEEE open journal of power electronics.

[18] Wikipedia, "Renault Megane E-Tech Electric" — scheda tecnica, motore dichiarato EESM.

[19] Wikipedia, "Renault Scenic E-Tech" — piattaforma AmpR Medium condivisa con Mégane E-Tech e Nissan Ariya.

[20] Wikipedia, "Renault Scenic E-Tech" — modelli correlati, tra cui Nissan Ariya.

[21] ScienceDirect, "Feasibility study of rare earth free electrically-excited machine technology for EV powertrain" (2025) — descrizione del sistema di raffreddamento a spruzzo d'olio della Nissan Ariya.

[22] Wikipedia, "BMW iX3" — motore posteriore "Externally-excited synchronous", piattaforma Neue Klasse, produzione da ottobre 2025.

[23] IEEE Spectrum, "German EV Motor Could Break Supply-Chain Deadlock" — motore ZF a doppia eccitazione da 220 kW privo di terre rare.

[24]] IEEE Spectrum, "EV Motors Without Rare Earth Permanent Magnets" — collaborazione Stellantis / General Motors / Niron Magnetics, con finanziamento parziale del Dipartimento dell'Energia USA.

[25] Interesting Engineering, "Why electric motor makers are moving away from rare-earth magnets" — progetto congiunto Stellantis–Niron per motori a induzione/EESM su modelli Chrysler e Jeep.

[26] ScienceDirect, "Feasibility study of rare earth free electrically-excited machine technology for EV powertrain" (2025) — assenza in letteratura di soluzioni di raffreddamento diretto lato albero per EESM; panoramica delle strategie di raffreddamento a spruzzo attualmente impiegate.

Technical data sheets

[27] AmpCool AC-100 series dielectric coolants, doc. Id: AC1XX-ENG-TDS-20260505.

[28] 3M Novec 7000 Engineered fluid, 3M Novec Brand.

[29] R134a, Climalife.

[30] 3M Fluorinet electronic liquid FC-72, 3M science applied to life.

[31] Opteon SF-33, Speciality & two-phase direct-to-chip fluid, Chemours.