FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO

Development and Simulation of a Motor Controller for a Formula Student Prototype

Francisco José Rodrigues de Lemos



Mestrado em Engenharia Eletrotécnica e de Computadores

Supervisor: Prof. Armando Luís Sousa Araújo

October 28, 2026

© Francisco José Rodrigues de Lemos, 2026

Resumo

Esta tese apresenta o desenvolvimento de um controlador de motor para um protótipo de Formula Student. Este tipo de projectos desafia as equipas universitárias a conceber e fabricar veículos avançados, exigindo uma combinação de conhecimentos de engenharia e aplicações práticas. O controlador desenvolvido nesta tese visa melhorar o desempenho e a competitividade do veículo e reforçar a qualidade técnica do projeto com componentes do grupo motopropulsor fabricados internamente, bem como contribuir para o *know-how* da equipa e para futuros desenvolvimentos. O sistema centra-se no controlo de um motor síncrono de ímanes permanentes (PMSM) utilizando um conversor CC-CA previamente desenvolvido por um membro da equipa, de acordo com as regras da competição e os objetivos da equipa.

A investigação iniciou-se com uma revisão bibliográfica aprofundada dos aspectos teóricos e do estado da arte. Focando-se nos motores eléctricos, os seus tipos e aplicações, os conversores de potência e as suas diferentes tipologias e métodos de controlo adequados aos objectivos do projeto.

O processo de desenvolvimento incluiu a modelação do motor em software de simulação e a conceção e otimização do controlador específico adaptado. Este trabalho estabelece as bases para futuros desenvolvimentos internos, destinados a melhorar as capacidades de inovação da equipa e o sucesso a longo prazo na competição de Formula Student.

Palavras-chave: Formula Student, Controlador, Motor Síncrono de Ímanes Permanentes, Controlo Orientado por Campo, Modulação por Vetor Espacial, Binário Máximo por Ampere.

Abstract

This thesis presents the development of a motor controller for a Formula Student prototype. This type of projects challenges university teams to design and manufacture advanced vehicles, requiring a combination of engineering expertise and practical applications. The controller designed in this thesis aims to improve the performance and competitiveness of the vehicle and enhance the project technical quality with in-house made powertrain components, as well as contributing to the know-how within the team and to future developments. The system focuses on controlling a permanent magnet synchronous motor (PMSM) using a previously self-developed DC-AC converter, in accordance with the competition rules and the team's objectives.

The research began with an in-depth literature review of theoretical aspects and state-of-the-art. Focusing on electric motors, its types and applications, power converters and their different typologies and control methods suitable for the objectives of the project.

The development process included modelling the motor in simulation software and designing and optimizing the specific tailored controller. This work lays the foundation for future in-house developments, aimed at enhancing the team's innovation capabilities and long-term success in Formula Student competition.

Keywords: Formula Student, Motor Controller, Permanent Magnet Synchronous Motor (PMSM), Field-Oriented Control (FOC), Space-Vector Modulation (SVM), Maximum Torque Per Ampere (MTPA).

UN Sustainable Development Goals

The United Nations Sustainable Development Goals (SDGs) provide a global framework to achieve a better and more sustainable future for all. It includes 17 goals to address the world's most pressing challenges, including poverty, inequality, climate change, environmental degradation, peace, and justice.

This work develops and validates a motor controller for a permanent magnet synchronous motor (PMSM) intended for use in a Formula Student electric vehicle prototype. By contributing to the study of electric drivetrain modelling and control, the work is related to the development of more efficient electric mobility technologies and the advancement of engineering knowledge. The modelling, simulation, and controller validation methods developed in this dissertation may also support further research into electric motor drives and their application in sustainable transport systems.

The specific Sustainable Development Goals addressed by this work are the following:

SDG 7 Ensure access to affordable, reliable, sustainable and modern energy for all

SDG 9 Build resilient infrastructure, promote inclusive and sustainable industrialisation and foster innovation

SDG 11 Make cities and human settlements inclusive, safe, resilient and sustainable

SDG	Target	Contribution	Performance Indicators and Metrics
7	7.a	The modelling and simulation methods support further research into electric drives and sustainable energy technologies.	Model validation and applicability to other electric drive systems.
9	9.5	The development and validation of the PMSM model and controller contribute to research and innovation in electric drivetrain engineering.	Model accuracy and controller performance under simulated conditions.
11	11.2	The work indirectly supports sustainable transport through research into electric drivetrain technology.	Simulated traction and regenerative braking performance.

Acknowledgements

I would like to thank my supervisor, Professor Armando Araújo, for his help during the development of this thesis, and the availability to guide a project on a theme I am interested in.

I would also like to thank Vasco da Gama CoLAB, especially Engineer Jorge Pinto, for always finding the time to help me and see this work through to the end.

Thanks to Formula Student FEUP for giving me the opportunity to learn a tremendous amount about engineering and make some great friends.

Special thanks to my family. Thank you to my sister for putting up with me ever since I came to study in Porto (and since I was born) and for always being ready to do anything for me. Thank you to my parents for showing me what it means to be a good person, how to have principles and values and for dedicating their lives to my sister and me.

A huge thanks to the friends I made during this journey. I feel extremely fortunate to have crossed paths with the friends I made at this university.

Thank you to Tourigo and all the *Touriguenses* for being the happy home my heart will never leave.

Lastly, a special thank you to the *Antiga, Mui Nobre, Sempre Leal e Invicta Cidade do Porto* and to all the *Portuenses* I met here and will take with me for life.

Francisco Lemos

"Este é o Porto de todos os barcos
Chegam os loucos voltam encantados"

Pedro Abrunhosa

Contents

List of Symbols xii

1 Introduction 1
- 1.1 Document Structure . . . 1
- 1.2 Formula Student . . . 1
- 1.3 Motivation and Objectives . . . 2

2 State of the Art and Design Requirements 4
- 2.1 Electric Motors . . . 4
  - 2.1.1 DC Motors . . . 5
  - 2.1.2 Brushless DC Motors . . . 5
  - 2.1.3 Switched Reluctance Motors . . . 6
  - 2.1.4 Synchronous Reluctance Motor . . . 7
  - 2.1.5 Induction Motors . . . 7
  - 2.1.6 Permanent Magnet Synchronous Motors . . . 8
- 2.2 Power Converters . . . 10
  - 2.2.1 Two-level three-phase topology . . . 11
  - 2.2.2 Three-level three-phase topology . . . 12
  - 2.2.3 Open-Winding . . . 13
  - 2.2.4 T-type . . . 14
- 2.3 Modulation Techniques . . . 15
  - 2.3.1 Sinusoidal PWM . . . 15
  - 2.3.2 Space-Vector Modulation . . . 16
- 2.4 Control Techniques . . . 17
  - 2.4.1 Scalar Control (V/f) . . . 18
  - 2.4.2 Direct Torque Control . . . 19
  - 2.4.3 Field-Oriented Control . . . 19
  - 2.4.4 Current Reference Strategies . . . 21

3 System Modelling and Controller Design 23
- 3.1 Modelling of the Permanent Magnet Synchronous Motor . . . 23
  - 3.1.1 FEM-Parameterized PMSM block . . . 23
  - 3.1.2 Processing of the finite element data . . . 26
  - 3.1.3 Application of the data to the block . . . 28
  - 3.1.4 Motor model validation . . . 29
  - 3.1.5 FEM-Parameterized PMSM vs Linear Model . . . 33
  - 3.1.6 Constant Parameters Used Outside the Machine Model . . . 35
- 3.2 Inverter . . . 35

3.3 Space-Vector Modulation PWM . . . 36
3.4 Current Controller . . . 40
3.4.1 Control Design . . . 40
3.5 Maximum Torque per Ampere . . . 45
3.5.1 Trajectory . . . 45
3.5.2 Parameter Values . . . 46
3.5.3 Implementation . . . 46

4 Results and Discussion **50**
4.1 Open-Loop V/f Control . . . 50
4.1.1 Results . . . 50
4.2 Current Controller . . . 52
4.2.1 Validation Tests . . . 52
4.3 Maximum Torque per Ampere . . . 59
4.3.1 Verification . . . 59
4.3.2 Scope and Limitations . . . 63
4.3.3 Complete System . . . 63
4.3.4 Simulation with Vehicle . . . 68

5 Conclusions and Future Work **72**
5.1 Achievement of Objectives . . . 72
5.2 Future Work . . . 73

References **74**

A Motor datasheet TI085-052-070-04B7S-07S04BE2 **77**

B Excerpt TI085-052-070-04B7S-07S04BE2 FEA spreadsheet **81**

List of Figures

2.1 Electric motors [3] . . . 4
2.2 Comparison between (a) BLDC motors and (b) [3] . . . 6
2.3 Rotation principle of permanent magnet synchronous motor [12] . . . 8
2.4 (A) Radial-flux type and (B) axial-flux type [3] . . . 9
2.5 (A) Surface mounted, (B) interior (parallel), and (C) interior (perpendicular) [3] . . . 9
2.6 Inverter topologies [13] . . . 11
2.7 Two-level 3-phase [2] . . . 11
2.8 NPC (a) and NPP (b) [15] . . . 13
2.9 Single source open-winding [15] . . . 13
2.10 Three-phase T-type topology [17] . . . 14
2.11 Sinusoidal PWM operating principle and corresponding harmonic spectrum [18] 15
2.12 Space-vector diagram of the two-level converter [18] . . . 16
2.13 Space-vector PWM: modulating signal, triangular carrier and resulting gate pulses [20] . . . 17
2.14 Control techniques for PMSM drives [21] . . . 18
2.15 Open-loop V/f control structure [23] . . . 18
2.16 Direct torque control [25] . . . 19
2.17 Field-oriented control [25] . . . 20
2.18 Current combinations producing the same torque and the MTPA operating point [3] 22

3.1 Interior PMSM steady-state phasor diagram . . . 25
3.2 Parameter maps loaded into the *FEM-Parameterized PMSM* block. . . . 28
3.3 FEM-Parameterized PMSM block on *Simulink* . . . 29
3.4 Torque comparison against the FEA reference. . . . 31
3.5 Relative torque error over the current plane. . . . 32
3.6 Simulated back-EMF versus mechanical speed. . . . 33
3.7 Torque along the $i_d = 0$ arc. . . . 34
3.8 Relative error versus current magnitude. . . . 34
3.9 Inverter model on *Simulink* . . . 36
3.10 Simulink implementation of the SVM modulator and of the PWM gate signal generation. . . . 37
3.11 Three-phase sinusoidal voltage references applied to the modulator. . . . 38
3.12 SVM modulating waveforms and triangular carrier. . . . 38
3.13 Complementary gate signals SW1 (upper switch) and SW2 (lower switch). . . . 39
3.14 Phase-to-neutral inverter output voltages generated by the SVM gate signals. . . . 39
3.15 FOC block diagram on *Simulink* . . . 45
3.16 Variation of k across the finite element dataset and its effect on the trajectory. . . . 47
3.17 Reference generation chain implementing the MTPA trajectory. . . . 48

3.18 MTPA trajectory in the dq current plane, compared with the $i_d = 0$ axis. 48
3.19 Polynomial inversion of the torque-to-current relation and the corresponding error. 49

4.1 Commanded and measured shaft speed during the open-loop V/f ramp. 51
4.2 Phase currents during the open-loop V/f ramp. 51
4.3 Current step response at 500 rpm. 53
4.4 Feedforward subsystem on *Simulink*. 54
4.5 Current step at 8000 rpm with and without decoupling feedforward. 55
4.6 Voltage limiter on *Simulink*. 55
4.7 Current step into sustained voltage saturation, with and without anti-windup. . . . 56
4.8 Reversal of the quadrature current reference at 500 rpm. 57
4.9 Torque tracking and speed response to a load step on a free shaft. 58
4.10 Direct and quadrature currents during the system-level test. 58
4.11 Coincidence of the analytical trajectory with the minimum of the current magnitude, and the resulting current reduction. 60
4.12 References produced by the block chain compared with the analytical trajectory. . 61
4.13 Delivered torque against command over the torque range, and the resulting relative error. 62
4.14 Phase current over several fundamental periods and magnified to four switching periods. 64
4.15 Electromagnetic torque and dq currents over four switching periods. 65
4.16 DC Bus model on *Simulink* 66
4.17 Torque step and the resulting drop in DC bus voltage. 67
4.18 Spectrum of the phase current, showing the fundamental and the switching families. 68
4.19 Mechanical Load on *Simulink* 69
4.20 Torque, vehicle speed and dq currents over an acceleration, coasting and regenerative braking manoeuvre. 71

List of Tables

1.1 Maximum points awarded [1] 2

2.1 Switching states of the two-level three-phase converter 12
2.2 Comparison of the presented inverter topologies 14

3.1 Configuration of the *FEM-Parameterized PMSM* block in Simscape Electrical. . 29
3.2 Simulated no-load line-to-line RMS voltage as a function of imposed mechanical speed. 32
3.3 Relative torque error against the FEA reference, linear model versus FEM-parametrised Simscape model. 33
3.4 Switching device model parameters (Infineon IMZA120R014M1H). 36
3.5 Phase margin of the d-axis loop under the modulator and computation delays. . . 43
3.6 Current regulator gains. 43
3.7 Coefficients of the polynomial inversion $i_q^*(T_e^*)$, with T_e in N·m and i_q in A peak. 47

4.1 Conditions of the current loop validation tests. 52
4.2 Delivered torque against command at 500 rpm. 62

List of Acronyms

BLAC	Brushless Alternating Current
BLDC	Brushless Direct Current
DC	Direct Current
DPWM	Discontinuous Pulse Width Modulation
DTC	Direct Torque Control
DV	Driverless
EMF	Electromotive Force
EV	Electric Vehicle
FEA	Finite Element Analysis
FEUP	Faculdade de Engenharia da Universidade do Porto
FEM	Finite Element Method
FOC	Field-Oriented Control
FSG	Formula Student Germany
FSFEUP	Formula Student FEUP
IGBT	Insulated-Gate Bipolar Transistor
IM	Induction Motor
MOSFET	Metal-Oxide-Semiconductor Field-Effect Transistor
MTPA	Maximum Torque Per Ampere
NPC	Neutral-Point-Clamped
NPP	Neutral-Point-Piloted
PI	Proportional-Integral
PMSM	Permanent Magnet Synchronous Motor
PWM	Pulse Width Modulation
RMS	Root Mean Square
SiC	Silicon Carbide
SPWM	Sinusoidal Pulse Width Modulation
SRM	Switched Reluctance Motor
SVM	Space-Vector Modulation
SynRM	Synchronous Reluctance Motor
VDC	Voltage Direct Current

List of Symbols

b	Viscous damping coefficient
f_{sw}	Switching frequency
i_a, i_b, i_c	Three-phase stator currents
i_α, i_β	Stationary frame ($\alpha\beta$) currents
i_d	D-axis stator current
i_q	Q-axis stator current
I_s	Stator current magnitude (RMS)
J	Rotor moment of inertia
k	MTPA trajectory parameter, $\psi_{PM}/(2\Delta L)$
k_e	Back-EMF constant (phase-to-phase)
k_t	Torque constant
K_i	PI regulator integral gain
K_p	PI regulator proportional gain
L_d	D-axis inductance
L_q	Q-axis inductance
L_0	Zero-sequence inductance
ΔL	Inductance saliency, $L_q - L_d$
p	Number of pole pairs
R_s	Stator resistance per phase
T_e	Electromagnetic torque
T_{load}	Load torque
T_s	Switching period
v_d	D-axis stator voltage
v_q	Q-axis stator voltage
V_{DC}	DC link voltage
θ_e	Electrical rotor position
ξ	Current-loop damping ratio
τ	PI zero time constant, K_p/K_i
ψ_d	D-axis flux linkage
ψ_q	Q-axis flux linkage
ψ_{PM}	Permanent-magnet flux linkage (peak)
ω_0	Current-loop natural frequency
ω_e	Electrical angular velocity
ω_m	Mechanical angular velocity

Chapter 1

Introduction

1.1 Document Structure

This dissertation is organized into five main chapters. Following this introductory chapter, Chapter 2 presents a comprehensive state-of-the-art review, covering the main types of electric motors, power converter topologies, modulation strategies, and control techniques applicable to electric vehicles. Chapter 3 details the methodology and implementation of the proposed system, including the FEA-based modeling of the permanent magnet synchronous motor, the space-vector modulation stage, the design of the current regulators, and the MTPA trajectory formulation. Chapter 4 presents and discusses the results obtained, covering open-loop V/f control, current controller validation tests, complete system simulations with the switching inverter, and vehicle-level maneuvers. Finally, Chapter 5 summarizes the main conclusions drawn from this work and provides recommendations for future developments.

1.2 Formula Student

Formula Student is an international university organization that challenges student teams to fully design and manufacture vehicles before taking their projects to compete in events that join teams, students, teachers and engineers from all around the world. This results in a highly demanding work from the teams over the progress of their project, it develops their engineering knowledge and more importantly its practical application. Soft skills such as short deadlines and intense workflow, team work and communication are also improved within the students.

The competitions are divided in static events and dynamic events [1]. In static events, the prototypes are put to the proof by professionals who evaluate the work done from an engineering point of view and also in terms of business and finance. The dynamic events is where the on road prototype's performance matters, with four events, namely acceleration, skidpad, autocross and endurance. The time marked by the car on track defines the classification, except in the endurance event, where its efficiency is also an evaluation factor.

Formula Student FEUP is a team from Faculty of Engineering of the University of Porto. It was created in 2021 and has had some very good results for a team with such a short experience, collecting awards in various competitions: Portugal, United Kingdom, Italy, Switzerland and Germany. The team works with cycles of two years, where the first year is dedicated to the prototype's design and the second to its manufacturing. Therefore every two years a new innovative car comes to life. Not only it is important to focus on the under development project, but it is also indispensable to define long term objectives, in order to guarantee continuity and success for the years coming.

The scoring of the competition is split between two classes: the Electric Vehicle (EV) class, in which all electric prototypes compete, and the Driverless Cup (DC), dedicated solely to autonomous driving, although EV teams may also collect points in the autonomous events. The maximum score attainable in each event is summarised in Table 1.1.

Table 1.1: Maximum points awarded [1]

	EV	**DC**
Static Events:		
Business Plan Presentation	75 points	-
Cost and Manufacturing	100 points	-
Engineering Design	150 points	150 points
Dynamic Events:		
Skidpad	50 points	-
Driverless (DV) Skidpad	75 points	75 points
Acceleration	50 points	-
Driverless (DV) Acceleration	75 points	75 points
Autocross	100 points	-
Driverless (DV) Autocross	-	100 points
Endurance	250 points	-
Efficiency	75 points	-
Trackdrive	-	200 points
Overall	1000 points	600 points

1.3 Motivation and Objectives

Currently on its second prototype, the FSFEUP team uses an off-the-shelf inverter and controller on its powertrain. However, the goal is to implement a system made in-house for the vehicles that will be designed in the future. The first step is taken, as a DC-AC converter has already been designed by a team member [2]. This thesis will focus on developing the controller to be applied on the referred converter and implemented in a future Formula Student FEUP car with one permanent magnet synchronous motor on each of the four wheels. The system must meet the

requirements not only of the competition rules but also of the technical and performance objectives aimed at by the team.

As the general objective is stated, there are smaller goals that must be achieved:

- Study and model the system of a Formula Student prototype according to the team design requirements on Matlab®/Simulink environment and all its subsystems, mainly the motor and inverter;
- Implement a simulation model of the motor;
- Implement and tune a simulation model of the motor controller;

This thesis is important for the Formula Student FEUP team future success not only in terms of performance of the prototype it will be applied on but also to create a strong base for other projects that may be developed on this area and possibly to follow-up this specific one.

Chapter 2

State of the Art and Design Requirements

2.1 Electric Motors

Figure 2.1 shows a diagram of the different types of electric motors [3]

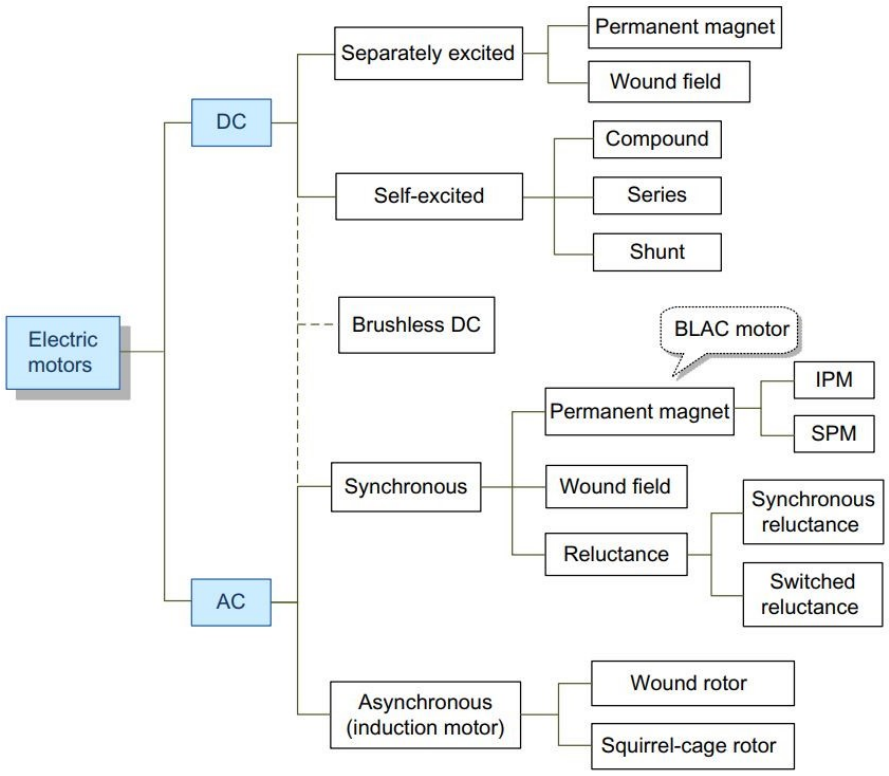


Figure 2.1: Electric motors [3]

The main electric motors used in electrical mobility applications will be described and compared in the following sections [4], [5].

2.1.1 DC Motors

In a direct-current motor, torque is produced by the interaction between the armature current and the magnetic field established by the stator, which may be created by permanent magnets or a field winding. As the rotor turns, each armature coil would reverse the direction of its torque each time it passes from one magnetic pole to the next if the direction of the current remained unchanged. To keep the torque acting in a single direction, the commutator — a segmented ring mounted on the rotor shaft, over which the fixed brushes slide — reverses the direction of the current in each coil at the instant it crosses between magnetic poles. This mechanical switching, synchronised with the rotor position by the geometry of the segments, ensures that the torque is continuously kept in the same direction [3].

DC motors have the advantage of being easily controlled, since flux and torque are inherently decoupled, and are able to produce a high starting torque at low speed [6]; however, due to the development and growth of control methods applied to other types of machines, preference for DC motors has decreased [7].

With low efficiency and power density, these motors are not a good option for high power applications. The brushes and commutator result in low reliability, as the commutation process causes sparking that wears down the contacts and generates electromagnetic interference, which is their main disadvantage [6].

2.1.2 Brushless DC Motors

Compared to brushed DC motors, brushless DC motors are more compact and reliable because the traditional mechanical commutation is replaced with electronic commutation, eliminating the brushes and commutator and therefore reducing maintenance needs.

Structurally, a BLDC motor can be seen as an inverted brushed DC motor: the permanent magnets are mounted on the rotor and the armature winding is placed on the stator, so no electrical connection to the rotating part is required. The role of the commutator is instead performed by a three-phase inverter switched according to the rotor position. One electrical revolution is divided into six sectors of 60 electrical degrees, typically identified by three Hall-effect sensors, and in each sector the inverter energizes two phases while the third is left open. By advancing the conducting phase pair from sector to sector — a scheme known as six-step commutation — the stator magnetic field is kept approximately perpendicular to the rotor field, so that torque is continuously produced in the same direction [3].

These machines have high efficiency and power density due to the rare-earth permanent magnets on the rotor.

Although BLDC motors are constructively very similar to permanent magnet synchronous machines, the two differ in the shape of their back-EMF: BLDC motors are designed with concentrated stator windings that produce a trapezoidal back-EMF and are consequently fed with rectangular currents. Since commutation between phases occurs at only six discrete instants per electrical cycle, low-resolution position feedback from Hall-effect sensors is sufficient, instead

of the continuous high-resolution position information required for the sinusoidal excitation of PMSMs [3]. Figure 2.2 illustrates this difference.

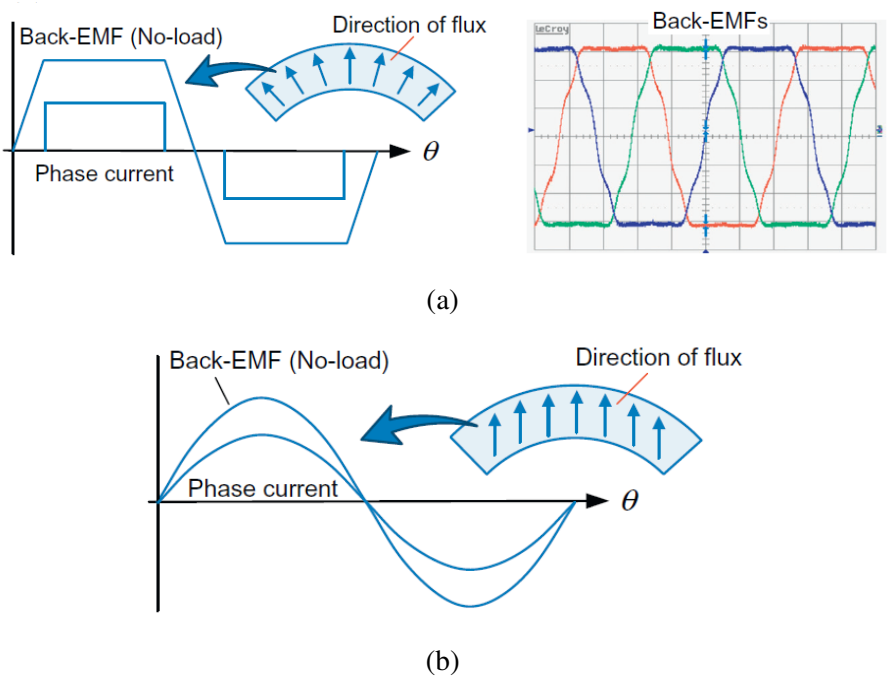


Figure 2.2: Comparison between (a) BLDC motors and (b) [3]

The main drawback of this drive principle is torque ripple. In addition to the ripple caused by the non-ideal shape of the back-EMF, every commutation between phases produces a transient torque disturbance, known as commutation torque ripple, because the phase inductances prevent the currents from switching instantaneously. This ripple component varies with speed and may reach 50% of the average torque [8]. The limited flux-weakening capability of this configuration results in a small constant power speed range, and for that reason it is not a good option for high speed applications.

2.1.3 Switched Reluctance Motors

The switched reluctance motor has a doubly salient structure: both the stator and the rotor have salient poles, but only the stator carries windings, while the rotor is a simple stack of laminations with no windings or permanent magnets. When a stator phase is energized, the rotor tends to move to the position of minimum reluctance of the magnetic circuit, where the inductance of the excited phase is maximum, producing a reluctance torque. By energizing the phases sequentially, according to the rotor position, a continuous rotation is obtained [9].

SRMs have a simple and robust construction and are more fault tolerant than most motors, since each phase operates independently. Dynamically speaking, these motors have a wide constant power region and a good torque-speed characteristic [6].

The fact that there are no coils on the rotor reduces its heat generation and makes it capable of operating at high speed, which results in a higher versatility.

When compared to other options, SRMs have high acoustic noise and torque ripple, caused by the doubly salient structure and the pulsed nature of the phase currents. These motors also need a special converter topology, which limits their application range [6].

2.1.4 Synchronous Reluctance Motor

The synchronous reluctance motor uses the same stator as conventional AC machines: a three-phase distributed winding that produces a sinusoidally rotating magnetic field. The rotor features flux barriers punched into the rotor laminations create a high-permeance path along the d-axis and a low-permeance path along the q-axis. As the rotor tends to align its minimum-reluctance axis with the rotating stator field, it locks to the field and rotates synchronously, producing a reluctance torque proportional to the inductance difference $(L_d - L_q)$ [10].

Unlike the SRM, whose phases are excited sequentially with current pulses, the SynRM is fed with sinusoidal currents from a standard three-phase inverter. When compared to SRMs, SynRMs are therefore more efficient, have smoother low speed torque and lower acoustic noise, and do not require a special converter topology, while still being small and robust. In general these motors also have high torque density and lower cost [10].

The main disadvantages are control complexity and low power factor, due to the fact that the machine magnetization is entirely supplied by the stator [10].

2.1.5 Induction Motors

In an induction motor, the three-phase stator winding produces a magnetic field rotating at synchronous speed, exactly as in synchronous AC machines. The most common rotor is the squirrel cage, consisting of conducting bars short-circuited at both ends by end rings, with no electrical connection to the outside. The relative motion between the rotating field and the rotor conductors induces voltages, and therefore currents, in the rotor, whose interaction with the stator field produces torque. Since no torque is produced without this relative motion, the rotor always runs slightly below synchronous speed, with the difference expressed by the slip — hence its designation as an asynchronous machine [3].

As the dominant machine in industrial applications, the induction motor is the most mature and widespread of the candidate technologies, and has been adopted in several commercial electric vehicles [11]. In fact, its reliability, low cost, low maintenance and robustness make it, for some authors, the most suitable option for electric propulsion systems [6].

Induction motors have good reliability and low cost because there are no magnets in the rotor; however, this makes them less efficient when compared to motors with permanent magnets, since the rotor flux must be created by a magnetizing current drawn from the stator and the induced rotor currents produce additional Joule losses.

Torque and flux may be independently controlled with decoupled vector control techniques. Flux-weakening is also possible to increase the operating speed range over the rated maximum motor speed [3].

On the other hand, at higher speeds, IMs have higher losses and consequently low efficiency. The breakdown torque limit in the constant power region is another main disadvantage when compared to permanent magnet machines [6].

2.1.6 Permanent Magnet Synchronous Motors

PMSMs are one of the most common options in electric vehicles [11]. The rotor is composed of permanent magnets and the stator carries, normally, a three-phase winding, fed with sinusoidal currents synchronous with the rotor position.

The rotation principle is based on the interaction between the field of the rotor magnets and the rotating magnetic field of the stator. When the three stator phases are fed with sinusoidal currents displaced by 120°, the resulting magnetomotive force is a field of constant amplitude that rotates at synchronous speed, even though the stator itself is stationary. The field of the permanent magnets tends to align with this rotating field, so the rotor locks to it and rotates synchronously, with torque being produced by the interaction between the two fields, as illustrated in figure 2.3 [12].

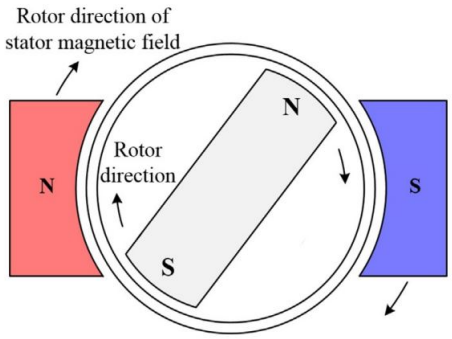


Figure 2.3: Rotation principle of permanent magnet synchronous motor [12]

The fact that the rotor flux position is the position of the rotor itself makes the vector control of these motors relatively simple. The conditions that torque control of a PMSM has to meet are the following:

- The stator current vector must be oriented with respect to the rotor flux at the angle imposed by the control strategy, which requires the rotor position to be known at every instant;
- The components of the current responsible for the flux and for the torque must be controllable independently of one another;
- The current must be regulated with a bandwidth well above the rate at which the torque demand and the operating point change.

As the field flux is produced by the permanent magnets and the torque is produced by the stator current, these may be independently controlled. The last condition is met by the combination of a current regulator in the synchronous reference frame with a PWM inverter, whose achievable bandwidth is limited by the switching frequency and by the delays of the digital implementation, as examined in section 3.4.1. The first condition is met by transforming the three-phase stator currents, from sinusoidal *abc* into two d-q axes currents, where the d-axis is aligned with the position of the rotor flux and the q-axis is orthogonal to it. The currents are firstly transformed into a stationary frame and then into a rotating reference frame that rotates along with the rotor, which means its position must be known, either by sensors (encoder or resolver) or estimation [3]. In a machine without saliency the torque is produced solely by the q-axis current, and the optimum orientation is therefore the one that keeps the current vector orthogonal to the rotor flux; in a salient machine, as discussed below and in section 3.5, the optimum departs from that orientation.

The motors may be of different topologies, essentially varying on the rotor structure [3]. When the magnetic flux crosses the air gap in the axial direction, parallel to the shaft, it is an axial-flux configuration; when the flux crosses the air gap radially, with the rotor typically rotating inside the stator, it is a radial-flux configuration, this one being the most common. Figure 2.4 shows the difference between the two configurations.

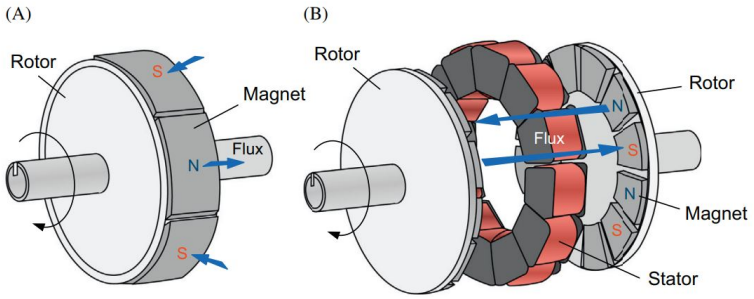


Figure 2.4: (A) Radial-flux type and (B) axial-flux type [3]

Radial-flux configurations may also differ in terms of the position of the permanent magnets: these may be surface mounted, when the magnets are fixed on the surface of the rotor, or interior mounted, when they are buried inside the rotor, in a parallel or perpendicular orientation, as shown in figure 2.5.

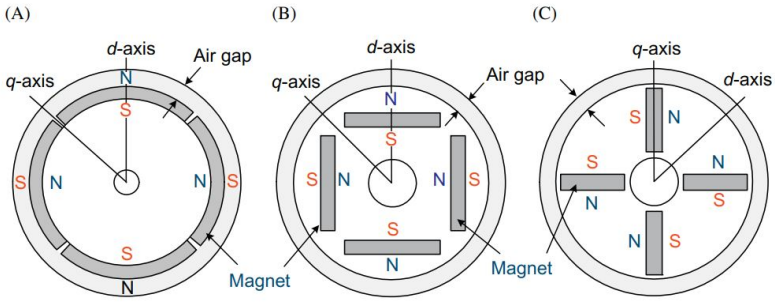


Figure 2.5: (A) Surface mounted, (B) interior (parallel), and (C) interior (perpendicular) [3]

The position of the magnets has a direct impact on the electrical behaviour of the machine. In surface mounted PMSMs, since the permeability of the magnets is close to that of air, the effective air gap is large and uniform, so the d- and q-axis inductances are practically equal. In interior PMSMs, the magnets lie in the d-axis flux path, which lowers the d-axis inductance and creates a magnetic saliency ($L_q > L_d$). This saliency produces an additional reluctance torque component and improves the flux-weakening capability of the machine, extending its constant power speed range; the buried magnets are also mechanically protected, making the interior configuration better suited for high speed operation [3]. The presence of this reluctance component is also what makes the choice of current references a design problem in its own right, since the required torque no longer determines a single pair of currents, a question addressed in section 3.5.

Although the radial-flux configuration is defined as the rotor rotating inside the stator, which is called the inner rotor type, the rotor may also rotate outside the stator, in the outer rotor type. The main difference is that the inertia is higher in outer rotor type motors, making them favorable for constant speed operation, while the inner rotor type is better suited for a quick speed response [3].

PMSMs have many advantages: firstly, as the rotor flux is provided by the magnets and not generated by an external source, there are no rotor excitation losses, which makes these machines highly efficient; moreover, they have high torque and power density, as they are normally very compact; these motors also offer a wide range of power ratings, making them useful for various applications [6].

As these need rare-earth permanent magnets on the rotor, the main disadvantage is the high cost, although this has been decreasing over time. There is also the risk of demagnetization of the magnets, caused by excessive heat or armature reaction [6].

The objectives of the application in question make this motor the most suitable, because of its main advantages of torque and power density and efficiency. The machine used in this work is of the interior magnet type ($L_q > L_d$), so the saliency described above is present and is taken into account throughout the control design.

2.2 Power Converters

As said before, the controller shall be implemented on an already designed two-level three-phase topology converter. However, this section compares the most common topologies used in PMSM drives for electric vehicles [13], shown in figure 2.6.

Regardless of the topology, the switching devices are typically silicon IGBTs or, increasingly, silicon carbide (SiC) MOSFETs. Since it is wide bandgap semiconductor, SiC supports higher blocking voltages, higher switching frequencies and lower switching losses than silicon, resulting in more efficient and more compact traction inverters [14]. The converter used in this work employs SiC MOSFETs switching at 40 kHz.

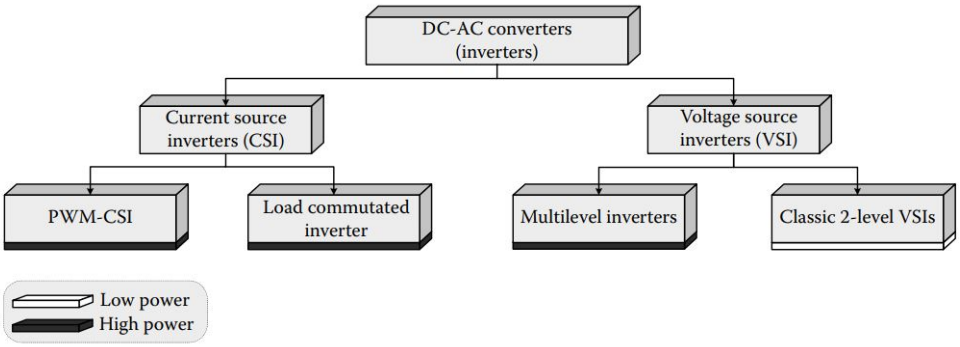


Figure 2.6: Inverter topologies [13]

2.2.1 Two-level three-phase topology

This is the most used topology among formula student vehicles, mainly for its simplicity [2], and is shown in figure 2.7. It consists of three parallel legs, one per phase, each with two switches and its own binary control signal. Within each leg, the gate signals are complementary, since turning both switches on would short-circuit the DC link. For this reason, a small dead time is inserted between the turn-off of one switch and the turn-on of the other, during which the output voltage is imposed by the freewheeling diodes, according to the direction of the load current [13].

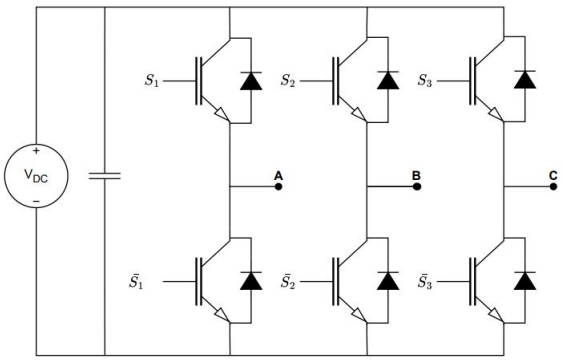


Figure 2.7: Two-level 3-phase [2]

As the name indicates, this is a two-level topology, which means that each phase output, measured with respect to the negative terminal of the DC link, may have two values: 0 or V_{DC}. Being controlled by three binary signals, the result is 8 (2^3) possible switching states, listed in table 2.1: six active states, which apply voltage to the machine, and two zero states (states 1 and 8), in which the three outputs are connected to the same DC link terminal and the line-to-line voltages are zero. These eight states are the basis of the space-vector modulation technique, described in section 2.3.2 [13].

Table 2.1: Switching states of the two-level three-phase converter

State	**Switches**			**Output Voltage**		
	S_1	S_2	S_3	V_{a-n}	V_{b-n}	V_{c-n}
1	0	0	0	0	0	0
2	1	0	0	V_{DC}	0	0
3	1	1	0	V_{DC}	V_{DC}	0
4	0	1	0	0	V_{DC}	0
5	0	1	1	0	V_{DC}	V_{DC}
6	0	0	1	0	0	V_{DC}
7	1	0	1	V_{DC}	0	V_{DC}
8	1	1	1	V_{DC}	V_{DC}	V_{DC}

2.2.2 Three-level three-phase topology

In this inverter topology each phase output may assume three values: 0, $V_{DC}/2$ or V_{DC}, with the intermediate level obtained from the midpoint of a split DC link. The main topologies of the three-level converter are the neutral-point-clamped (NPC) and the neutral-point-piloted (NPP) [15], shown in figure 2.8.

On the NPC topology, each arm of the leg has two series connected semiconductors, clamped to the DC link midpoint by diodes. Since each device only has to block half of the DC link voltage, this topology is able to work with higher DC voltages.

On the other hand, the NPP topology has two anti-series-connected semiconductors connecting the output directly to the DC link midpoint for voltage clamping.

Compared with the two-level converter, the additional output level produces voltage waveforms with lower harmonic distortion and halves the voltage steps applied to the machine, reducing the filtering needs and the stress on the winding insulation. This comes at the cost of a larger number of semiconductors and gate drivers, a more complex modulation, and the need to keep the two DC link capacitor voltages balanced [13], [15]. Other multilevel families, such as the flying capacitor and the cascaded H-bridge converters, extend this principle to a higher number of levels, but are mostly used in high power and high voltage applications, not being common in electric vehicle drives [13].

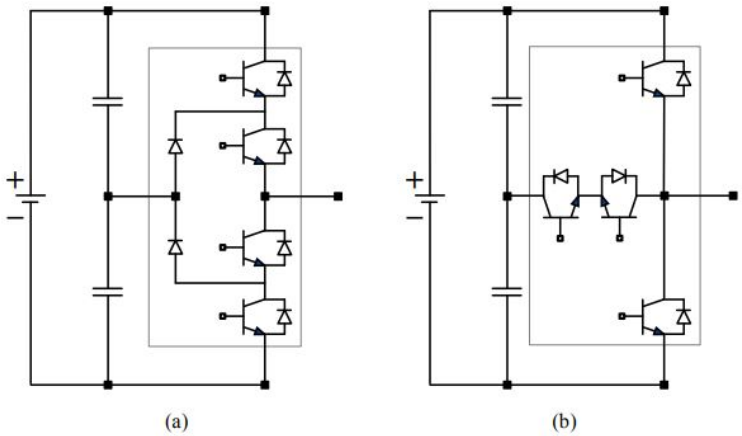


Figure 2.8: NPC (a) and NPP (b) [15]

2.2.3 Open-Winding

In the open-winding configuration, the star point of the machine is opened and both terminals of each stator phase are made accessible, so that each winding is fed from both ends by two two-level inverters. The voltage applied to each phase is the difference between the outputs of the two inverters, which means the winding can be excited with up to the full DC link voltage in either polarity, and the combined converter produces output waveforms equivalent to those of a three-level inverter [15].

The two inverters may be supplied by two isolated DC sources or by a single, shared DC link, the latter shown in figure 2.9. The single source solution is simpler and lighter, but creates a closed path for zero-sequence currents, which must be suppressed by the modulation strategy or by additional filtering. The main advantages of this topology are the better utilization of the DC link voltage and its inherent fault tolerance, since a faulty inverter can be bypassed and the machine kept running from the healthy side. The main drawbacks are the doubled number of semiconductors and gate drivers and the increased control complexity [15].

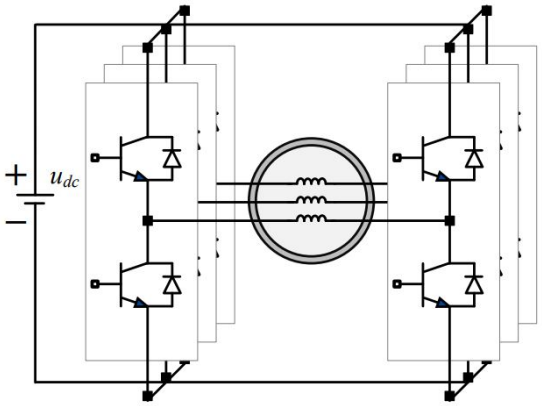


Figure 2.9: Single source open-winding [15]

2.2.4 T-type

The T-type converter is a three-level topology derived from the two-level leg, closely related to the NPP concept presented above. Each phase keeps the two main switches connected between the positive and negative DC link terminals, and adds a bidirectional switch — two anti-series-connected devices — linking the output to the midpoint of the split DC link, forming the characteristic "T" shape, as shown in figure 2.10. The output can therefore be connected to V_{DC}, $V_{DC}/2$ or 0, producing the same three-level waveform as the NPC converter with fewer semiconductors, since no clamping diodes are required [16].

Unlike the NPC, the two main switches must block the full DC link voltage, while the midpoint devices only block half of it. Since the conduction path for the upper and lower levels involves a single device, the T-type combines the low conduction losses of the two-level converter with the low switching losses and superior output quality of three-level operation, making it an attractive solution for low voltage, high efficiency applications at moderate switching frequencies [16].

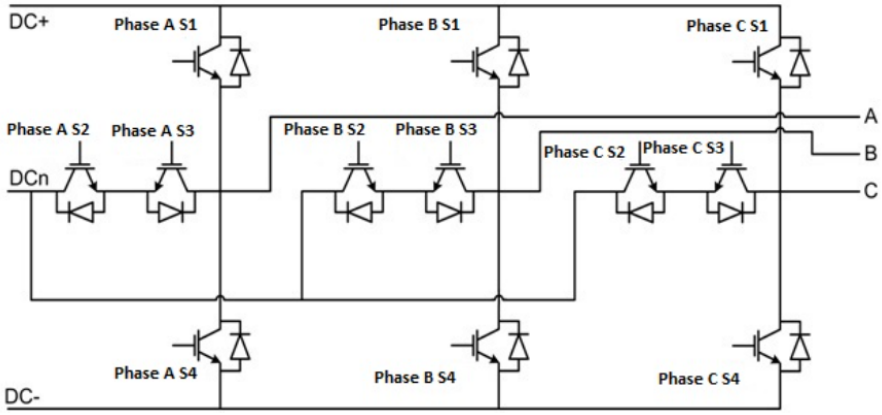


Figure 2.10: Three-phase T-type topology [17]

Table 2.2 summarizes the main characteristics of the presented topologies. For the voltage and power levels of a Formula Student powertrain, the simplicity and reliability of the two-level converter keep it as the dominant solution [2].

Table 2.2: Comparison of the presented inverter topologies

Topology	**Levels**	**Switches**	**Clamping diodes**	**Blocking voltage**
Two-level	2	2	0	V_{DC}
NPC	3	4	2	$V_{DC}/2$
NPP / T-type	3	4	0	V_{DC} and $V_{DC}/2$
Open-winding	3	4	0	V_{DC}

Switches and clamping diodes are given per phase.

2.3 Modulation Techniques

The control techniques described in section 2.4 provide the voltage references to be applied to the machine. The modulation stage translates these references into the binary gate signals of the inverter switches, so that the average voltage applied to the machine over each switching period matches the reference. The chosen modulation technique determines how well the DC link voltage is utilized, the harmonic content of the output currents and the number of commutations, directly affecting the losses of the drive [13].

2.3.1 Sinusoidal PWM

Sinusoidal PWM is the simplest modulation technique. For each phase, a sinusoidal reference with the desired fundamental frequency and amplitude is compared with a high frequency triangular carrier common to the three phases: while the reference is higher than the carrier, the upper switch of the leg conducts, and while it is lower, the lower switch conducts. The result is a train of pulses whose average value reproduces the sinusoidal reference, with the harmonic content concentrated around multiples of the carrier frequency, as shown in figure 2.11. Since these components lie far above the fundamental, they are naturally attenuated by the inductance of the machine windings, so the resulting currents are practically sinusoidal [13], [18].

Its main limitation is the DC link utilization: in the linear operating region, the maximum amplitude of the fundamental phase voltage is $V_{DC}/2$. This limit can be extended by injecting a zero-sequence (third harmonic) component into the references, which does not appear in the line-to-line voltages applied to the machine [13].

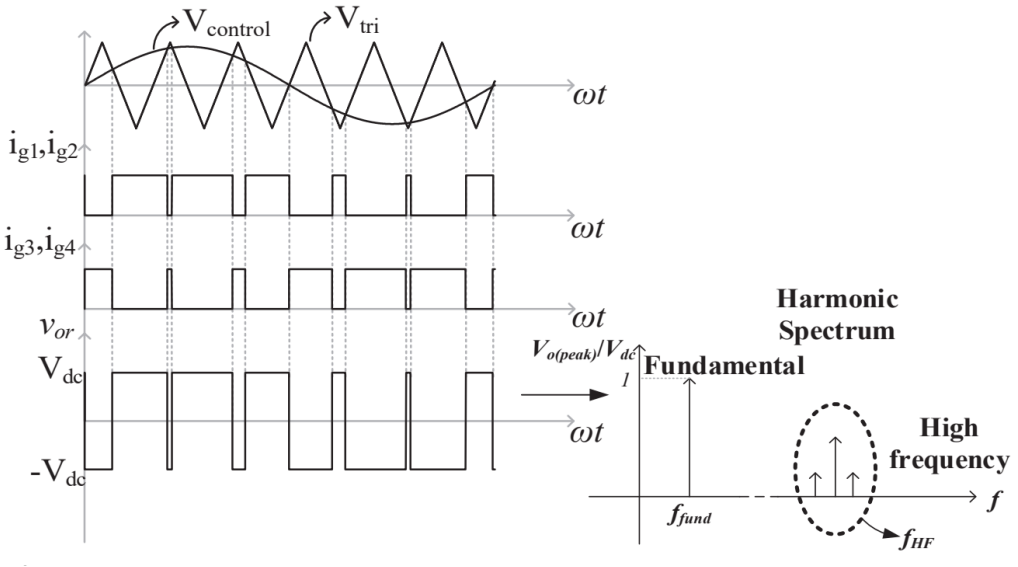


Figure 2.11: Sinusoidal PWM operating principle and corresponding harmonic spectrum [18]

2.3.2 Space-Vector Modulation

Space-vector modulation treats the inverter as a whole, instead of modulating each phase independently. The eight switching states of the two-level converter, presented in table 2.1, are represented as voltage vectors in the stationary $\alpha\beta$ plane: the six active vectors form a hexagon, divided into six sectors, and the two zero vectors lie at its origin, as shown in figure 2.12. The voltage reference produced by the controller is a rotating vector in this plane and, in each switching period, it is synthesized by applying the two active vectors adjacent to its sector, together with the zero vectors, during appropriately calculated fractions of the period, so that the average output voltage equals the reference [19].

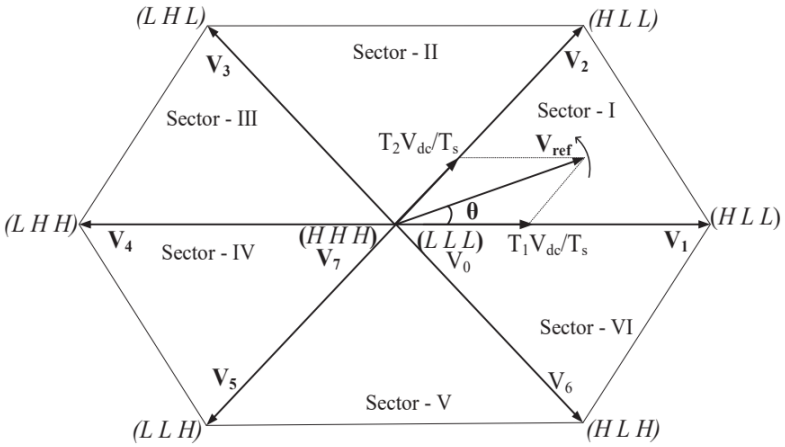


Figure 2.12: Space-vector diagram of the two-level converter [18]

When compared with SPWM, SVM increases the maximum fundamental phase voltage in the linear region to $V_{DC}/\sqrt{3}$, approximately 15% higher, produces lower current harmonic distortion for the same switching frequency, and gives direct control over the switching sequence, allowing the number of commutations to be minimized [19]. Moreover, since the technique operates directly on the $\alpha\beta$ representation of the reference voltage, it integrates naturally with the reference frame transformations used by vector control, making it well suited for digital implementation together with FOC [13]. SVM can also be realized in a carrier-based fashion, by comparing a triangular carrier with equivalent modulating signals, shown in figure 2.13.

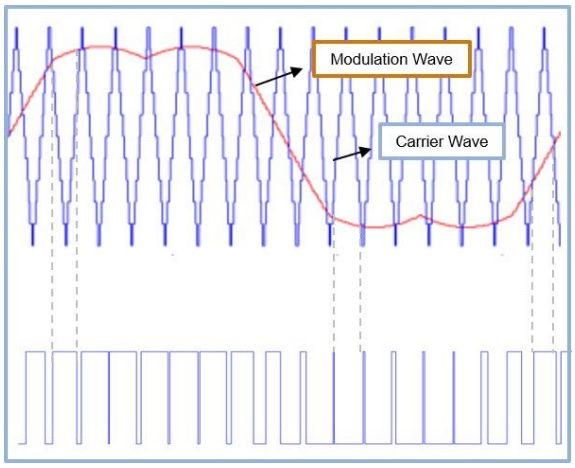


Figure 2.13: Space-vector PWM: modulating signal, triangular carrier and resulting gate pulses [20]

The way the zero vectors are distributed within each switching period implicitly adds a zero-sequence component to the sinusoidal references — equivalent to a third harmonic injection — giving the modulating signals their characteristic double-hump shape. Since the machine windings are star-connected, without a neutral conductor, this zero-sequence component drives no current and does not appear in the line-to-line voltages applied to the machine; its effect is to flatten the peaks of the modulating signals, which is what allows the fundamental component to be increased by the aforementioned 15% before entering overmodulation [20]. Furthermore, the lower harmonic distortion of the output currents reduces the harmonic losses induced in the machine, and the optimized switching sequence, with a single commutation per state transition, reduces the switching losses of the converter, improving the overall efficiency of the drive [19].

For these reasons — higher DC link utilization, lower current distortion and natural integration with the adopted control technique — SVM is the modulation technique selected for this work. Its detailed operation and implementation are described in the methodology chapter.

2.4 Control Techniques

The main focus of this dissertation is the motor control. The applicable techniques may be divided into scalar control, which acts only on the magnitude and frequency of the applied voltages, and vector control, which acts on the instantaneous position of the voltage and current space vectors. Figure 2.14 shows the main vector control techniques for PMSM applications.

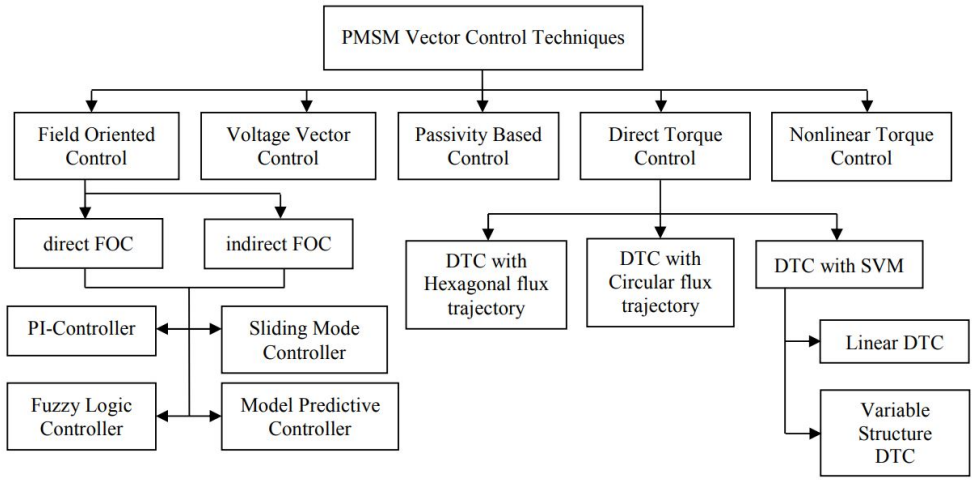


Figure 2.14: Control techniques for PMSM drives [21]

This section addresses the main techniques applicable to the system in question.

2.4.1 Scalar Control (V/f)

Scalar control, commonly known as V/f control, is one of the simplest ways of driving a PMSM. The machine is fed in open loop with sinusoidal voltage references whose amplitude varies proportionally to the applied frequency, keeping the V/f ratio, and therefore the stator flux, approximately constant. Since it is based only on the steady-state model of the machine, no position or speed feedback and no current controllers are required, which makes its implementation very simple and computationally light [22]. Figure 2.15 shows the structure of this technique.

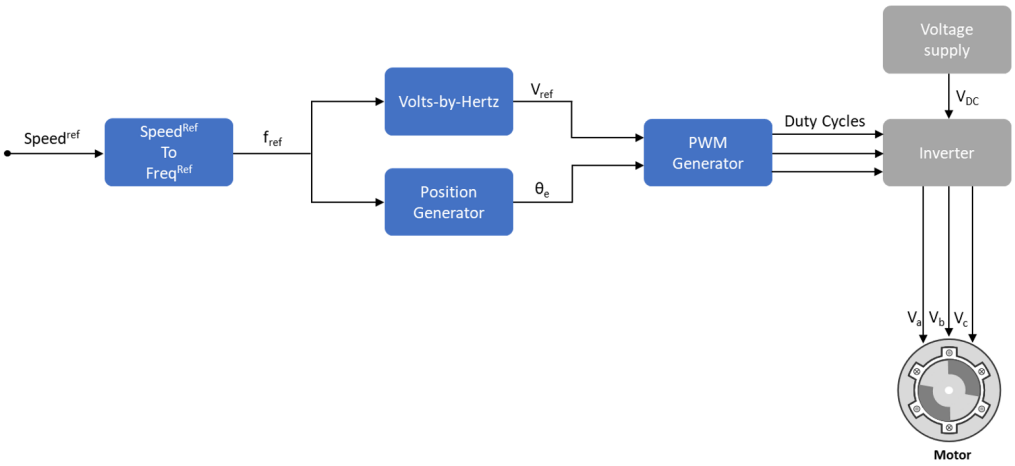


Figure 2.15: Open-loop V/f control structure [23]

The main limitations are the poor dynamic performance and the absence of direct torque control. Moreover, PMSMs without damper windings are inherently unstable in open-loop V/f operation above a certain applied frequency, so practical drives usually require additional stabilization loops [22]. For these reasons, V/f control is mostly restricted to low performance applications. In

this work it is used only as an open-loop method for the initial validation of the motor and inverter models.

2.4.2 Direct Torque Control

DTC is a relatively simple control technique that uses the measured motor voltages and currents, from which the stator flux linkage and the electromagnetic torque are estimated. Both estimates are compared with their references in hysteresis controllers, whose outputs, together with the angular sector of the stator flux, select the inverter switching state directly from a predefined switching table [24]. Since the switching states are selected directly, no modulator and no inner current control loops are needed, and the rotor position sensor is also unnecessary. Figure 2.16 shows the control diagram of this technique [25].

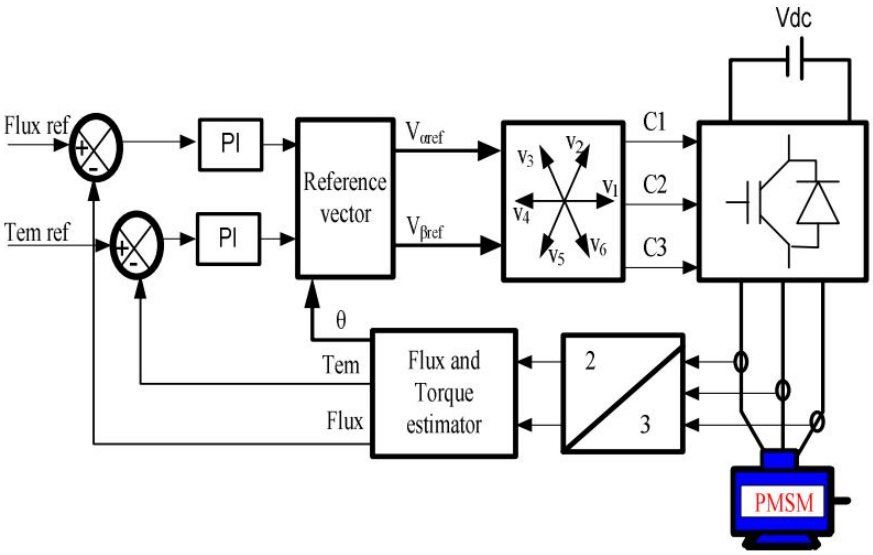


Figure 2.16: Direct torque control [25]

This method has some advantages such as circuit and structure simplicity, no need for advanced reference frame transformations, which reduces the computational capacity needs, has a fast dynamic response and a good load adaptability.

However, the hysteresis-based selection of the switching states results in a variable switching frequency and in high torque and flux ripple in steady-state operation, making it a bad option for high precision applications [21].

2.4.3 Field-Oriented Control

The principle of FOC that makes it different from other control techniques is the transformation of the three-phase stator currents into a two-phase reference frame aligned with the rotor flux, the *dq* reference frame. The d-axis corresponds to the direction of the magnetic flux of the rotor magnets and the q-axis, also called the torque-producing axis, is perpendicular to it, with the current along

this axis (i_q) directly determining the motor torque [26]. In this synchronous frame, the currents become DC quantities in steady state, which allows simple PI controllers to regulate them without steady-state error; the voltage references produced by these controllers are then synthesized by the modulator, typically SVM, as presented in section 2.3.2 [3]. Figure 2.17 shows the control diagram of this technique.

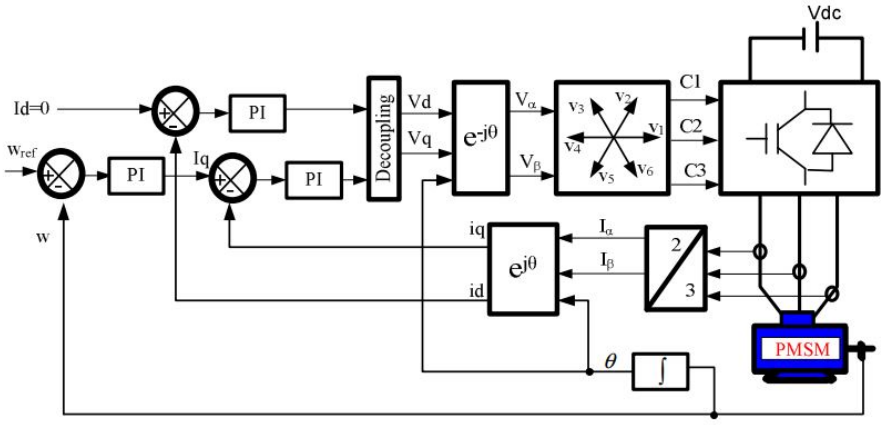


Figure 2.17: Field-oriented control [25]

The change of reference frame is carried out in two stages. The Clarke transformation first reduces the three phase quantities to two orthogonal components in a frame that remains fixed to the stator, the $\alpha\beta$ frame, and the Park transformation then rotates those components into the frame attached to the rotor [3].

The Clarke transformation maps the three phase currents onto the stationary $\alpha\beta$ axes,

$$\begin{bmatrix} i_\alpha \\ i_\beta \end{bmatrix} = \frac{2}{3} \begin{bmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & \frac{\sqrt{3}}{2} & -\frac{\sqrt{3}}{2} \end{bmatrix} \begin{bmatrix} i_a \\ i_b \\ i_c \end{bmatrix}, \tag{2.1}$$

where the factor $2/3$ makes the transformation amplitude invariant, so that the two-axis quantities retain the peak value of the phase quantities. In a balanced three-wire machine the phase currents sum to zero, so the α axis coincides with phase a and only two of the three currents need to be measured. The two components obtained are still alternating quantities, at the electrical frequency of the machine.

The Park transformation rotates them into the frame that turns with the rotor,

$$\begin{bmatrix} i_d \\ i_q \end{bmatrix} = \begin{bmatrix} \cos\theta_e & \sin\theta_e \\ -\sin\theta_e & \cos\theta_e \end{bmatrix} \begin{bmatrix} i_\alpha \\ i_\beta \end{bmatrix}, \tag{2.2}$$

where θ_e is the electrical position of the rotor, obtained from the measured mechanical position through $\theta_e = p\,\theta_m$, with p the number of pole pairs. Because this frame rotates synchronously with the rotor flux, the alternating quantities of the $\alpha\beta$ frame become constant in steady state, which is

the property that allows PI controllers to regulate them without steady-state error and is the reason the rotor position must be known at every instant.

The reverse path is followed by the voltage references produced by the controllers. The inverse Park transformation returns them to the stationary frame,

$$\begin{bmatrix} v_\alpha \\ v_\beta \end{bmatrix} = \begin{bmatrix} \cos\theta_e & -\sin\theta_e \\ \sin\theta_e & \cos\theta_e \end{bmatrix} \begin{bmatrix} v_d \\ v_q \end{bmatrix}, \tag{2.3}$$

and the inverse Clarke transformation restores the three phase references,

$$\begin{bmatrix} v_a \\ v_b \\ v_c \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -\frac{1}{2} & \frac{\sqrt{3}}{2} \\ -\frac{1}{2} & -\frac{\sqrt{3}}{2} \end{bmatrix} \begin{bmatrix} v_\alpha \\ v_\beta \end{bmatrix}, \tag{2.4}$$

which are the references delivered to the modulator. The choice of the amplitude invariant scaling in (2.1) is carried through consistently, and is what fixes the factor $3/2$ in the torque expression of the machine.

The decoupling of torque and flux of the motor makes FOC a very precise control technique with a fast response to input at load changes. This control method also contributes to the system efficiency by precisely optimizing the current distribution and, in consequence, reducing losses. It can operate on a wide speed range without losing its performance and precision, and on various different applications.

Some disadvantages of FOC are the high computation requirements, due to its complex algorithm and mathematical operations, and the need for a position sensor (although it may be implemented with a sensorless algorithm). Both these factors increase the cost of implementation of this type of controller. FOC may also be sensitive to motor parameter changes or external disturbances [21].

Given the requirements of the application in question — precise torque control, high efficiency and a wide speed range — FOC, combined with the SVM modulation presented in the previous section, is the control technique adopted in this work.

2.4.4 Current Reference Strategies

In a FOC scheme, the same electromagnetic torque may be produced by different combinations of i_d and i_q, so a strategy is needed to translate the torque request into the pair of current references to be tracked by the current controllers. The relevance of this choice becomes clear from the electromagnetic torque expression of the PMSM [3]:

$$T_e = \frac{3}{2} p \left[\psi_{PM} i_q + (L_d - L_q) i_d i_q \right] \tag{2.5}$$

where p is the number of pole pairs and ψ_{PM} the permanent magnet flux linkage. The first term is the magnet torque, produced by the interaction between the magnet flux and i_q; the second is the reluctance torque, which only exists in machines with magnetic saliency ($L_d \neq L_q$), as discussed in section 2.1.

The simplest strategy is to impose $i_d = 0$, aligning the stator current vector with the q-axis. The torque becomes directly proportional to i_q, resulting in a simple and linear torque control, and the absence of a demagnetizing d-axis current component protects the magnets. For surface mounted PMSMs, where $L_d \approx L_q$ and the reluctance term vanishes, this strategy also produces the required torque with the minimum stator current [3].

In salient machines, however, $i_d = 0$ leaves the reluctance torque unused. Since $L_q > L_d$ in an interior PMSM, a negative i_d makes the reluctance term positive, adding to the magnet torque. The maximum torque per ampere (MTPA) strategy exploits this by computing, for each torque request, the combination of i_d and i_q that produces the required torque with the minimum stator current magnitude. Operating with minimum current at every torque level minimizes the copper losses and the inverter current requirements, maximizing the overall efficiency, which makes MTPA the standard strategy for interior magnet traction drives [27]. Figure 2.18 illustrates this concept: among all the current vectors capable of producing a given torque, the MTPA operating point (M) is the one with the minimum magnitude. Its main difficulty is that the optimal current angle depends on the machine parameters (ψ_{PM}, L_d, L_q), which vary with magnetic saturation, so accurate parameter knowledge is required.

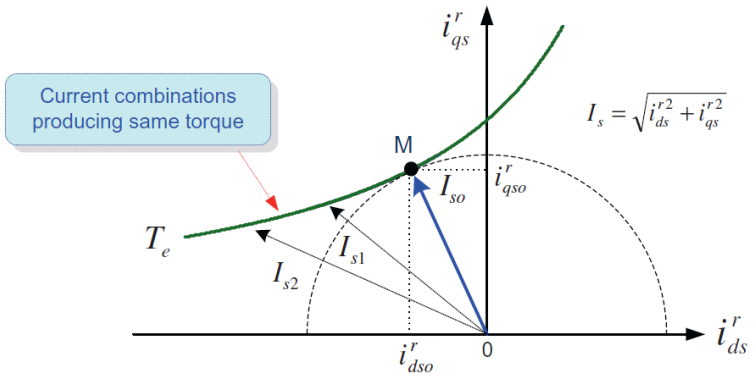


Figure 2.18: Current combinations producing the same torque and the MTPA operating point [3]

Above the base speed, the voltage limit of the inverter prevents these strategies from being followed, and the current references must be modified to weaken the machine flux, in the flux-weakening operation region, extending the speed range within the voltage and current limits of the drive [27]. Since the motor used in this work is of the interior magnet type ($L_q > L_d$), both the $i_d = 0$ and MTPA strategies are directly relevant to the control implemented in this dissertation.

Chapter 3

System Modelling and Controller Design

3.1 Modelling of the Permanent Magnet Synchronous Motor

The specific motor utilized in this research, the TI085-052-070-04B7S-07S04BE2 manufactured by Fischer Elektromotoren GmbH, was explicitly selected by the Formula Student FEUP team. This particular model was chosen because its performance characteristics and technical specifications align with the team's engineering objectives. The available data used to parameterize the motor model consisted of the technical datasheet A and a spreadsheet B with the results of a finite element analysis (FEA). The first one provided the nominal electromechanical ratings and basic constant electrical parameters of the machine, such as the number of pole pairs, the stator resistance per phase R_s, the zero-sequence inductance L_0. The second source provided the non-linear data - namely, the saturation of the inductances (which can be seen on figures 3.2a and 3.2b) and the cross-coupling between the d- and q-axes - used to make a more accurate model. The spreadsheet consists of a set of operating points spanning the motoring quadrant of the current plane ($i_d \leq 0$, $i_q \geq 0$). The L_d and L_q inductances and the permanent magnet flux linkage ψ_{PM} were the main dependent parameters used from the spreadsheet.

3.1.1 FEM-Parameterized PMSM block

Unlike the conventional *PMSM* block from the Simscape Electrical library and the *Interior PMSM* from the *Powertrain Blockset* library, which represent the machine with constant, lumped electrical parameters, the *FEM-Parameterized PMSM* block defines the motor behaviour through tabulated look-up tables of the d- and q-axis inductances and the permanent magnet flux linkage as functions of the d- and q-axis currents through table lookup, enabling a nonlinear, current-dependent representation of the machine's magnetic circuit [28]. This aligns with the FEA characterization data provided by the manufacturer.

Among the block's four modelling options, the *2-D partial derivative data* option with the *Assume sinusoidal back emf — tabulate with d- and q-axis currents* parameterization is the one whose

expected inputs match the structure of the manufacturer's characterization data. The manufacturer's file provides, for each operating point in the current plane, the d- and q-axis inductances L_d and L_q, the permanent magnet flux linkage ψ_{PM} (*FluxM_eff* on the spreadsheet), and the d- and q-axis flux linkages ψ_d and ψ_q. The two representations are equivalent, related by $\psi_d = L_d\, i_d + \psi_{PM}$ and $\psi_q = L_q\, i_q$, and both describe a purely two-dimensional dependence on (i_d, i_q): the file contains a single value of each quantity per operating point, without any tabulated dependence on the rotor angle θ_r. This immediately rules out the block's remaining modelling options, which all require an explicit angular dependence that the available data does not provide: the *3-D partial derivative data* and *3-D flux linkage data* options tabulate against (i_d, i_q, θ_r), and the *4-D partial derivative data* option additionally works in the phase-current frame. Adopting any of these would have required synthesising the angular dependence under an assumption of sinusoidal back-EMF — precisely the assumption that the 2-D option already builds in. In contrast, the 2-D option requires only $L_d(i_d, i_q)$, $L_q(i_d, i_q)$ and $\psi_{PM}(i_d, i_q)$, which are directly available in the manufacturer's file. According to the MathWorks documentation, the sinusoidal back-EMF assumption fits interior magnet PMSMs well and is therefore appropriate for the target machine. Under this option, the block still captures the two most important nonlinear effects of the magnetic circuit: magnetic saturation (variation of the parameters with current magnitude) and cross-coupling between axes (dependence of the d-axis parameters on i_q and vice versa).

Under this parameterization, the block solves the standard PMSM equations in the rotor-oriented (synchronous) dq reference frame. The flux linkages of the two axes are defined from the tabulated maps as

$$\begin{bmatrix} \psi_d \\ \psi_q \end{bmatrix} = \begin{bmatrix} L_d(i_d, i_q) & \\ & L_q(i_d, i_q) \end{bmatrix} \begin{bmatrix} i_d \\ i_q \end{bmatrix} + \begin{bmatrix} \psi_{PM}(i_d, i_q) \\ 0 \end{bmatrix} \tag{3.1}$$

and the electromagnetic torque produced by the machine is

$$T = \frac{3}{2}\, p(i_q(i_d\, L_d(i_d, i_q) + \psi_{PM}(i_d, i_q)) - i_d\, i_q\, L_q(i_d, i_q)) \tag{3.2}$$

where ψ_d and ψ_q are the d- and q-axis flux linkages, ψ_{PM} is the permanent magnet flux linkage, p is the number of pole pairs, and T is the electrical torque.

Figure 3.1 shows the steady-state phasor diagram of an interior PMSM on the dq reference frame.

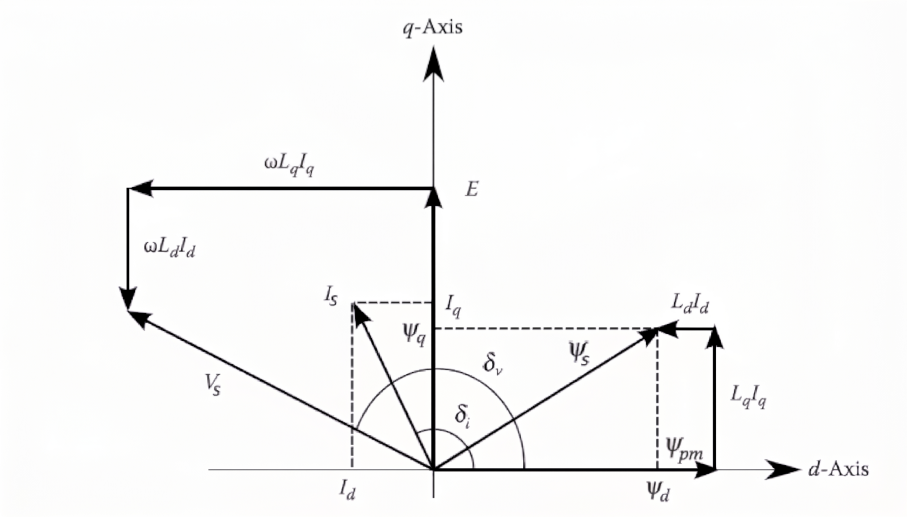


Figure 3.1: Interior PMSM steady-state phasor diagram

At each simulation time step, the block operates as follows. Given the instantaneous three-phase currents at its electrical ports (motor phases) and the rotor angle at its mechanical port (motor shaft), it first applies the Park transform to obtain the currents i_d and i_q in the rotor-oriented frame, using the internal convention in which the q-axis leads the d-axis and the rotor angle is measured from the A-phase magnetic axis to the d-axis. It then evaluates the three tabulated maps at the current operating point (i_d, i_q) by two-dimensional interpolation, retrieving instantaneous values of L_d, L_q and ψ_{PM} consistent with the local level of magnetic saturation. These values are substituted into equation 3.1 to obtain the flux linkages, from which the terminal voltages imposed at the electrical ports are computed via the standard PMSM voltage equations in the dq frame. In parallel, equation 3.2 yields the electromagnetic torque, which is applied at the mechanical port and drives the rotor dynamics.

The block requires two current vectors, i_d and i_q, defining a regular rectangular grid covering all four quadrants of the current plane; three matrices, $L_d(i_d, i_q)$, $L_q(i_d, i_q)$ and $\psi_{PM}(i_d, i_q)$, with dimensions consistent with those vectors, evaluated at every grid point; and all quantities expressed as peak (amplitude) values, consistent with the $3/2$ factor in equation 3.2. Note that the electromagnetic torque is not supplied as a tabulated input: the three maps L_d, L_q and ψ_{PM} fully determine the block's behaviour, and the torque is derived from Equation 3.2 at run time. In addition, the block accepts scalar values for the stator resistance R_s, the zero-sequence inductance L_0, and the number of pole pairs N, which are constant throughout the simulation and taken directly from the motor technical datasheet.

The spreadsheet data does not directly meet these requirements, and a number of incompatibilities must be resolved before it can be loaded into the block. First, the block expects L_d, L_q and ψ_{PM} to be defined on a *regular rectangular grid* — that is, the current vectors i_d and i_q must be uniformly spaced and monotonically increasing, and the matrices must contain a value at every node of the resulting grid. The manufacturer's FEA data, by contrast, consists of *scattered* operating points in the (i_d, i_q) plane: the simulation campaign did not follow a regular spacing, and

many combinations of i_d and i_q are absent from the dataset. Second, the FEA data covers only the motoring quadrant ($i_d \leq 0$, $i_q \geq 0$), whereas the block requires maps that span all four quadrants of the current plane. Third, all quantities in the manufacturer's file — currents and permanent magnet flux linkage alike — are expressed as *RMS* (effective) values, while the block operates internally with *peak* (amplitude) values; using RMS data directly would underestimate the flux linkage by a factor of $\sqrt{2}$ and lead to inconsistent results. A dedicated MATLAB script was therefore developed to address each of these incompatibilities and transform the original data into the format required by the block, as described in the following section.

3.1.2 Processing of the finite element data

After looking at the manufacturer's FEA file, the relevant columns are extracted:

```
i_d_rms = dados(:, 3);    % Id_eff [A] (RMS)
i_q_rms = dados(:, 4);    % Iq_eff [A] (RMS)
Ld_raw  = dados(:, 5);    % Ld [H]
Lq_raw  = dados(:, 6);    % Lq [H]
PM_rms  = dados(:, 7);    % FluxM_eff [Vs] (RMS)
```

Listing 3.1: Extraction of the FEM data columns.

The first processing step is the RMS-to-peak convention conversion. As established in Section 3.1.1, the block operates with peak (amplitude) values, consistent with the $3/2$ factor in the torque equation. The currents and the magnet flux linkage are therefore multiplied by $\sqrt{2}$. The inductances, being defined as ratios of flux linkage to current ($L = \psi/i$), are invariant under this scaling and require no conversion:

```
i_d    = SQRT2 * i_d_rms;
i_q    = SQRT2 * i_q_rms;
PM_raw = SQRT2 * PM_rms;
```

Listing 3.2: RMS-to-peak conversion.

The regular rectangular grid required by the block is constructed. A symmetric 51×51 mesh covering all four quadrants of the current plane was generated. The odd dimension was deliberately chosen to guarantee that the axes $i_d = 0$ and $i_q = 0$ fall exactly on grid nodes, with a node placed at the origin $(0,0)$. The mirroring procedure described below reflects the raw data about these axes, and this reflection introduces a slope discontinuity precisely along them. When grid nodes coincide with the axes, the discontinuity is captured by tabulated values and read directly by the block; an even grid dimension would place the axes between nodes, forcing the block's smooth interpolator to average across the discontinuity and distorting the maps near the origin — precisely the region through which the machine passes most frequently.

```
resolution = 51;                    % odd -> guarantees 0 at the centre
id_vec = linspace(-max_id, max_id, resolution);
```

```
iq_vec = linspace(-max_iq, max_iq, resolution);
[Id_grid, Iq_grid] = ndgrid(id_vec, iq_vec);
```

Listing 3.3: Construction of the symmetric interpolation grid.

The *ndgrid* function is used instead of *meshgrid* because the block expects the first dimension to correspond to i_d and the second to i_q — the convention used by *ndgrid* — rather than the transposed convention of *meshgrid*.

The original data, which consists of scattered points in the motoring quadrant, was interpolated onto this regular mesh using MATLAB's *scatteredInterpolant* function with natural interpolation and nearest-neighbour extrapolation. Since the original data covers only the motoring quadrant ($i_d \leq 0$, $i_q \geq 0$), the extension to the remaining three quadrants was performed through physical symmetry arguments: the functions $L_d(i_d, i_q)$, $L_q(i_d, i_q)$ and $\psi_{PM}(i_d, i_q)$ are even with respect to the sign of both currents, so each point of the complete mesh is evaluated at its image point $(-|i_d|, +|i_q|)$, which belongs to the domain of the original data:

```
F_Ld = scatteredInterpolant(i_d, i_q, Ld_raw, 'natural', 'nearest');
F_Lq = scatteredInterpolant(i_d, i_q, Lq_raw, 'natural', 'nearest');
F_PM = scatteredInterpolant(i_d, i_q, PM_raw, 'natural', 'nearest');

% Mirroring by physical symmetry (Ld, Lq, PM even in id and iq)
Id_sym = -abs(Id_grid);
Iq_sym =  abs(Iq_grid);
Ld = F_Ld(Id_sym, Iq_sym);
Lq = F_Lq(Id_sym, Iq_sym);
PM = F_PM(Id_sym, Iq_sym);
```

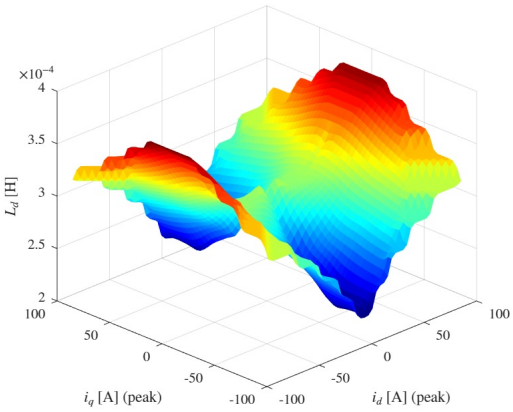
Listing 3.4: Interpolation and mirroring by physical symmetry.

Finally, a light Gaussian smoothing (3-point window) was applied to each of the three surfaces. This step addresses the smoothness requirement identified in Section 3.1.1: it eliminates numerical noise introduced by the interpolation of scattered data, without distorting the physical shape of the saturation characteristics. The resulting matrices were then exported to a *.mat* file for subsequent loading into the Simscape block:

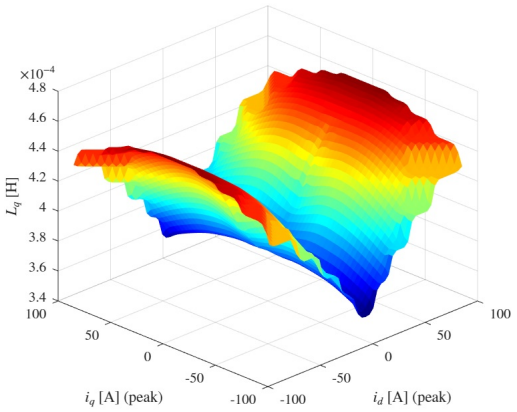
```
save('FEM_Parameters_Corrected.mat', 'id_vec', 'iq_vec', 'Ld', 'Lq', 'PM')
    ;
```

Listing 3.5: Export of the processed parameter maps.

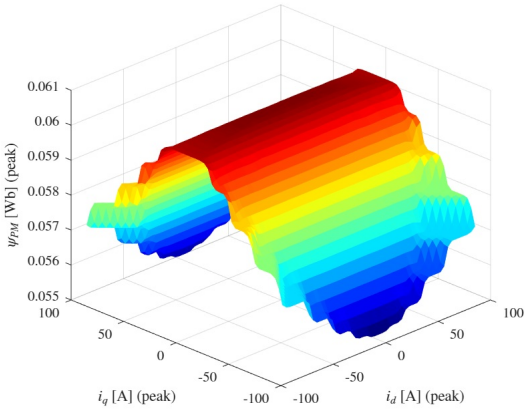
Figure 3.2 shows the graphical representations of the interpolations results.



(a) Direct-axis inductance $L_d(i_d, i_q)$.



(b) Quadrature-axis inductance $L_q(i_d, i_q)$.



(c) Permanent magnet flux linkage $\psi_{PM}(i_d, i_q)$.

Figure 3.2: Parameter maps loaded into the *FEM-Parameterized PMSM* block.

3.1.3 Application of the data to the block

With the processing and validation complete, the *FEM-Parameterized PMSM* block was configured as summarised in Table 3.1. The vectors *id_vec* and *iq_vec* were loaded into the *Direct-axis current vector* and *Quadrature-axis current vector* fields, respectively, and the 51×51 matrices *Ld*, *Lq* and *PM* into the corresponding inductance and flux linkage fields.

The interpolation method was set to *Linear*, which offers a reduction in simulation time relative to the *Smooth* option. The latter uses a modified Akima algorithm to produce a surface with continuous first-order derivatives, whereas the linear method guarantees continuity of the tabulated quantities themselves but not of their derivatives, which the block uses internally in its state equations. Since the parameter maps were smoothed during processing (Section 3.1.2) and the grid resolution is fine relative to the curvature of the surfaces, the discontinuities introduced at the grid nodes are small and did not compromise the numerical stability of the solver. The remaining parameters were entered directly from the manufacturer's datasheet..

Table 3.1: Configuration of the *FEM-Parameterized PMSM* block in Simscape Electrical.

Parameter	**Variable**	**Value / Size**	**Unit**
Modeling option	—	2-D partial derivative data	—
Parameterization	—	Assume sinusoidal back emf — tabulated with d- and q-axis currents	—
Winding type	—	Wye-wound	—
Direct-axis current vector, i_d	*id_vec*	$\langle 1 \times 51 \text{ double} \rangle$	A
Quadrature-axis current vector, i_q	*iq_vec*	$\langle 1 \times 51 \text{ double} \rangle$	A
L_d matrix, $L_d(i_d, i_q)$	*Ld*	$\langle 51 \times 51 \text{ double} \rangle$	H
L_q matrix, $L_q(i_d, i_q)$	*Lq*	$\langle 51 \times 51 \text{ double} \rangle$	H
PM flux linkage, $\psi_{PM}(i_d, i_q)$	*PM*	$\langle 51 \times 51 \text{ double} \rangle$	Wb
Interpolation method	—	Linear	—
Number of pole pairs	p	4	—
Stator resistance per phase, R_s	*R_s*	0.126	Ω
Zero-sequence inductance, L_0	*L_s*	0.000393	H

With this configuration, the motor model is fully parameterized from the manufacturer's finite element data, capturing magnetic saturation and cross-coupling between the d- and q-axes across the entire operating range. As a limitation, the chosen formulation — by assuming sinusoidal permanent magnet flux — neglects the spatial harmonics associated with the variation of the flux linkages with rotor position.

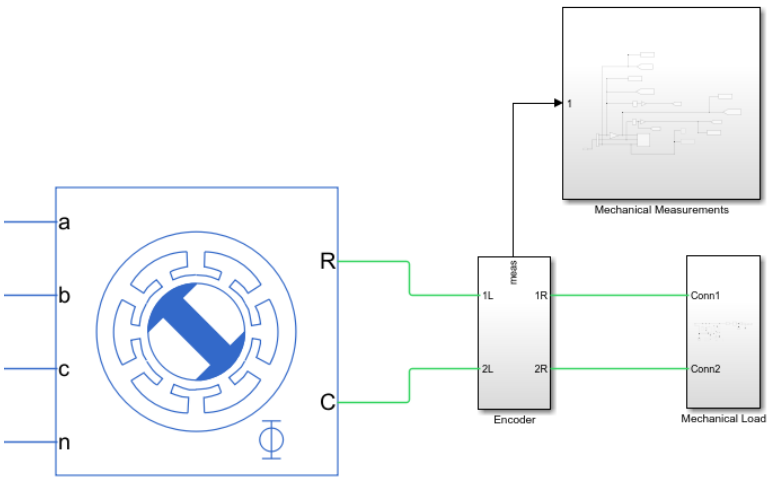


Figure 3.3: FEM-Parameterized PMSM block on *Simulink*

3.1.4 Motor model validation

Under the chosen parameterization, the electromagnetic torque is a function of the current pair (i_d, i_q) alone, which allows the static torque characteristic to be validated independently of any

rotational dynamics. The rotor was therefore held stationary and the currents imposed directly in the synchronous frame, so that the machine is subject to exactly the commanded current pair.

The 117 operating points of the FEA dataset were commanded as a sequence of steps, organised as 13 arcs of constant current magnitude I_s, each sweeping the nine tabulated current angles θ from -80° to 0°. The d- and q-axis currents were reconstructed from each (I_s, θ) pair according to

$$i_d = \sqrt{2} I_s \sin\theta, \qquad i_q = \sqrt{2} I_s \cos\theta \tag{3.3}$$

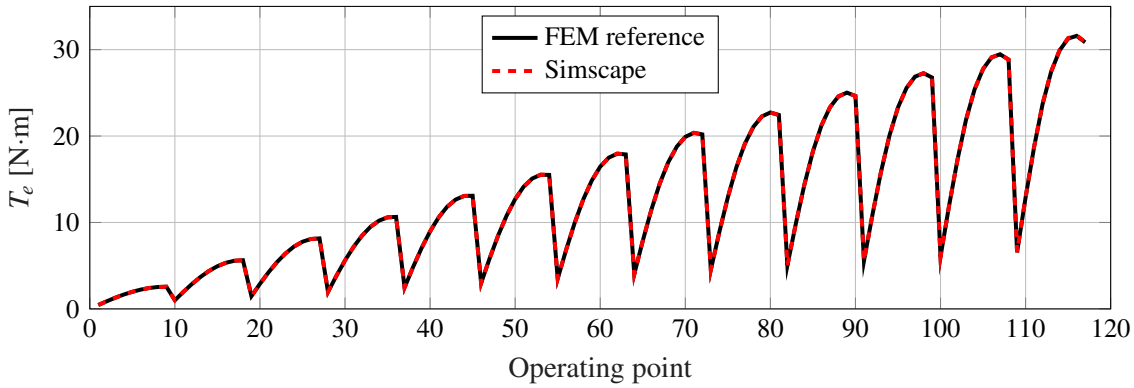
where the factor $\sqrt{2}$ converts the RMS values of the dataset to the peak convention adopted by the block. Each operating point was held for $1/9$ s, well beyond the electrical time constant of the machine ($L_q/R_s \approx 3\,\mathrm{ms}$), and the torque was averaged within each dwell, away from the transitions.

The model reproduces the FEA torque with a mean absolute error of $0.0001\,\mathrm{N\,m}$ and a root-mean-square error of $0.0123\,\mathrm{N\,m}$, against a maximum torque of approximately $31.5\,\mathrm{N\,m}$. The largest deviation at any single operating point is $0.0694\,\mathrm{N\,m}$. Restricting the relative comparison to the 112 points whose torque exceeds 5 % of the maximum, the mean error is $-0.010\,\%$ with a standard deviation of $0.053\,\%$ and a maximum of $0.220\,\%$. The coefficient of determination is $R^2 = 0.999998$. The comparison is shown in Figure 3.4.

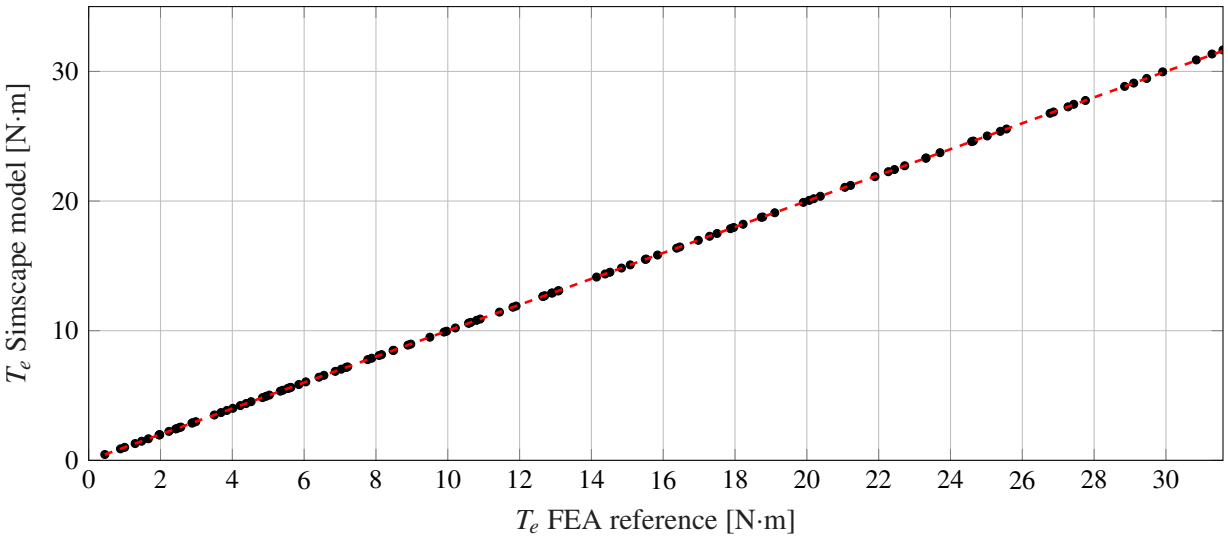
The mean error is two orders of magnitude smaller than its standard deviation, indicating that the residual is dispersion rather than bias: no systematic offset was introduced by the RMS-to-peak conversion, by the number of pole pairs, or by the Park convention, since an error in any of these would appear as a proportional deviation across all points. The magnitude of the dispersion is consistent with interpolation residuals, as the FEA points do not in general coincide with the nodes of the 51×51 grid. Figure 3.5 shows the distribution of the relative error over the current plane.

This confirms that the processing chain described previously reproduces the source data without introducing appreciable error. It should be emphasised, however, that this is a verification of the implementation rather than of the physical fidelity of the model: since the same dataset serves both as parameterization and as reference, close agreement is a necessary but not sufficient condition for the model to represent the machine correctly. The independent verification is provided by the back-EMF test described below.

Similarly to the static torque test, the back-EMF characteristic of the FEM-parametrised model was validated against the manufacturer datasheet A under a no-load condition. The rotor was driven at an imposed mechanical speed through an ideal angular velocity source, while the d- and q-axis stator currents were held at zero by ideal current sources, so that no reaction torque or saturation effect was present and the terminal voltage reflected only the permanent-magnet flux linkage. The floating star point was referenced to the electrical ground through a $1\,\mathrm{M\Omega}$ resistive star to avoid an ill-conditioned open-circuit topology, without drawing any measurable current. The line-to-line RMS voltage was recorded with a Voltage Sensor placed between two phase terminals, across six imposed speed levels from 2000 to 12000 rpm, each held for 1 s to allow the



(a) Torque over the sweep.



(b) Simulated against reference.

Figure 3.4: Torque comparison against the FEA reference.

measurement to settle.

Table 3.2 lists the resulting line-to-line RMS voltage V_{LL} at each imposed mechanical angular speed ω.

The back-EMF constant k_e was extracted from a least-squares linear fit of V_{LL} against ω, forced through the origin in accordance with its definition $V_{LL} = k_e\,\omega$. The resulting value, $k_e = 0.3005$ V_{rms}/(rad/s), deviates by $+1.52\,\%$ from the datasheet value of 0.296 V_{rms}/(rad/s), with a coefficient of determination $R^2 = 0.9917$ over the six points. As the datasheet is marked as preliminary and its parameters are reported as calculated rather than bench-measured values, this deviation is best interpreted as agreement between two FEA-derived characterisations rather than against an independent physical reference.

The point at 2000 rpm departs from the otherwise linear trend: dividing each measurement by its own speed gives a local back-EMF constant of 0.398 V_{rms}/(rad/s) at 2000 rpm, against a near-constant ≈ 0.30 V_{rms}/(rad/s) for the remaining five points. Expressed as a residual voltage

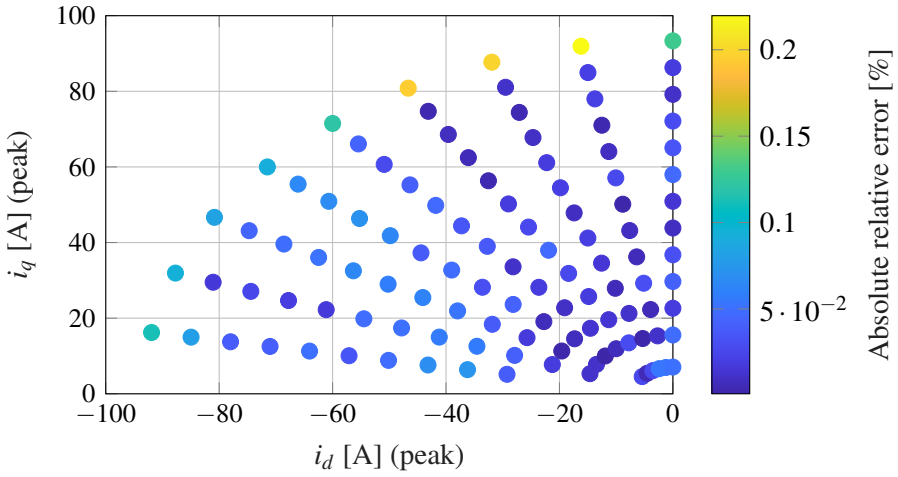


Figure 3.5: Relative torque error over the current plane.

relative to the fitted line, the deviation is of the order of 20 V at 2000 rpm and falls off towards zero at higher speeds (7.2 V at 4000 rpm, 2.0 V at 6000 rpm, and below 5 V in magnitude beyond that). A residual of approximately fixed magnitude that does not scale with ω is inconsistent with a physical effect proportional to the fundamental voltage, such as a spatial flux harmonic or a saturation-related variation of ψ_{PM}, both of which would scale with the fundamental component itself and would therefore be most pronounced, not least pronounced, at high speed. The pattern is instead consistent with a fixed-magnitude transient still present in the RMS voltage computed over the 1 s dwell window, which represents a much larger fraction of the true back-EMF amplitude at 2000 rpm than at higher speeds. Excluding the 2000 rpm point from the fit changes the extracted constant by only 0.37 % (from 0.3005 to 0.2994 V_{rms}/(rad/s), lowering the datasheet deviation to $+1.15\,\%$) while raising R^2 to 0.9979, confirming that the discrepancy is concentrated in this single low-speed measurement rather than reflecting a systematic bias of the model across the operating range. This explanation has not been confirmed by inspection of the time-domain voltage waveform at 2000 rpm and is reported here as the most consistent account given the available data, rather than as a verified cause.

Figure 3.6 shows the simulated points together with the fitted line and the datasheet reference.

Table 3.2: Simulated no-load line-to-line RMS voltage as a function of imposed mechanical speed.

n [rpm]	ω [rad/s]	V_{LL} [V_{rms}]
2000	209.5	83.35
4000	418.9	133.06
6000	628.4	190.82
8000	837.8	250.97
10000	1047.2	312.02
12000	1256.4	373.55

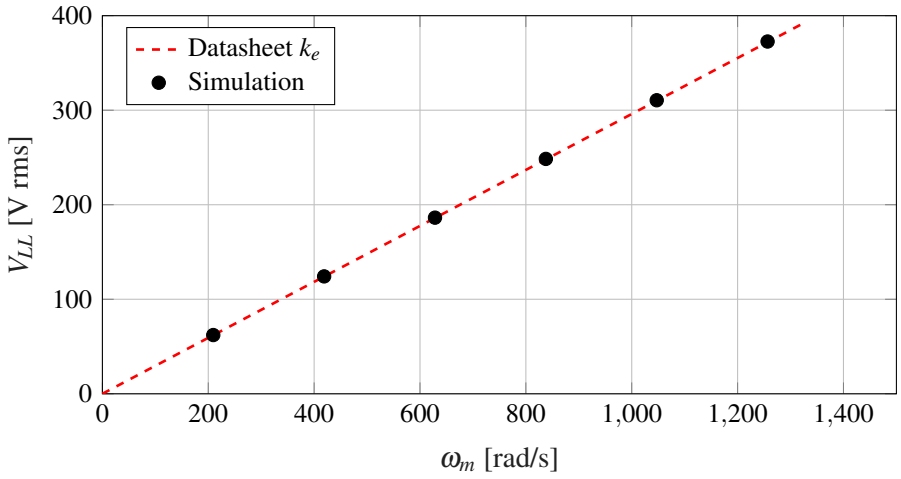


Figure 3.6: Simulated back-EMF versus mechanical speed.

3.1.5 FEM-Parameterized PMSM vs Linear Model

To assess the contribution of magnetic saturation to the machine's torque characteristic, the FEM-parametrised Simscape model was compared against a simple linear model evaluated at the same 112 operating points used in the static torque validation (excluding the five points with $|T_e|$ below 5 % of the maximum torque). The linear model uses fixed parameters extracted at the lowest-current point of the FEA dataset at zero d-axis current ($I_s = 5$ A$_{rms}$, $\theta = 0°$): $L_d = 0.290$ mH, $L_q = 0.372$ mH and $\psi_{PM} = 0.0605$ Wb (peak), calculating the torque from equation 2.5, evaluated directly from the FEA current pairs (i_d, i_q) without requiring simulation. These pairs are reconstructed from the FEA current magnitude I_s and angle θ through the same rotation and RMS-to-peak conversion used in the Simulink harness, and were checked to reproduce the $I_{d,eff}/I_{q,eff}$ columns of the FEA spreadsheet. Table 3.3 compares the relative error of the linear model and of the FEM-parametrised Simscape model against the FEA reference torque, over the 112 significant operating points. The FEM-parametrised model reproduces the reference torque essentially

Table 3.3: Relative torque error against the FEA reference, linear model versus FEM-parametrised Simscape model.

Model	Mean error [%]	Std. dev. [%]	Max. |error| [%]
Linear	0.93	2.44	9.84
FEM (Simscape)	−0.01	0.05	0.22

exactly, consistent with the static torque validation reported above. The linear model shows an order-of-magnitude larger spread, with a maximum deviation of 9.84 %. This is expected: its parameters are fixed at a single operating point and cannot represent the variation captured by the FEA data, over which L_d ranges from 0.217 to 0.384 mH, L_q from 0.355 to 0.461 mH, and ψ_{PM} from 0.0551 to 0.0606 Wb (peak).

It is worth noting that this variation is not governed by the magnitude of the current alone. Along the $i_d = 0$ arc, where the magnet flux is unopposed, saturation does increase monotonically with current: ψ_{PM} falls by approximately 9 % (from 0.0428 to 0.0390 Wb_{rms}) as I_s rises from 5 to 66 A_{rms}. Away from that arc, however, a negative i_d opposes the magnet flux and reduces the total flux in the machine, so that the least saturated condition of the entire dataset occurs at the highest current magnitude with the most strongly demagnetising current angle ($I_s = 66\ \mathrm{A}_{rms}$, $\theta = -80°$), where the total flux linkage is less than half its value at the reference point. The chosen reference point is therefore one of low current, not one of low saturation.

Figure 3.7 isolates the $i_d = 0$ arc, comparing the FEA reference, the FEM-parametrised Simscape model and the linear model as a function of i_q. Figure 3.8 shows the relative error of both models across the full set of significant operating points as a function of the current magnitude $|I_s|$, making it possible to identify the loading region in which the linear approximation degrades.

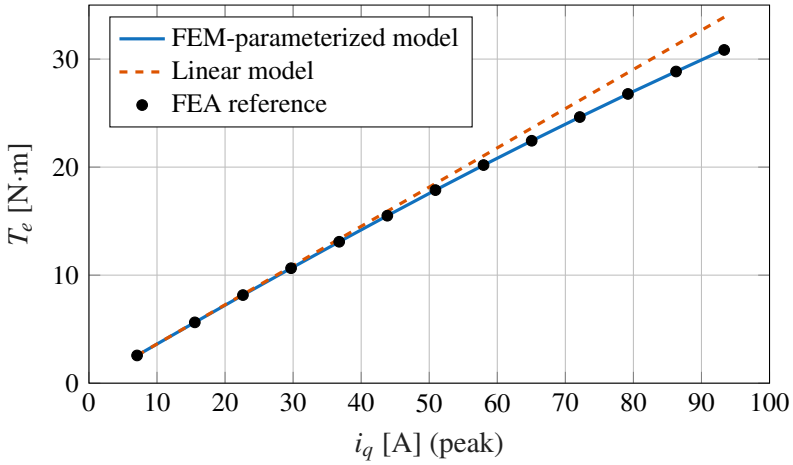


Figure 3.7: Torque along the $i_d = 0$ arc.

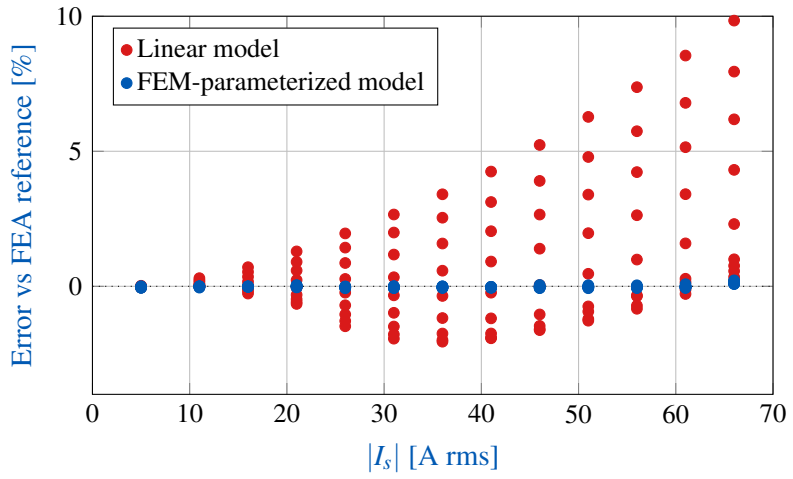


Figure 3.8: Relative error versus current magnitude.

3.1.6 Constant Parameters Used Outside the Machine Model

The machine model itself uses the full FEA maps, but several parts of the work require single fixed values of L_d, L_q and ψ_{PM}: the linear torque model of Section 3.1.5, the current regulator gains, the scalar V/f voltage profile, and the analytical MTPA trajectory. A single set is used throughout, taken from the lowest-current point of the FEA dataset at zero d-axis current ($I_s = 5$ A$_{rms}$, $\theta = 0°$):

$$L_d = 0.290 \text{ mH}, \qquad L_q = 0.372 \text{ mH}, \qquad \psi_{PM} = 0.0605 \text{ Wb (peak)}. \tag{3.4}$$

This point is not the least saturated of the dataset. Ordered by total flux linkage it lies among the more saturated points, because with $i_d = 0$ nothing opposes the magnet flux; the least saturated condition occurs at the highest current magnitude with the most strongly demagnetising angle. It is, however, the point of lowest current at which the machine still produces torque, and therefore the natural origin of the operating envelope: along the MTPA trajectory the current angle stays between 0 and $-6.6°$, so the region the drive actually visits lies in the neighbourhood of $i_d = 0$, and current only increases from this point onwards.

Anchoring at this point gives each use a deviation in a known and favourable direction. For the regulator gains, the inductances fall as the machine loads, which raises both the natural frequency and the damping ratio, so the design overshoot is an upper bound rather than a nominal figure. For the linear torque model, the error grows monotonically with load, which is the most direct possible reading of the effect being measured. For the V/f voltage profile, the magnet flux is the highest of the range, so the machine is slightly over-excited, and under-excitation is what causes an open-loop scalar scheme to lose synchronism. In the feedforward terms the choice is one of consistency rather than optimality: a value nearer the mean of the operating region would leave a smaller residual, but the residual is a disturbance at the regulator input which integral action removes in steady state, and using one parameter set throughout avoids maintaining two.

3.2 Inverter

The converter is modelled following the general block-based architecture described in [2] for a two-level three-phase SiC-based inverter, adapted from its original PLECS implementation into the Simscape environment used throughout this dissertation. It consists of six switching devices arranged as three half-bridge legs, consistent with the two-level three-phase topology already described in Section 2.2. Each switching device is a SiC MOSFET (Infineon IMZA120R014M1H [29]) parametrised directly from the manufacturer datasheet through its on-state resistance, $R_{DS(on)} = 14\,\text{m}\Omega$, and gate-source threshold voltage, $V_{th} = 4.2$ V, both taken at $T_{vj} = 25\,°\text{C}$.

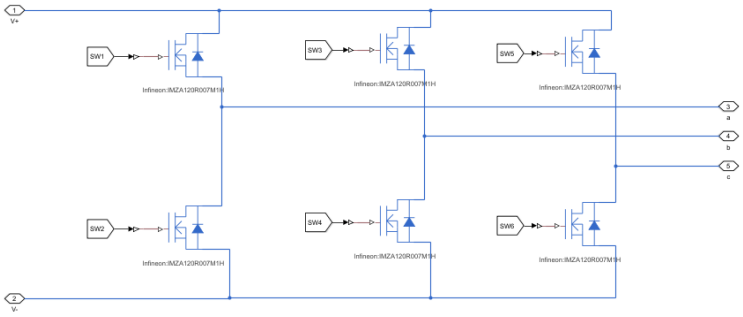


Figure 3.9: Inverter model on *Simulink*

The integral body diode, providing the freewheeling path for the reverse current, is modelled as a linear $V = V_f + R_{on} I$ characteristic fitted to two points of the diode's reverse conduction curve ($I_{SD} = f(V_{SD})$, $V_{GS} = 0$ V) from the same datasheet: the tabulated value $V_{SD} = 3.8$ V at $I_{SD} = 108$ A, and an approximate reading of $V_{SD} \approx 3.2$ V at $I_{SD} \approx 5$ A taken from the conduction knee of the curve. This gives $R_{on} \approx 6\,\text{m}\Omega$ and $V_f \approx 3.2$ V. Since the second point was read graphically rather than tabulated, these diode parameters are an approximation rather than an exact datasheet extraction.

Table 3.4 summarises the parameters used to configure the switching device block.

Table 3.4: Switching device model parameters (Infineon IMZA120R014M1H).

MOSFET		
Parameter	Value	Unit
Drain-source on-state resistance, $R_{DS(on)}$	0.014	Ω
Threshold voltage, V_{th}	4.2	V
Integral diode		
Parameter	Value	Unit
Forward voltage, V_f	3.2	V
On resistance, R_{on}	0.006	Ω

3.3 Space-Vector Modulation PWM

Space-Vector Modulation is implemented in MATLAB/Simulink using the native *PWM Generator (Three-phase, Two-level)* block from the Simscape Electrical library, configured in *Continuous* space-vector mode. Sector determination, dwell-time calculation and the assignment of switching instants to the three inverter phases are therefore performed internally by this library block, rather than by a custom-built subsystem chain.

Figure 3.10 presents the complete Simulink implementation of the modulator, from the voltage references to the six gate-drive signals. The three-phase voltage references v_a, v_b, v_c and the DC-link voltage ($V_{dc} = 600$ V) are supplied to the *PWM Generator* block. The three modulating waveforms available at the *ModWave* port, split into the individual signals *sine1*, *sine2* and *sine3*, are the ones actually used to build the gate-drive signals.

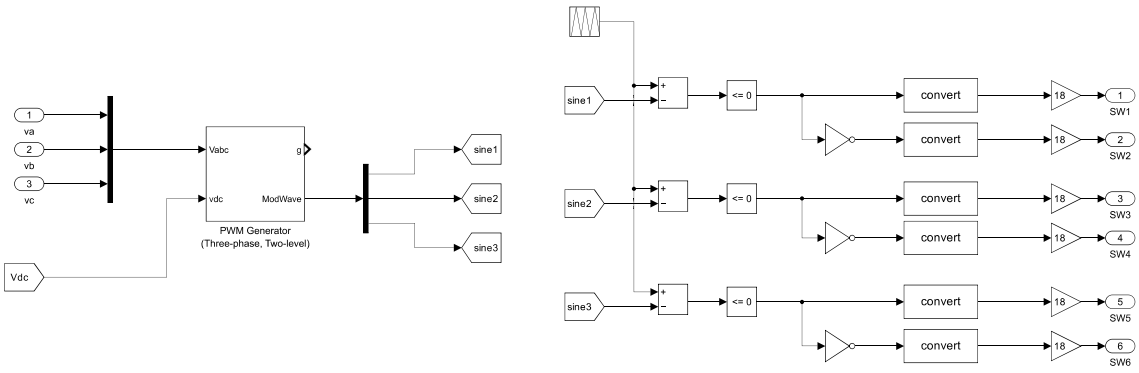


Figure 3.10: Simulink implementation of the SVM modulator and of the PWM gate signal generation.

In the second half of the diagram, each modulating waveform is compared with a common triangular carrier operating at the switching frequency: the carrier is subtracted from the modulating waveform, and the result is passed through a relational (*<=0*) block, producing the logical gating signal for the upper switch of the corresponding inverter leg. Since each leg is composed of two complementary switches, the gating signal of the lower switch is obtained through a logical NOT of the upper one. Both logical signals are then converted from boolean to double (*convert* blocks) and scaled by a gain of 18 before being routed to the six switch references, *SW1*–*SW6*.

3.3.0.1 Operating Principle of the Modulator

According to MathWorks' own description of the technique [20], and consistently with the original formulation of the algorithm [19], the block internally treats the three-phase reference as a single rotating vector in the stationary $\alpha\beta$ plane. Its angular position is obtained from

$$\theta = \operatorname{atan2}(V_\beta, V_\alpha), \tag{3.5}$$

which determines the active sector, i.e. the 60° interval of the space-vector plane in which the reference vector currently lies, and consequently the pair of adjacent active voltage vectors used to synthesise it. The application times of these two active vectors within the sector, together with that of the zero vectors, are then computed as

$$T_1 = k \sin\left(\theta - (S-1)\tfrac{\pi}{3}\right), \qquad T_2 = k \sin\left(S\tfrac{\pi}{3} - \theta\right), \qquad T_0 = T_s - T_1 - T_2, \tag{3.6}$$

with S the identified sector, T_s the switching period and k a constant proportional to the reference voltage magnitude and T_s, and inversely proportional to V_{dc}.

In the continuous mode selected for this block, both zero vectors (000 and 111) are applied within every switching period, split symmetrically as $T_0/2$ at the start and end of the interval, giving the conventional seven-segment switching sequence. This differs from discontinuous (DPWM) variants, where one inverter leg is clamped to a DC rail for part of the fundamental period to reduce switching losses at the cost of increased current ripple in that leg.

Rather than exposing these internal dwell-time calculations, the block outputs the modulating waveforms directly, and the gating signals are obtained, as described by MathWorks [20], by the carrier comparison described above. Figure 3.11 shows the three sinusoidal voltage references applied to the modulator over one fundamental period, and Figure 3.12 shows the corresponding modulating waveforms together with the triangular carrier against which they are compared. Comparing the two figures makes the effect of the algorithm visible: the modulating waveforms are no longer purely sinusoidal, but exhibit the characteristic "double-hump" shape reported by MathWorks for this technique [20], and it is this additional component that accounts for the improved DC-bus utilisation over sinusoidal PWM.

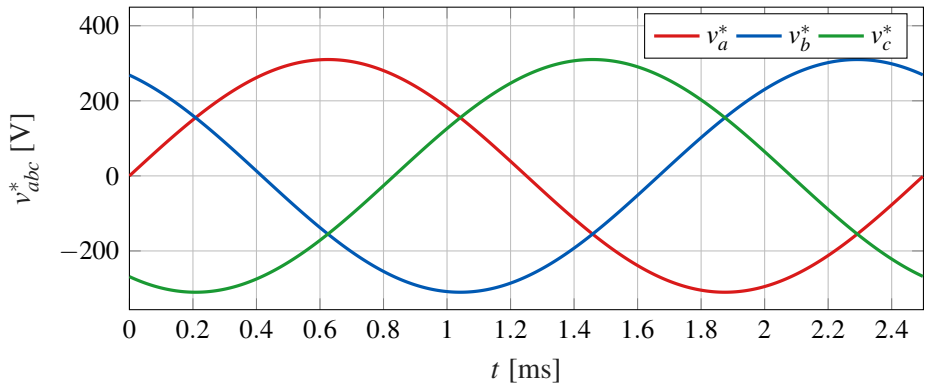


Figure 3.11: Three-phase sinusoidal voltage references applied to the modulator.

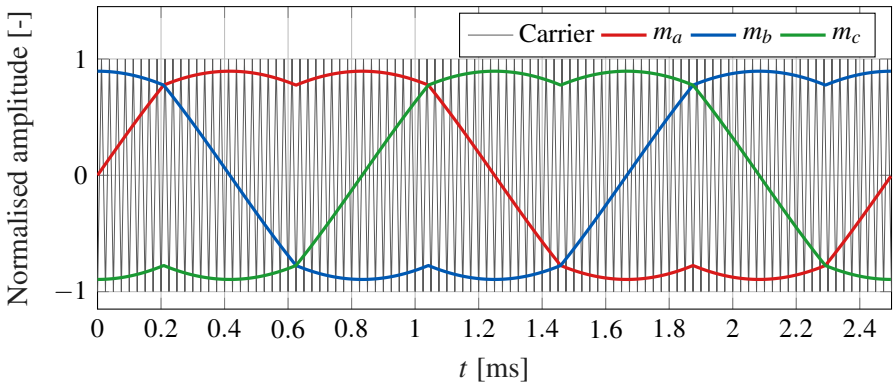


Figure 3.12: SVM modulating waveforms and triangular carrier.

Figure 3.13 shows the resulting gate signals *SW1* and *SW2* for a single inverter leg over a window spanning a few switching periods, plotted as stacked subplots, confirming their complementary switching behaviour.

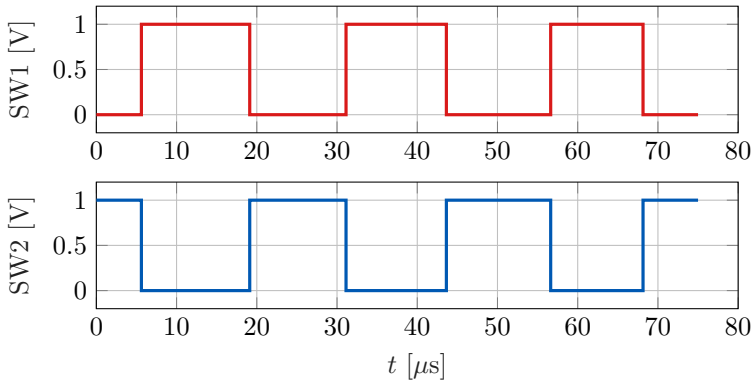


Figure 3.13: Complementary gate signals SW1 (upper switch) and SW2 (lower switch).

It is worth noting that this implementation assumes ideal switching devices: dead-time insertion, finite switching delays, and semiconductor voltage drops are not modelled here. The objective is to evaluate the modulation strategy itself, independently of these hardware-related effects.

Figure 3.14 shows the resulting phase-to-neutral output voltages of the inverter over two fundamental periods, measured with respect to the floating star point of the load. Because the star point is not tied to the DC-link midpoint, its potential varies with the common-mode component of the modulation, and the phase-to-neutral voltage therefore takes five discrete levels (0, $\pm V_{dc}/3$ and $\pm 2V_{dc}/3$) rather than the two levels of the pole voltages. The resulting stepped waveform confirms that the generated switching sequence synthesises the intended sinusoidal reference.

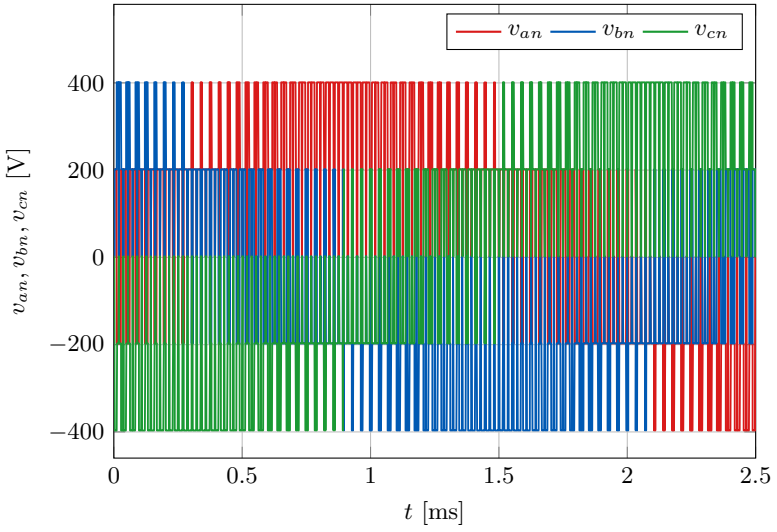


Figure 3.14: Phase-to-neutral inverter output voltages generated by the SVM gate signals.

3.4 Current Controller

3.4.1 Control Design

The gains of the two current regulators are obtained by writing the closed-loop transfer function of each axis, recognising that it has the form of a second-order system, and identifying its coefficients with those of the canonical second-order form [30]. The damping ratio and natural frequency of that canonical form are then fixed from the desired step response, and the resulting expressions are solved for the gains in terms of quantities known from the machine.

3.4.1.1 Plant and Closed-Loop Transfer Function

The starting point for the current controller design is the dynamic voltage equations of the PMSM expressed in the synchronous rotating reference frame [31]:

$$v_d = R_s i_d + \frac{\mathrm{d}\psi_d}{\mathrm{d}t} - \omega_e \psi_q, \tag{3.7}$$

$$v_q = R_s i_q + \frac{\mathrm{d}\psi_q}{\mathrm{d}t} + \omega_e \psi_d, \tag{3.8}$$

where the d- and q-axis stator flux linkages are defined as $\psi_d = L_d i_d + \psi_{PM}$ and $\psi_q = L_q i_q$.

For the purpose of linear current regulator design and tuning via canonical second-order identification, a linear magnetic circuit is assumed. Consequently, the stator inductances (L_d, L_q) and the permanent magnet flux linkage (ψ_{PM}) are considered constant over time, meaning their derivatives with respect to time are zero. Under these linear assumptions, the time derivatives of the flux linkages in Equation (3.7) and Equation (3.8) reduce directly to the derivatives of the currents. This mathematical simplification yields the standard plant equations:

$$v_d = R_s i_d + L_d \frac{\mathrm{d}i_d}{\mathrm{d}t} - \omega_e L_q i_q, \tag{3.9}$$

$$v_q = R_s i_q + L_q \frac{\mathrm{d}i_q}{\mathrm{d}t} + \omega_e \left(\psi_{PM} + L_d i_d\right). \tag{3.10}$$

Once the speed-dependent terms are cancelled by the feedforward of 4.2.1.2, both equations reduce to the same first-order form, which may be written for a generic axis x as

$$v_x = R_s i_x + L_x \frac{\mathrm{d}i_x}{\mathrm{d}t}, \tag{3.11}$$

so that a single design carried out on (3.11) applies to both axes, differing only in the inductance. In the Laplace domain the plant relating the applied voltage to the current follows a decoupled first-order structure [32] and is therefore

$$P(s) = \frac{I(s)}{V(s)} = \frac{1}{sL + R_s}. \tag{3.12}$$

Each axis is regulated by a proportional-integral controller placed in the forward path of a unity-feedback loop,

$$C(s) = K_p + \frac{K_i}{s} = \frac{sK_p + K_i}{s}, \tag{3.13}$$

so that the transfer function from current reference to current is

$$H(s) = \frac{I(s)}{I_{ref}(s)} = \frac{C(s)P(s)}{1 + C(s)P(s)}. \tag{3.14}$$

Substituting (3.12) and (3.13) and simplifying gives

$$H(s) = \frac{K_p s + K_i}{Ls^2 + (R_s + K_p)\, s + K_i}. \tag{3.15}$$

The denominator is a second-order polynomial, so (3.15) can be compared term by term with the canonical second-order transfer function. Dividing numerator and denominator by L to make the coefficient of s^2 unity,

$$H(s) = \frac{\dfrac{K_p}{L} s + \dfrac{K_i}{L}}{s^2 + \dfrac{R_s + K_p}{L} s + \dfrac{K_i}{L}}, \tag{3.16}$$

which is to be identified with

$$H_c(s) = \frac{\omega_0^2}{s^2 + 2\xi\omega_0 s + \omega_0^2}. \tag{3.17}$$

3.4.1.2 Identification and Design Equations

Comparing the denominators of (3.16) and (3.17) term by term gives two conditions,

$$\omega_0^2 = \frac{K_i}{L}, \qquad 2\xi\omega_0 = \frac{R_s + K_p}{L}, \tag{3.18}$$

which solved for the gains yield the design equations

$$K_i = L\omega_0^2, \qquad K_p = 2\xi\omega_0 L - R_s. \tag{3.19}$$

Both gains are thus expressed in terms of two quantities known from the machine, the winding resistance and the axis inductance, and two design parameters, the damping ratio and the natural frequency, whose selection is discussed next.

Comparing the numerators, however, shows that the identification is not complete. The canonical form (3.17) has a constant numerator ω_0^2, which corresponds to the term K_i/L in (3.16), but the actual closed loop carries an additional term $(K_p/L)\, s$ that has no counterpart in (3.17). The consequence of this extra term, and the means of removing it, are treated in Section 3.4.1.6.

3.4.1.3 Choice of the Design Parameters

The damping ratio adopted is

$$\xi = \frac{1}{\sqrt{2}} \approx 0.707, \tag{3.20}$$

which is the conventional choice for a second-order closed loop. Its significance is a frequency-domain one: the closed-loop magnitude response of (3.17) exhibits a resonant peak only for $\xi < 1/\sqrt{2}$, so this value is precisely the boundary at which the peak disappears and the response becomes maximally flat [33]. It corresponds to the second-order Butterworth response, and at this damping the -3 dB bandwidth of the closed loop coincides exactly with ω_0, which makes the natural frequency of the design directly interpretable as the bandwidth of the current loop.

The resulting overshoot follows from the standard expression for the peak of the step response of a second-order system, which depends on the damping ratio alone [33],

$$M_p = \exp\left(\frac{-\pi\xi}{\sqrt{1-\xi^2}}\right), \tag{3.21}$$

and amounts to 4.3 % for the value adopted. A lower damping ratio would introduce peaking and a larger overshoot, while a higher one would slow the response without removing any peaking, since none remains to be removed.

The natural frequency is bounded from above by the switching frequency of the converter. The regulator acts through a modulator that updates once per switching period, so the closed-loop dynamics must be slow enough for the converter to be treated as an ideal voltage source over the timescale of the response; the requirement is $\omega_0 \ll \omega_{sw}$. The value adopted here is

$$\omega_0 = \frac{\omega_{sw}}{20} = \frac{2\pi \times 40000}{20} = 12566 \text{ rad/s}, \tag{3.22}$$

that is a natural frequency of 2 kHz against a switching frequency of 40 kHz.

With ξ and ω_0 fixed, the remaining response characteristics follow from the standard approximations for a second-order system [33],

$$t_r \approx \frac{2.16\,\xi + 0.60}{\omega_0}, \qquad t_s \approx \frac{4}{\xi\,\omega_0}, \tag{3.23}$$

the first valid for $0.3 \leq \xi \leq 0.8$ and referred to the 10–90 % interval, the second to the 2 % band. They give a rise time of 169 μs and a settling time of 450 μs. Both are estimates rather than exact values: evaluating the canonical step response directly for these parameters gives 171 μs and 475 μs, the settling time approximation being the coarser of the two because it is derived from the exponential envelope rather than from the oscillation itself.

3.4.1.4 Verification of the Bandwidth

The condition $\omega_0 \ll \omega_{sw}$ is a qualitative one, and the factor of twenty adopted in (3.22) deserves to be checked against the phase available at the crossover frequency, since the continuous-time identification of Section 3.4.1.2 takes no account of the delays of the digital implementation. The modulator holds its output for one period, contributing an average delay of $T_s/2$, and the sampled current is only acted upon in the following period, adding a further T_s. Such delays leave the magnitude of the open-loop response unchanged but subtract phase in proportion to frequency [34].

Table 3.5 collects the result. At the chosen bandwidth the crossover lies at 3053 Hz and the phase margin falls from 65.9° to 24.7° once both delays are included, which is low but workable. Doubling the natural frequency to 4 kHz moves the crossover to 6161 Hz, where the same delay consumes more phase than is available and the margin becomes negative. The factor of twenty is therefore not merely conventional: it is close to the limit of what this sampling rate permits.

Table 3.5: Phase margin of the d-axis loop under the modulator and computation delays.

Condition	Crossover [Hz]	Phase margin [°]
No delay	3053	65.9
Modulator delay, $0.5\,T_s$	3053	52.2
Modulator and computation, $1.5\,T_s$	3053	24.7
As above at 4 kHz natural frequency	6161	−17.5

For comparison, the open-loop electrical poles of the machine lie at $R_s/L_d = 434$ rad/s and $R_s/L_q = 339$ rad/s, that is at 69 and 54 Hz, so the regulators accelerate the natural current response by roughly a factor of thirty.

3.4.1.5 Resulting Gains

Evaluating (3.19) with the constant parameters of Section 3.1.6 gives the gains of Table 3.6. The two axes differ only through their inductance, the quadrature axis being the larger and therefore requiring proportionally larger gains to obtain the same closed-loop dynamics.

Table 3.6: Current regulator gains.

Axis	K_p [V/A]	K_i [V/(A·s)]	$\tau = K_p/K_i$ [μs]
d	5.03	4.58×10^4	110
q	6.48	5.87×10^4	110

3.4.1.6 Reference Filter

The identification carried out in Section 3.4.1.2 matched the two denominators, but the numerators do not match: the canonical form (3.17) has the constant numerator ω_0^2, whereas (3.16) carries an

additional term $(K_p/L)\,s$. Because of it the numerator is a first-degree polynomial, and the closed loop therefore has a zero, located where that polynomial vanishes:

$$\frac{K_p}{L}s+\frac{K_i}{L}=0 \iff K_p s+K_i=0 \iff s=-\frac{K_i}{K_p}. \tag{3.24}$$

The origin of this term is the proportional action. The controller acts on the error $e=i_{ref}-i$, so its proportional part contributes $K_p\,i_{ref}$ directly to the commanded voltage, and a step in the reference produces an immediate step in the applied voltage before the current has responded.

Writing $\tau=K_p/K_i$ and using $K_i=L\omega_0^2$, the numerator factors as $\omega_0^2\,(1+\tau s)$, so that the closed loop is exactly the canonical response multiplied by the zero,

$$H(s)=(1+\tau s)\,H_c(s). \tag{3.25}$$

For the gains designed here the zero lies at -9110 rad/s, close to the closed-loop poles at $\omega_0=12566$ rad/s, and its effect on the step response is large: the overshoot becomes $19.8\,\%$ instead of the intended $4.3\,\%$.

The unwanted factor in (3.25) is removed by applying its inverse to the reference, which is a first-order low-pass filter whose pole is placed on the zero,

$$F(s)=\frac{1}{\tau s+1}, \qquad \tau=\frac{K_p}{K_i}, \tag{3.26}$$

so that

$$F(s)H(s)=\frac{1}{1+\tau s}\,(1+\tau s)\,H_c(s)=H_c(s), \tag{3.27}$$

which is the canonical form the design equations were derived from. The overshoot, rise time and settling time predicted in Section 3.4.1.3 are thereby recovered.

Two remarks complete the description. The filter acts on the reference only and lies outside the feedback loop, so the closed-loop poles, the stability of the loop and the rejection of disturbances are unaffected; it changes only the trajectory the loop is asked to follow. And since the filter is part of the controller rather than of the reference generation, the response figures quoted in this section are measured from the commanded reference, upstream of the filter.

For the gains of Table 3.6 the time constant is $110\ \mu$s on both axes, the two values coinciding because

$$\tau=\frac{2\xi\,\omega_0 L-R_s}{L\omega_0^2}=\frac{2\xi}{\omega_0}-\frac{R_s}{L\omega_0^2}, \tag{3.28}$$

whose first term does not depend on the inductance and dominates the second by two orders of magnitude for this machine. The final *Simulink* implementation of the FOC is shown in Figure 3.15

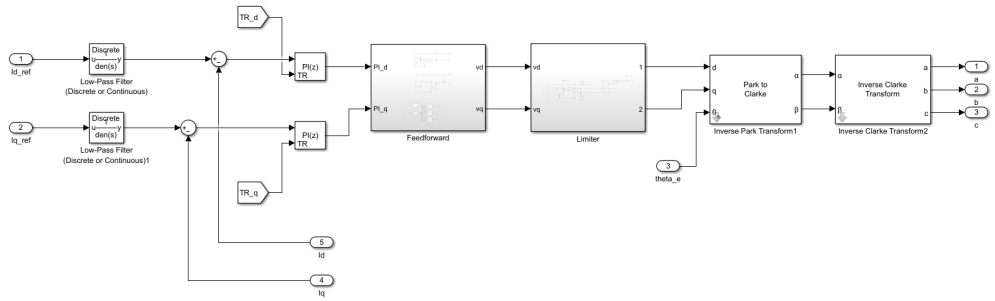


Figure 3.15: FOC block diagram on *Simulink*

The feedback signals required by the control system — namely the three-phase stator currents (i_{abc}), the electrical rotor position (θ_e), and the mechanical shaft speed (ω_m) — are obtained directly from ideal measurement blocks within the *Simulink* and *Simscape Electrical* libraries. The real phase currents are extracted via *Ideal Current Sensor* blocks at the machine terminals and are subsequently converted into the synchronous dq components through the Clarke and Park reference frame transformations. The high-resolution position and velocity measurements are provided by the *Ideal Rotational Motion Sensor* block.

3.5 Maximum Torque per Ampere

The rationale for the MTPA strategy and its contrast with the $i_d = 0$ alternative were established in Section 2.1. What remains is to derive the trajectory for this particular machine, decide which parameter values to evaluate it with, implement it in a form the drive can execute, and verify that the implementation does what it is supposed to do.

3.5.1 Trajectory

Controlling the stator current vector in magnitude and angle so as to obtain a given torque with the minimum current was introduced with the interior magnet machine itself [35], and the resulting current angle condition was formalised in [36]. The trajectory follows from minimising the current magnitude $|i_s|^2 = i_d^2 + i_q^2$ subject to the torque of (2.5) being held at a prescribed value. With a Lagrange multiplier λ,

$$\mathscr{L} = i_d^2 + i_q^2 + \lambda \left\{ T_e - \tfrac{3}{2} p \left[\psi_{PM} i_q + (L_d - L_q) i_d i_q \right] \right\}. \tag{3.29}$$

Setting $\partial \mathscr{L} / \partial i_d$ and $\partial \mathscr{L} / \partial i_q$ to zero and eliminating λ removes the torque level from the problem, leaving a condition that the optimal pair satisfies at every operating point,

$$\Delta L \, i_d^2 - \psi_{PM} \, i_d - \Delta L \, i_q^2 = 0, \qquad \Delta L = L_q - L_d, \tag{3.30}$$

whose negative root, the physically meaningful one since reluctance torque requires a demagnetising current, gives the trajectory in the closed form used throughout the literature [36]:

$$i_d = k - \sqrt{k^2 + i_q^2}, \qquad k = \frac{\psi_{PM}}{2\Delta L}. \tag{3.31}$$

A single parameter therefore defines the entire trajectory. Since k is the ratio of magnet flux to saliency, a small k corresponds to a trajectory that departs markedly from the i_q axis and a large k to one that stays close to it.

3.5.2 Parameter Values

Equation (3.31) is evaluated with the constant set declared in Section 3.1.6, giving $\Delta L = 0.082$ mH and $k = 368.9$ A. The magnet flux linkage in (3.31) is the same one that appears in the back-EMF, in the decoupling feedforward and in the torque constant, and it is the value validated against the datasheet in Section 3.1.4. Adopting a different value here, extracted from some saturated operating point, would require the inductances of that same point to be adopted with it, which would replace a model with a clear physical definition by one linearised about an arbitrarily chosen condition, and would leave the drive using two different values of ψ_{PM} in different places.

This choice must be qualified. That magnetic saturation alters the machine parameters and therefore displaces the optimal current angle is a long-established limitation of parameter-based MTPA [37], and in this machine it is k that carries almost all of that sensitivity. It is formed from the difference of two inductances that are close to one another, so moderate saturation of either produces a large relative change in ΔL. Across the finite element dataset ψ_{PM} varies by 9.0 %, but ΔL varies by 59.7 % and k ranges from 194 A to 507 A. Figure 3.16 shows this: the upper panel gives k evaluated at every point of the dataset against the constant adopted, and the lower panel the trajectories resulting from the extremes of that range. The scatter at any given current magnitude is itself substantial, because k depends on the current angle as well as on its magnitude, so there is not even a single well-defined k per current level. At the highest current in the dataset with zero direct-axis current, k falls to 206 A, which would call for -17.3 A of direct-axis current against the -10.0 A given by the constant.

The trajectory therefore remains valid throughout the operating range, in the sense that every point on it produces the commanded torque, but its optimality degrades as the machine saturates. The consequence is quantified in Section 4.3.1.3.

3.5.3 Implementation

The controller receives a torque reference and must produce two current references. Equation (3.31) yields i_d once i_q is known, but obtaining i_q from T_e^* requires substituting the trajectory into the torque expression, which produces a quartic. Rather than solve it online, the inverse relation is approximated by a least-squares polynomial of third degree,

$$i_q^* = a_3 T_e^3 + a_2 T_e^2 + a_1 T_e + a_0, \tag{3.32}$$

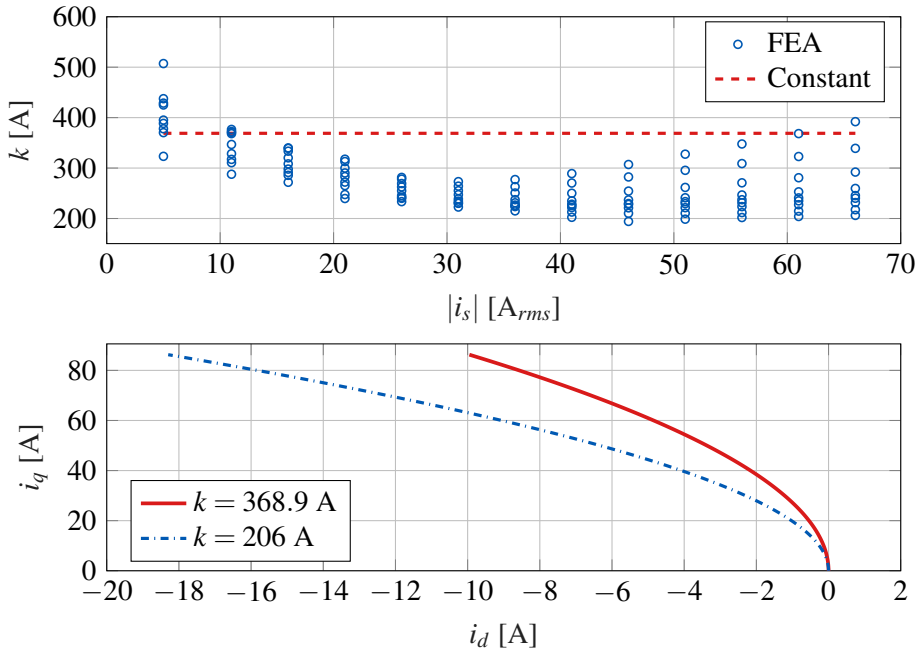


Figure 3.16: Variation of k across the finite element dataset and its effect on the trajectory.

with the coefficients of Table 3.7. The maximum deviation from the exact trajectory is 2.24 mA over the full torque range, three orders of magnitude below the resolution of any realistic current measurement, so a higher degree would bring no observable benefit. Once i_q^* is known, i_d^* follows from (3.31) using only a square and a square root.

Table 3.7: Coefficients of the polynomial inversion $i_q^*(T_e^*)$, with T_e in N·m and i_q in A peak.

Coefficient	Value
a_3	-3.2986×10^{-5}
a_2	-1.4419×10^{-4}
a_1	2.7559
a_0	-1.8645×10^{-3}

The reference generator is built from native Simulink blocks, without embedded code, so that it maps directly onto the arithmetic of an embedded target. In operation it executes the following sequence at every control period. The torque reference is first clipped to $\pm T_{e,\max}$, the upper end of the range over which the polynomial was fitted and beyond which the current limit of the machine would be exceeded. Its absolute value is passed through (3.32), and the sign of the original reference is then restored by multiplication. This construction extends a trajectory fitted only for positive torque to negative torque without refitting, and it is legitimate because the MTPA condition is symmetric in i_q: (3.31) depends on i_q^2, so reversing the torque reverses i_q and leaves i_d unchanged. It is this symmetry that allows traction and regenerative braking to be handled by the same trajectory, as observed in the reversal test of Section 4.2.1.4. Finally i_q^* is squared, added to

k^2, square-rooted and subtracted from k, which is (3.31) implemented literally. Figure 3.17 shows the resulting chain.

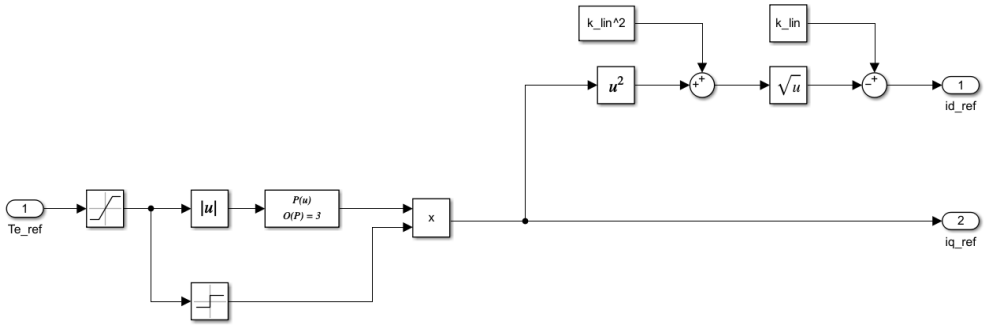


Figure 3.17: Reference generation chain implementing the MTPA trajectory.

The operating limits follow from the peak current rating of the machine, 61 A_{rms}, giving $i_{q,\max} = 86.3$ A peak and $T_{e,\max} = 31.7$ N·m under the linear model. The latter exceeds the 29.1 N·m of the datasheet [38], in the expected direction: the linear model uses the unsaturated magnet flux and therefore overestimates torque at high current, consistently with Section 3.1.5.

Figure 3.18 shows the resulting trajectory against the $i_d = 0$ axis it replaces, and Figure 3.19 the polynomial inversion together with its error. The two remain close throughout: the current angle grows from zero at no load to only 6.6° at peak current, which is the graphical expression of the weak saliency of this machine. In the lower panel of Figure 3.19 the error oscillates between -2.2 and $+1.0$ mA with three extrema, the signature of a well-conditioned third-degree fit.

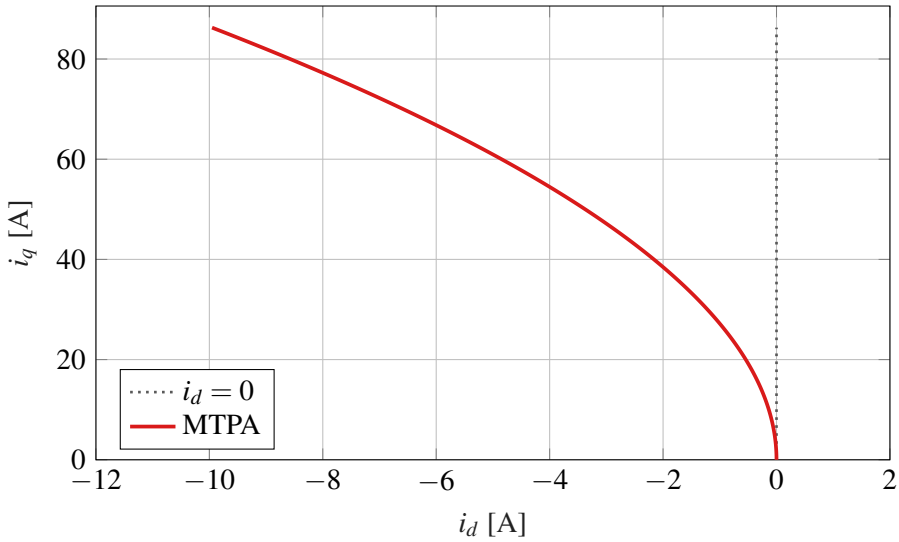


Figure 3.18: MTPA trajectory in the dq current plane, compared with the $i_d = 0$ axis.

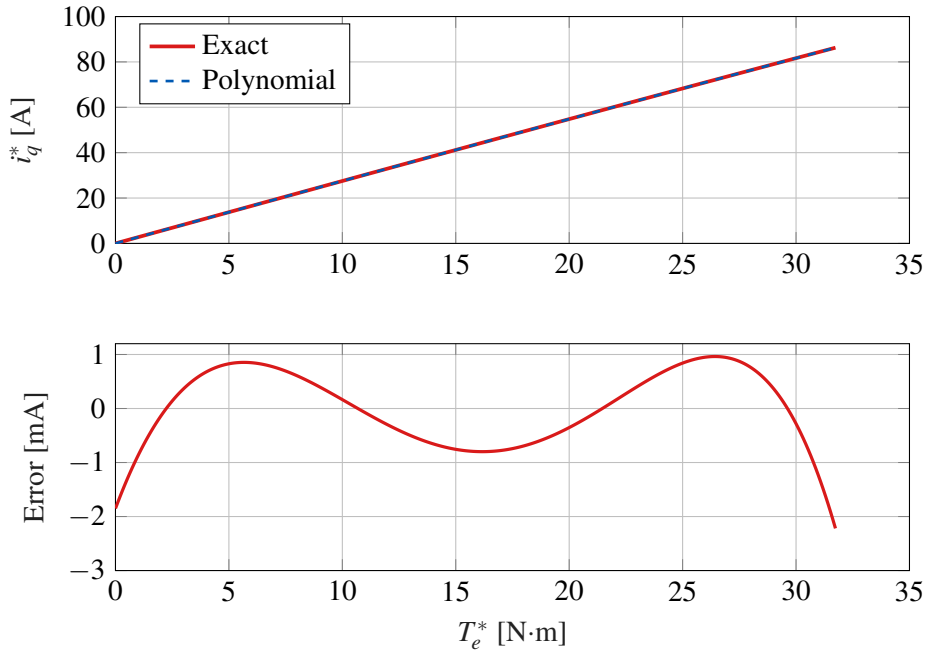


Figure 3.19: Polynomial inversion of the torque-to-current relation and the corresponding error.

Chapter 4

Results and Discussion

4.1 Open-Loop V/f Control

Before implementing field-oriented control, an open-loop scalar (V/f) scheme was applied to the drive, in order to verify that the complete simulation chain — voltage references, space-vector modulator, inverter and FEM-parametrised machine model — operates correctly as an integrated system.

V/f control suits this purpose because it requires neither rotor position feedback nor current regulation: the stator voltage amplitude is imposed as a fixed function of the applied electrical frequency, so that the stator flux magnitude remains approximately constant, and the machine follows the applied frequency with no closed loop acting on it. A fault anywhere in the chain — in the modulator, in the routing of the gate signals, or in the parametrisation of the machine — would prevent the machine from turning.

4.1.1 Results

Figure 4.1 shows the commanded and measured shaft speed during the frequency ramp, and Figure 4.2 the corresponding phase currents.

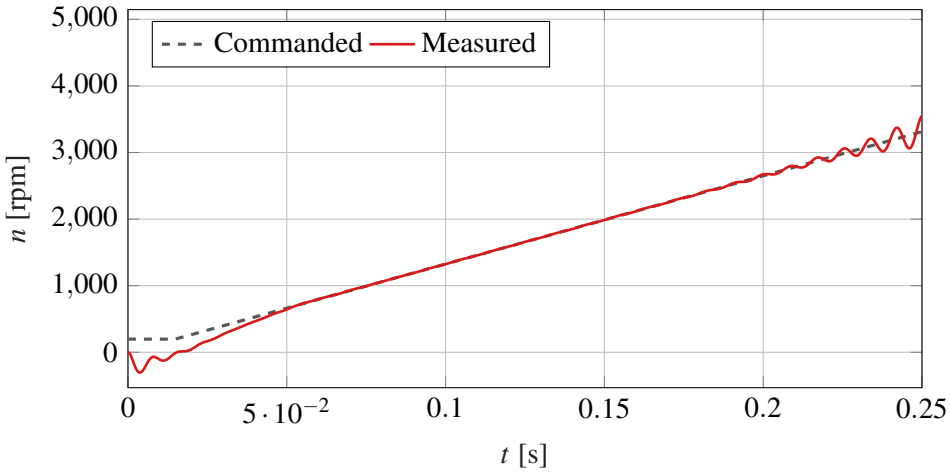


Figure 4.1: Commanded and measured shaft speed during the open-loop V/f ramp.

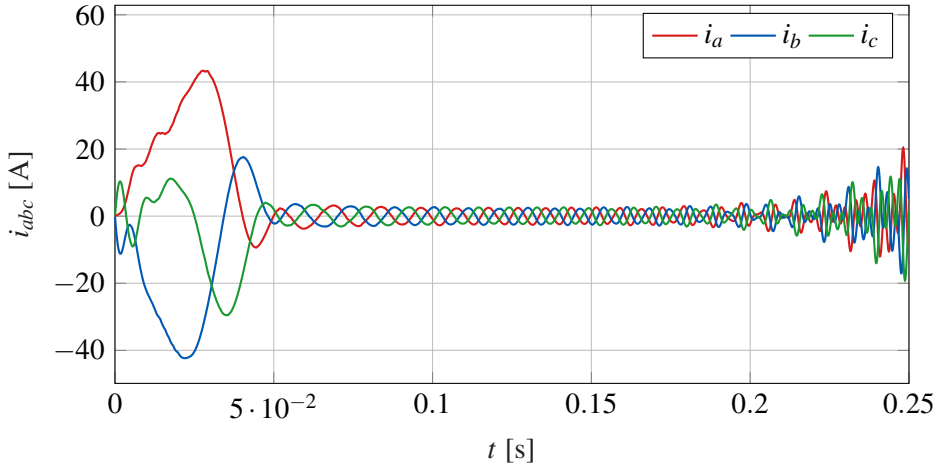


Figure 4.2: Phase currents during the open-loop V/f ramp.

Three regions are visible. The first is the start-up transient, lasting some 50 ms, in which the rotor briefly turns backwards and the phase currents reach 40 A, roughly an order of magnitude above their steady-state value. This is expected of an open-loop start: the initial rotor position is unknown to the modulator, so the stator field is applied at an arbitrary angle relative to the magnets and the rotor is pulled into alignment before it begins to follow the rotating field. Since no current regulator is present, the current that flows during this alignment is set only by the impedance of the machine.

In the second region the machine follows the commanded ramp closely, and the phase currents are balanced and sinusoidal at approximately 3 A. This is the result the test exists to establish, and it confirms three things at once. The modulator produces a correct rotating voltage vector, since an incorrect one would not sustain synchronous rotation. The gate signals are routed consistently to the six inverter switches, since an inverted or transposed pair would apply the complement of the intended voltage. And the FEM-parametrised machine model responds to a three-phase supply

as a synchronous machine should, developing torque and holding synchronism with the applied frequency. The drive chain is therefore validated as an integrated system, and any behaviour observed in the closed-loop tests that follow can be attributed to the controller rather than to the underlying model.

In the third region, above approximately 3000 rpm, an oscillation appears in both the speed and the currents, and grows as the ramp proceeds. This is a known property of open-loop V/f control of permanent-magnet machines rather than a defect of the implementation: the rotor oscillates about its synchronous position, and a permanent-magnet machine has no damper winding to dissipate that oscillation, so above a certain applied frequency the effective damping becomes negative and the oscillation grows [22]. Stable scalar control of a PMSM is possible, but requires an additional stabilising loop, typically a correction of the applied frequency derived from the variation of active power [22].

4.2 Current Controller

4.2.1 Validation Tests

The current loops were validated through five simulation tests, each isolating a single property of the design. All of them replace the inverter by controlled voltage sources, so that the regulator behaviour is observed without the switching ripple of the modulator; the consequences of this simplification are addressed in Section 4.2.1.6. In every test the direct-axis reference is held at zero, and the step is applied part-way through the window so that the steady state preceding it is visible.

Table 4.1 summarises the conditions that distinguish the tests from one another. Tests E1 to E4 impose the shaft speed with an *Ideal Angular Velocity Source* block, which decouples the electrical transient from any mechanical dynamics and makes the measured response attributable to the regulator alone. Test E5 releases the shaft and drives it against a mechanical load, and is therefore the only one in which speed is an outcome rather than a condition.

Table 4.1: Conditions of the current loop validation tests.

Test	Speed	DC bus	Current step	Runs
E1	500 rpm imposed	600 V	$i_q : 0 \rightarrow 5$ A	1
E2	8000 rpm imposed	600 V	$i_q : 0 \rightarrow 5$ A	2
E3	10000 rpm imposed	450 V	$i_q : 0 \rightarrow 20$ A	2
E4	500 rpm imposed	600 V	$i_q : +5 \rightarrow -5$ A	1
E5	free shaft	600 V	$T_e^* : 0 \rightarrow 5$ N·m	1

4.2.1.1 E1 — Nominal Response at Low Speed

The first test establishes the reference behaviour of the loop in the absence of any disturbance. At 500 rpm the back-EMF amounts to only 12.7 V against the 346.4 V available from the bus, and

the cross-coupling term is of the same order as the resistive drop; the loop therefore operates far from any voltage constraint, and the measured response can be compared directly with the design specification.

The measured step response is shown in Figure 4.3. The rise time is 140 μs, the overshoot 4.2 % and the settling time 395 μs, against design values of 169 μs, 4.3 % and 450 μs. The agreement in overshoot is the significant result, since overshoot is set by the damping ratio alone: it confirms that the discretisation preserves the designed $\xi = 0.707$. The choice of discretisation is not incidental here. The forward Euler approximation maps the left half-plane onto a region that shrinks as the sampling interval grows, and consequently displaces the closed-loop poles, an effect that the trapezoidal transformation avoids by mapping the stability boundary correctly [39]. Applied to these gains at $T_s = 25$ μs, forward Euler would have reduced the effective damping to $\xi = 0.606$ and produced an overshoot of 9.1 %, more than twice the specification. The direct-axis current remains at zero throughout, disturbed only by a small transient at the instant of the step.

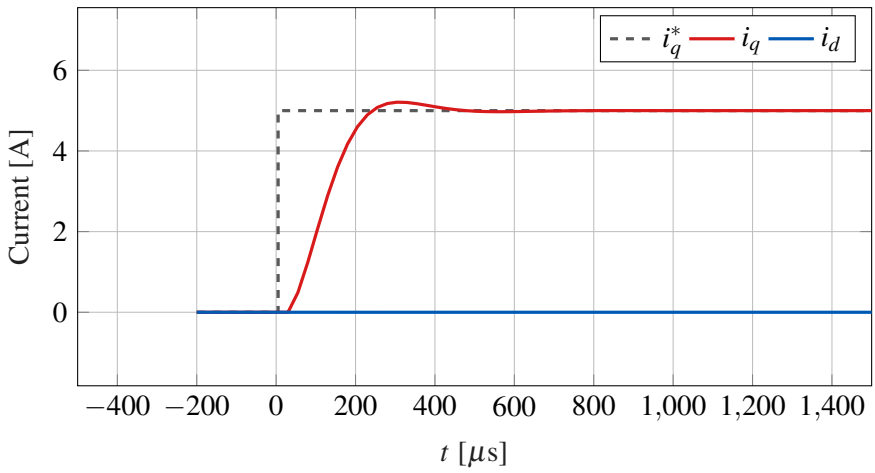


Figure 4.3: Current step response at 500 rpm.

4.2.1.2 E2 — Cross-Coupling and Decoupling Feedforward

The second test isolates the effect of the decoupling feedforward. The step is identical to that of E1, but the speed is raised to 8000 rpm, where the coupling term $\omega_e L_q i_q$ reaches 39.8 V against a resistive drop of 4.0 V, and the back-EMF reaches 203 V. The test is run twice, with and without the feedforward, and nothing else is changed between the two runs.

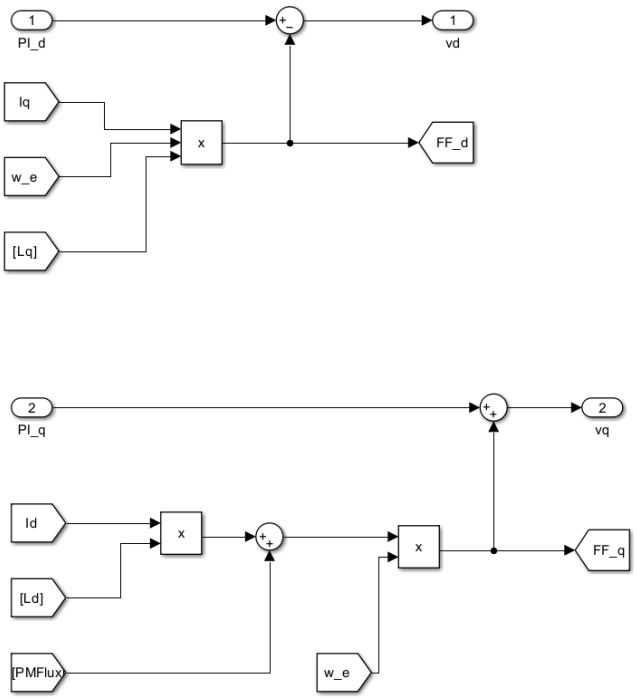


Figure 4.4: Feedforward subsystem on *Simulink*.

The anti-windup is disabled in both runs. Its tracking signal is formed by subtracting the feedforward from the limited voltage, so removing the feedforward while leaving the tracking path active would corrupt the integrator and introduce a second variable into the comparison. At this operating point the voltage does not saturate, and the anti-windup would in any case remain dormant.

Figure 4.5 presents both runs. The maximum direct-axis deviation is 0.000 A with decoupling and 0.062 A without it. The quadrature-axis responses are almost indistinguishable, which is the expected outcome: the back-EMF is a constant disturbance at fixed speed, and the integral action of the regulator rejects it in steady state whether or not it is cancelled in advance. The role of the feedforward is thus to remove the disturbance from the transient rather than from the steady state [31], and its effect is visible not on the q axis but on the d axis, where it prevents a step in torque-producing current from disturbing the flux-producing one. The deviation scales with both current and speed, so the figure quoted here is a lower bound on the effect at rated conditions.

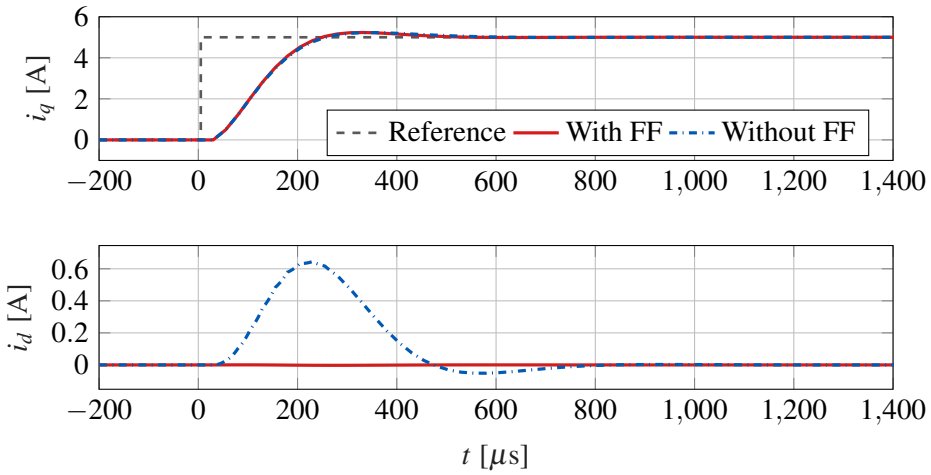


Figure 4.5: Current step at 8000 rpm with and without decoupling feedforward.

4.2.1.3 E3 — Voltage Saturation and Anti-Windup

The third test exercises the voltage limiter and the anti-windup, which only become relevant when the regulator demands more voltage than the bus can supply. At the nominal 600 V bus this condition is difficult to sustain: at 8000 rpm a step to 80 A requires only 213 V in steady state, and the loop leaves saturation after some 50 μs, too brief for the integrator to accumulate any appreciable error.

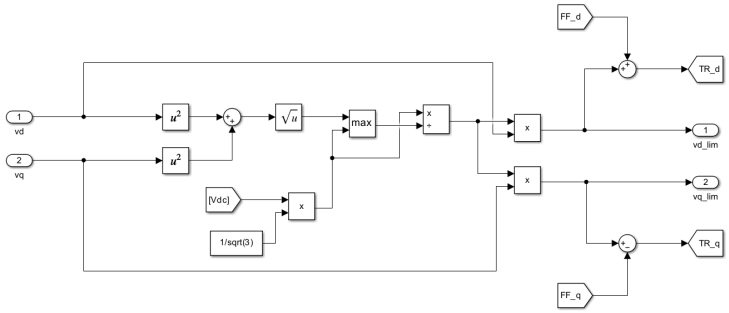


Figure 4.6: Voltage limiter on *Simulink*.

The bus is therefore reduced to 450 V for this test, giving $V_{max} = 259.8$ V, and the speed raised to 10000 rpm. This is a deliberate stress test: it emulates the voltage-limited operation encountered near base speed, where the back-EMF consumes most of the available voltage and the regulator is forced against its limit for an extended period [31]. At this operating point the back-EMF alone amounts to 253 V, leaving approximately 6 V of headroom, which bounds the steady-state current at roughly 26 A. The step is set to 20 A so that it remains reachable in steady state while still saturating the loop throughout the transient.

The reduced headroom is also what makes the current rise slowly, over approximately 2 ms, in Figure 4.7. This is not a property of the regulator but the signature of voltage-limited operation:

with only 6 V across the inductance, the attainable di/dt is correspondingly small. The applied voltage magnitude in the lower panel remains pinned at V_{max} for 1825 μs and releases the limit precisely when the current reaches its reference.

The test is run twice, with and without anti-windup, the feedforward remaining active in both. The anti-windup follows the tracking formulation of [40], in which the difference between the commanded and the actually applied output is fed back around the integrator through a tracking gain, driving the integrator towards a value consistent with the saturated output. The tracking gain is set to $K_t = K_i/K_p \approx 9000$, that is, a tracking time constant equal to the integral time of the regulator, which is the value recommended for a PI controller in [40]. This places the tracking dynamics within the bandwidth of the loop; with the default unity gain the tracking path acts far too slowly to restrain the integrator over the timescale of the transient, and the two runs become indistinguishable.

With anti-windup the current settles on the reference with an overshoot of 1.8 %. Without it the integrator accumulates throughout the saturation interval, and the current crosses the reference and continues to diverge without recovering within the observation window; no overshoot figure is quoted for that run, since the response does not converge and the quantity is undefined.

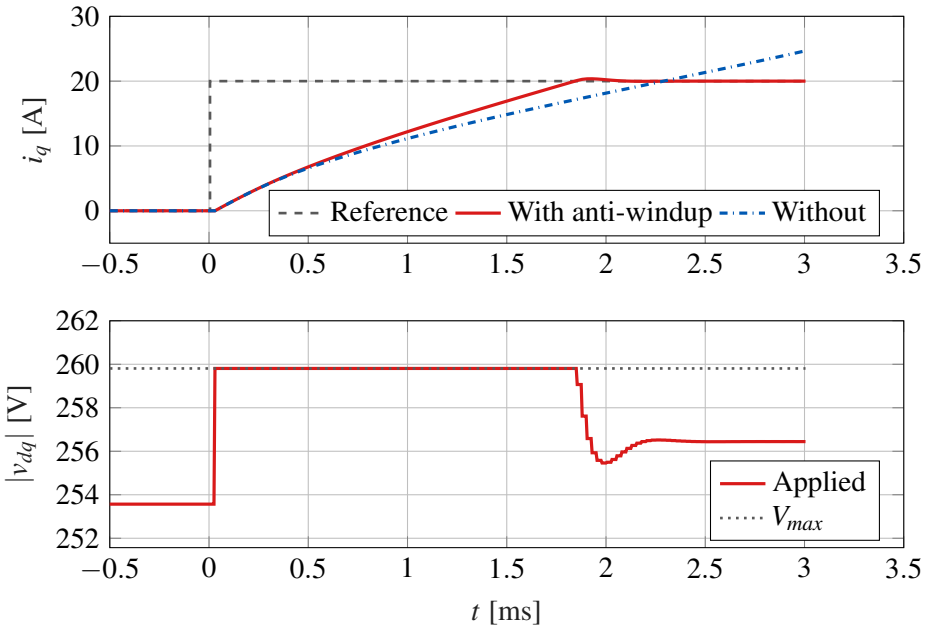


Figure 4.7: Current step into sustained voltage saturation, with and without anti-windup.

4.2.1.4 E4 — Current Reversal

The fourth test reverses the quadrature reference from +5 A to −5 A at 500 rpm, taking the loop through zero current and from motoring into generating operation. The low speed leaves ample voltage headroom, so the transition is limited by the regulator rather than by the bus, and the same conditions as E1 apply in every other respect.

Figure 4.8 shows the result. The reversal is fast and symmetric about zero, with no hesitation at the crossing, and the direct-axis current is held at zero throughout. This behaviour is what allows the drive to transition between traction and regenerative braking without a discontinuity in the delivered torque.

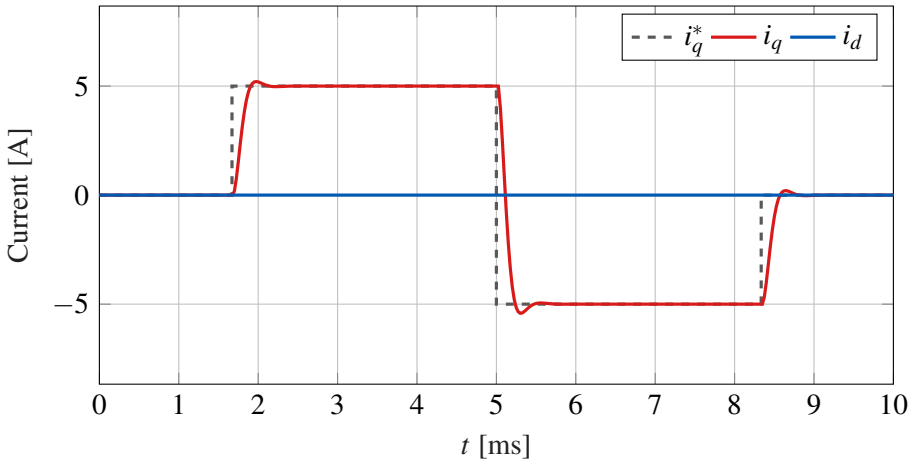


Figure 4.8: Reversal of the quadrature current reference at 500 rpm.

4.2.1.5 E5 — System-Level Behaviour and Load Rejection

The final test releases the shaft and evaluates the loop as part of the drive. The mechanical arrangement consists of the rotor inertia of 0.33×10^{-3} kg·m^2 [38] and a rotational damper of $b = 8 \times 10^{-3}$ N·m·s/rad. The damper does not represent a physical friction: it bounds the shaft speed to an observable range within the simulation window, giving a mechanical time constant $J/b = 41$ ms and an equilibrium speed of T_e/b. Torque is commanded rather than current, the reference being converted by $i_q^* = T_e^*/k_t$ with $k_t = 1.5\, p\, \psi_{PM} = 0.363$ N·m/A at zero direct-axis current, which follows from the torque expression of the dq model [30].

A torque reference of 5 N·m is applied at 20 ms, and a load of 2 N·m is imposed between 120 ms and 170 ms. Figure 4.9 shows the torques and the resulting speed, and Figure 4.10 the corresponding currents.

The developed torque reaches 5 N·m immediately at 20 ms and remains at that value throughout, including across the load step. Only the speed responds to the disturbance, falling from approximately 5400 rpm to 4100 rpm as the equilibrium moves from T_e/b to $(T_e - T_{load})/b$, and recovering once the load is removed. The quadrature current stays at 14 A and the direct-axis current at zero for the entire test.

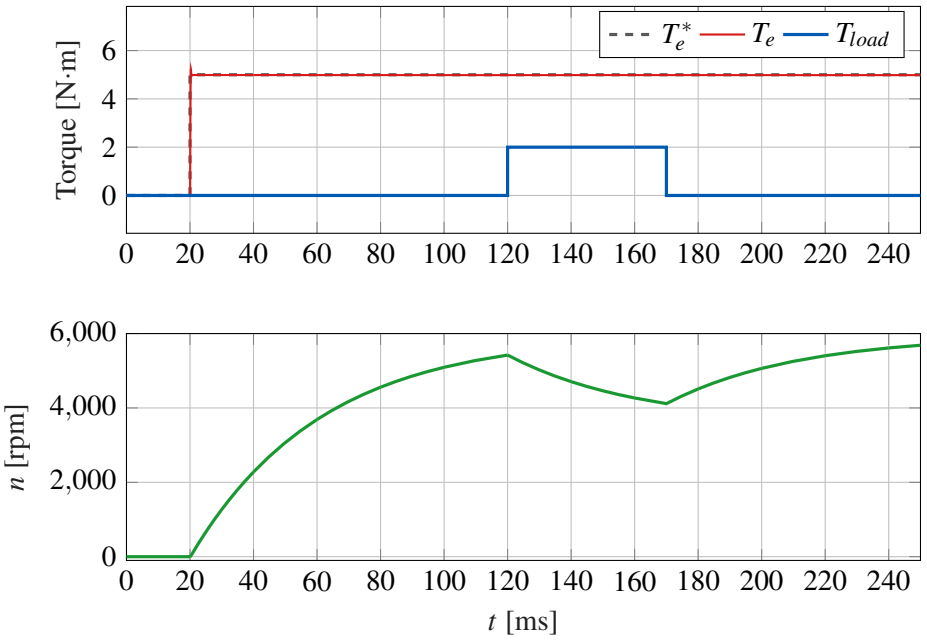


Figure 4.9: Torque tracking and speed response to a load step on a free shaft.

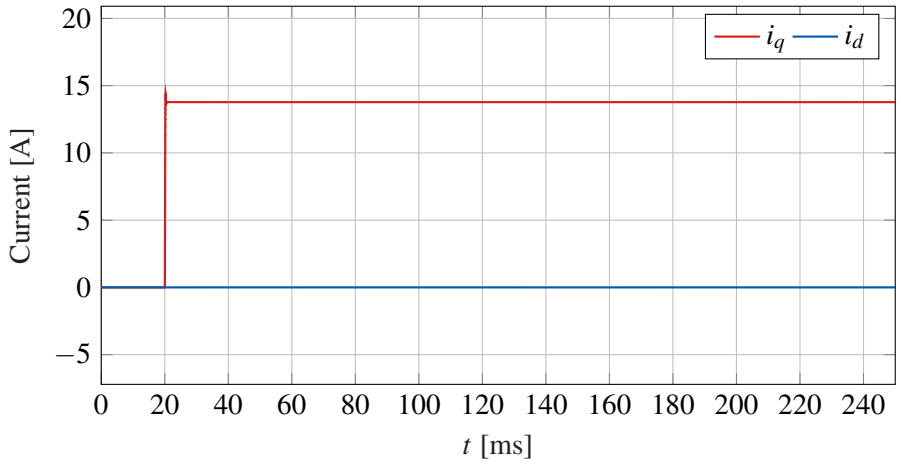


Figure 4.10: Direct and quadrature currents during the system-level test.

4.2.1.6 Switching Ripple

All the tests reported above use ideal controlled voltage sources in place of the inverter, so the currents they show are free of switching ripple. Assessing the ripple requires the switching inverter and its modulator in the loop, and is therefore reported separately in section 4.3.3. The regulator design is unaffected by this simplification, since the modulator delay it introduces was already accounted for in the phase margin of Table 3.5.

4.3 Maximum Torque per Ampere

4.3.1 Verification

Three tests were carried out, addressing successively the derivation, the implementation and the interaction with the machine.

4.3.1.1 M1 — Optimality of the Trajectory

Setting the derivatives of (3.29) to zero identifies a stationary point but does not by itself establish that it is a minimum. Since the whole strategy rests on that claim, it is checked by an independent route. Solving (2.5) for the quadrature current gives the locus of all current pairs producing a given torque,

$$i_q = \frac{T_e}{\frac{3}{2}p\left[\psi_{PM} + (L_d - L_q)\, i_d\right]}, \tag{4.1}$$

so sweeping i_d traces every admissible operating point for that torque, and the minimum of $|i_s|$ along the sweep can be located numerically without reference to the derivation.

The upper panel of Figure 4.11 shows this sweep for three torque levels, with the analytical trajectory point marked on each curve. The marked point coincides with the numerical minimum in all cases, the discrepancy being below 0.25 mA, of the order of the resolution of the sweep. The curves are very flat around the optimum, and increasingly so at high torque: at 30 N·m the current magnitude stays within 1 % of its minimum over some 20 A of i_d, which is why the minima are barely perceptible on the absolute scale of the figure and are identified by the markers. This flatness is useful in practice, since a moderate error in the direct-axis reference costs very little in current, but it is also what limits the benefit of the strategy.

The lower panel quantifies that benefit as the reduction in current magnitude relative to operating at $i_d = 0$. It grows monotonically with load, reaching 0.68 % at peak torque. The reason is the weak saliency: the reluctance term contributes 0.42 N·m of the 31.7 N·m available, and the direct-axis current adds to the magnitude in quadrature, so the current saved on the q axis is largely offset by the current spent on the d axis.

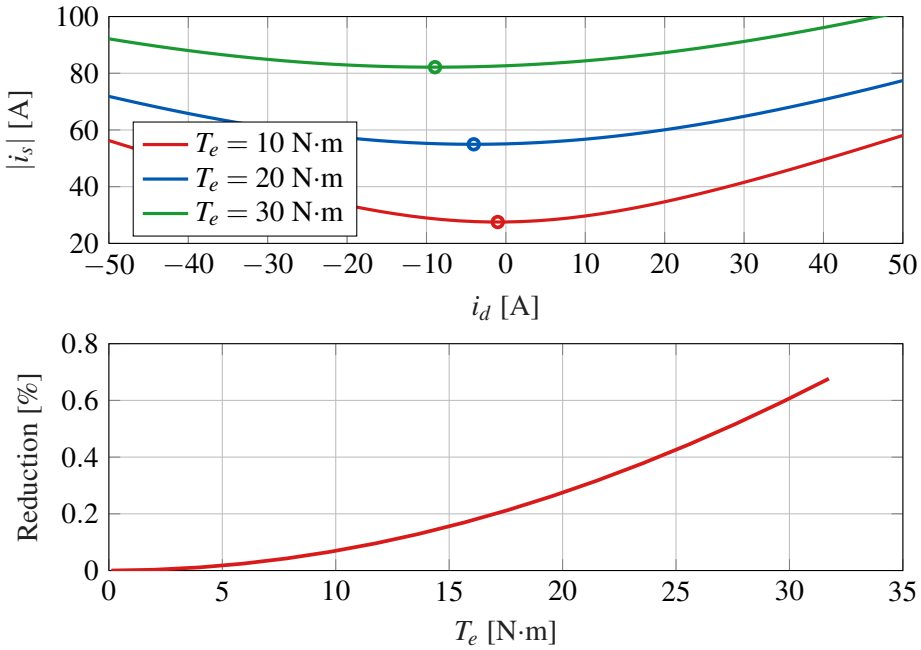


Figure 4.11: Coincidence of the analytical trajectory with the minimum of the current magnitude, and the resulting current reduction.

4.3.1.2 M2 — Reference Generation

The second test verifies the block implementation in isolation, with no machine and no current loops. A slow torque ramp from -35 to $+35$ N·m is applied to the generator and the two references recorded, then compared with the analytical trajectory evaluated at the same torque. The ramp deliberately exceeds $\pm T_{e,\max}$ so that the saturation blocks are exercised at both ends.

Figure 4.12 shows the analytical references, those produced by the chain, and the difference between them. The maximum deviation is 2.24 mA in i_q^* and 0.51 mA in i_d^*; the former is exactly the polynomial fitting error established above, which means the chain introduces no error of its own. The saturation plateaus agree to 2.2 mA, confirming that the limits are correctly placed.

The symmetry properties were checked numerically, since they are what the construction for negative torque depends on. The direct-axis reference is even in T_e to within 10^{-4} mA and never becomes positive, confirming that the $|\cdot|$ and sign operations are correctly arranged; a sign error here would produce a magnetising current under braking and would not be evident from the trajectory plot alone. The quadrature reference is odd to within 0.93 mA rather than exactly, a residue that arises because the fitted polynomial does not pass exactly through the origin, a_0 being -1.86 mA, so the reconstruction by $\operatorname{sign}(T_e)\, i_q^*(|T_e|)$ leaves a discontinuity of twice that value at zero torque. At sub-milliampere level this has no consequence for the drive.

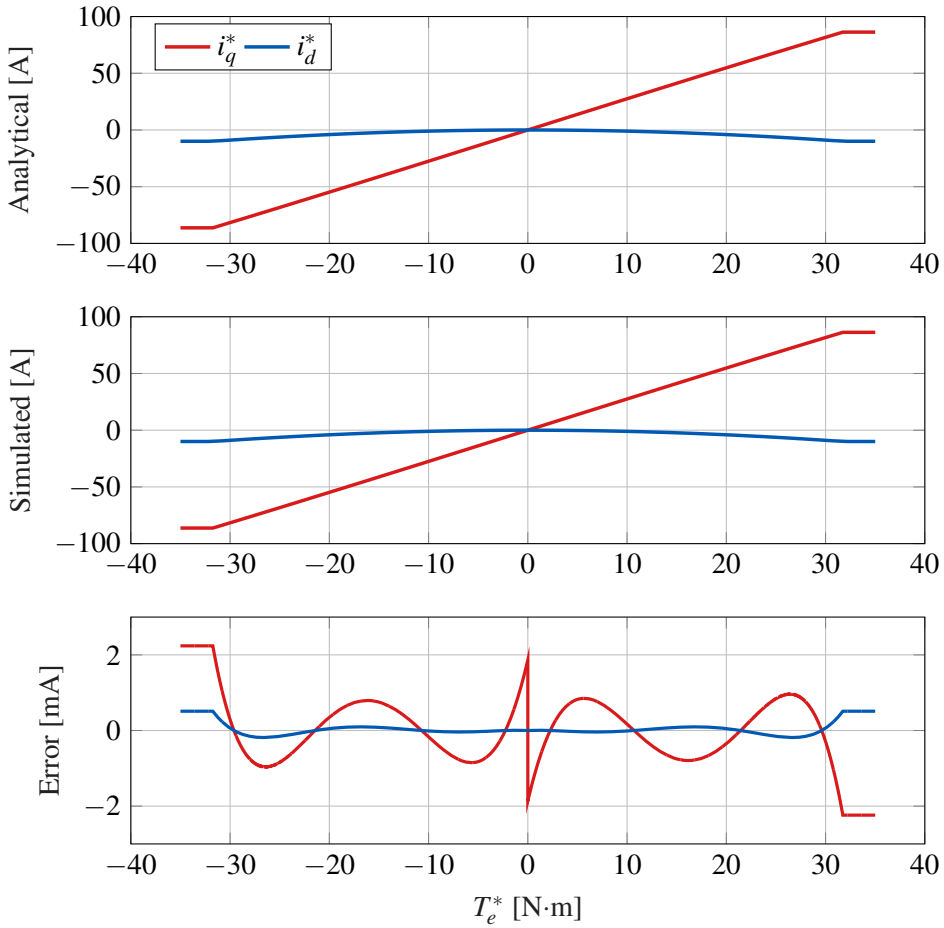


Figure 4.12: References produced by the block chain compared with the analytical trajectory.

4.3.1.3 M3 — Delivered Torque

The final test closes the loop around the machine model. The reference generator drives the current controllers of Section 4.2.1, which in turn drive the finite element machine model through ideal voltage sources, with the shaft speed imposed at 500 rpm so that no voltage limitation interferes with the measurement. A staircase of torque references from 5 to 30 N·m in steps of 5 N·m is applied, each level held for 50 ms, and the delivered torque measured in steady state at each level.

The results are given in Table 4.2 and Figure 4.13. Two observations follow. First, the measured currents match the analytical trajectory at every level: at 30 N·m the drive settles at $i_d = -8.93$ A and $i_q = 81.66$ A, against -8.93 A and 81.7 A predicted. The reference generator and the current loops are therefore doing exactly what they were designed to do. Second, the delivered torque falls consistently short of the command, by 0.23 % at 5 N·m and by 6.40 % at 30 N·m.

The error is systematic, always negative, and grows monotonically with load, which identifies its origin unambiguously. The trajectory and the torque constant are evaluated with the unsaturated magnet flux linkage, whereas the flux in the machine model falls by some 9 % across the current

range. Since torque is approximately proportional to the flux linkage, the machine produces less torque than the linear model predicts, and the shortfall tracks the saturation. This is not a defect of the reference generator but the practical measure of the parameter sensitivity anticipated in Section 3.5.2, and it is the quantity that a trajectory derived from the finite element data would recover.

Table 4.2: Delivered torque against command at 500 rpm.

T_e^* [N·m]	T_e [N·m]	Error [N·m]	Error [%]	i_d [A]	i_q [A]
5	4.988	−0.012	−0.23	−0.26	13.77
10	9.895	−0.105	−1.05	−1.02	27.51
15	14.675	−0.325	−2.17	−2.29	41.19
20	19.330	−0.670	−3.35	−4.05	54.80
25	23.809	−1.191	−4.76	−6.27	68.29
30	28.081	−1.919	−6.40	−8.93	81.66

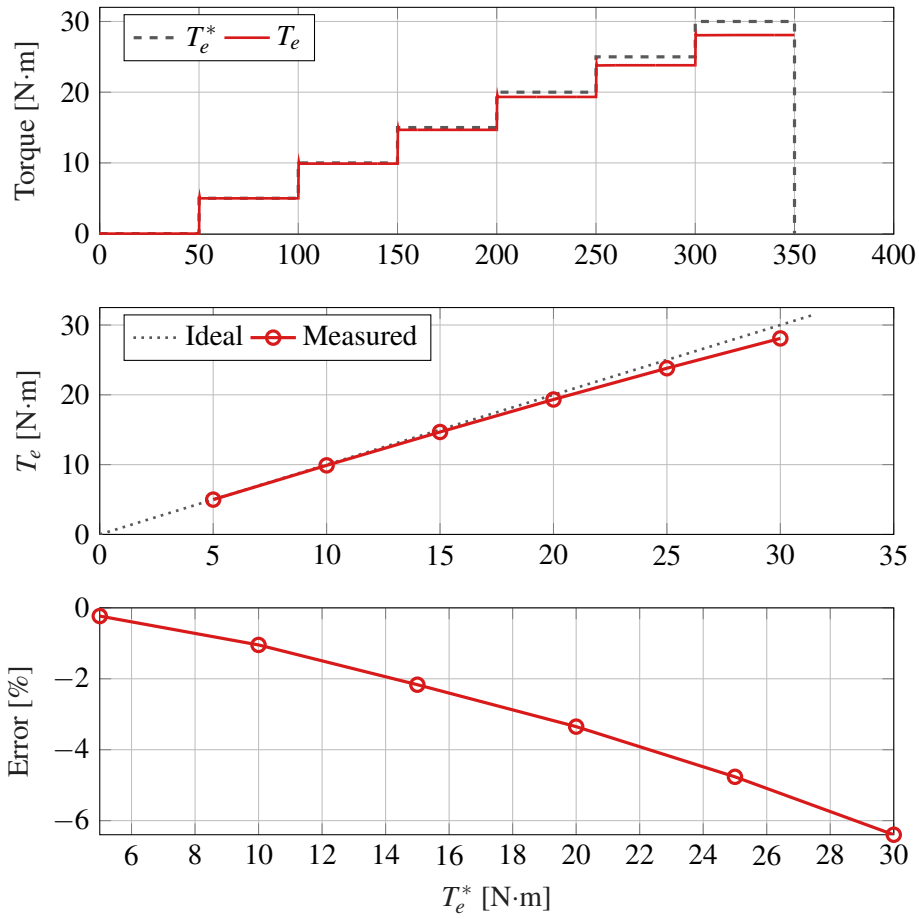


Figure 4.13: Delivered torque against command over the torque range, and the resulting relative error.

4.3.2 Scope and Limitations

The implementation described here is the classical analytical formulation of MTPA. Within its own assumptions it is exact, as M1 and M2 confirm.

Its limitation is the one M3 measures. The constant-parameter trajectory is optimal only near the condition at which the parameters were extracted, and the departure grows precisely where it matters most, at high current. Two remedies are established in the literature. The trajectory may be derived numerically from the finite element data and stored as a one-dimensional map indexed by torque, which would preserve the interface used here and substitute directly for the analytical block; or the optimum may be tracked online by signal injection methods, which locate the MTPA point from the machine response and are therefore insensitive to parameter drift altogether [41]. The former is the natural choice here, since the finite element dataset is already available and has been used for the machine model itself. It was not implemented in this work and is identified as its natural continuation.

The practical significance of these results should be interpreted in terms of torque accuracy rather than current minimisation. The torque shortfall reaches 6.4 % at peak, whereas the current-minimisation benefit of MTPA itself is only 0.68 % on this machine. Therefore, the main practical contribution of accounting for saturation is not an improvement in efficiency, but greater accuracy in the torque delivered. A saturated trajectory would consequently be valuable primarily for reducing torque prediction errors, rather than for improving the already marginal current-minimisation benefit.

4.3.3 Complete System

All the tests reported so far replaced the inverter by controlled voltage sources, which isolates the regulator from the modulator but produces currents that are free of switching ripple. This test closes that gap: the modulator and the two-level inverter are placed in the loop, supplied from the battery model according to a Formula Student Accumulator, and the drive is run at a fixed operating point so that the ripple can be measured in steady state.

The machine is driven at 3000 rpm with the speed imposed, which corresponds to an electrical frequency of 200 Hz, and a torque reference of 15 N·m is applied at 20 ms. The switching frequency is 40 kHz, so there are 200 switching periods per electrical cycle. All signals are recorded at 10^{-7} s, giving 250 samples per switching period, and none of them is filtered: the ripple is the quantity under measurement.

4.3.3.1 Current Ripple

Figure 4.14 shows the phase current, first over several fundamental periods and then magnified to four switching periods. At the scale of the upper panel the ripple is not distinguishable from the thickness of the line; the lower panel resolves it into the sawtooth characteristic of a two-level inverter, in which the current rises while an active vector is applied and falls during the zero vectors, with the underlying slope being that of the fundamental itself.

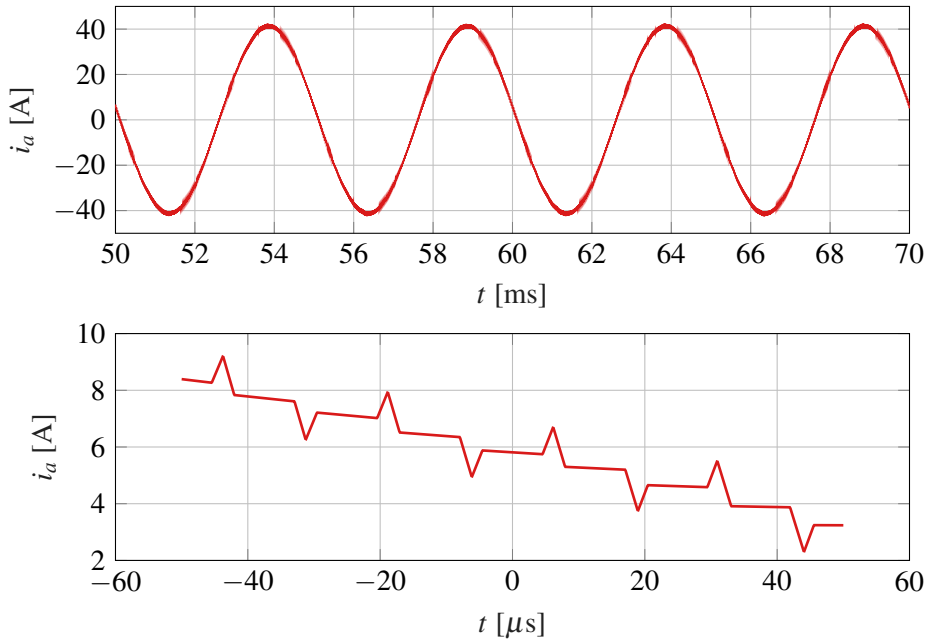


Figure 4.14: Phase current over several fundamental periods and magnified to four switching periods.

The measured peak-to-peak ripple is 6.85 A against a fundamental peak of 42.5 A. The magnitude is consistent with what the machine parameters predict: neglecting resistance, the current slope during a zero vector is set by the back-EMF, which at this speed is 76 V, giving $\mathrm{d}i/\mathrm{d}t = 204$ kA/s and therefore an excursion of 5.1 A over one switching period, while the usual estimates of the form $V_{dc}T_s/nL$ place the peak-to-peak value between 4.4 and 8.7 A for n between 8 and 4. The measured value falls inside that range.

The magnitude is nevertheless large in relative terms, and the reason is the very low inductance of this machine. With $L_q = 0.372$ mH the current responds quickly to any applied voltage, which is what makes the fast current loop of Section 3.4.1 possible in the first place, but it equally means the current follows the switching waveform closely between commutations. Low inductance and low ripple are conflicting requirements, and this machine is designed for the former.

It should be stressed that the ripple is an absolute quantity, essentially independent of load, so the relative figure depends entirely on the operating point at which it is quoted. At the 15 N·m of this test the ripple amounts to 16 % of the fundamental peak, but the same absolute ripple corresponds to about 8 % at rated current and to nearly 50 % at light load. The relative measure is therefore worst where the currents are smallest and the losses least significant.

That last observation is what places the ripple in perspective. Losses depend on the RMS value, and the ripple adds to the fundamental in quadrature rather than arithmetically. For a triangular ripple the RMS contribution is $\Delta i_{pp}/2\sqrt{3} = 1.98$ A, which raises the total RMS current from 30.03 to 30.10 A, an increase of 0.22 %. The associated additional conduction loss in the stator is 1.5 W against the 341 W produced by the fundamental, that is 0.4 %. A peak-to-peak ripple of 16 % therefore costs well under one percent in conduction loss, and the more relevant consequences of

switching ripple lie elsewhere, in the additional iron losses, in the acoustic and electromagnetic emissions, and in the torque ripple examined next.

4.3.3.2 Torque Ripple

Figure 4.15 shows the electromagnetic torque over the same four switching periods, together with the direct and quadrature currents. The torque ripple is considerably smaller than the phase current ripple would suggest, and the lower panel explains why: the quadrature current, which is what determines the torque, carries much less ripple than the phase current does. The Park transformation combines the three phase currents, and because the three switching functions are displaced from one another, a substantial part of the ripple is common to the three phases and cancels in the transformation. What survives appears in the dq frame at the switching frequency and its multiples, and it is only this residue that reaches the torque, scaled by k_t.

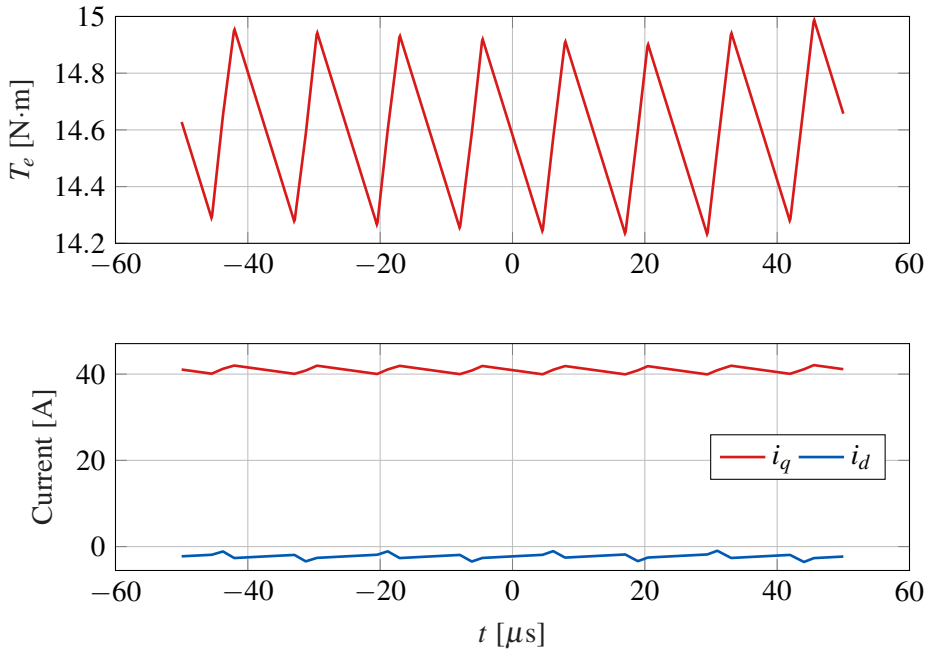


Figure 4.15: Electromagnetic torque and dq currents over four switching periods.

Two checks confirm that the measurement is internally consistent. The torque ripple of 0.754N·m agrees with the value $k_t \Delta i_q = 0.772$ N·m obtained from the quadrature ripple alone, the small difference being the contribution of the reluctance term. And the ripple divides between the two axes in inverse proportion to their inductances: the ratio $\Delta i_d / \Delta i_q$ is 1.20 against the $L_q / L_d = 1.28$ that would be expected, the direct axis carrying more ripple precisely because it is the less inductive of the two.

This is a useful result for the application. The current ripple is a matter for the inverter and the machine losses, whereas it is the torque ripple that the driveline and ultimately the driver experience, and the two are not the same quantity: the vector modulation of a three-phase machine attenuates the former considerably before it becomes the latter.

4.3.3.3 DC Bus

For this test the accumulator is modelled with the Simscape *Battery* block, with a nominal voltage of 600 V, a minimum voltage of 360 V and an internal resistance of 614.4 mΩ, the parameters of the FSFEUP03 accumulator. Charge is treated as infinite, since the state of charge does not change appreciably over the duration of the tests reported here. The bus capacitance is 80 μF [2], and supplies the high-frequency component of the inverter current so that the battery sees only the average.

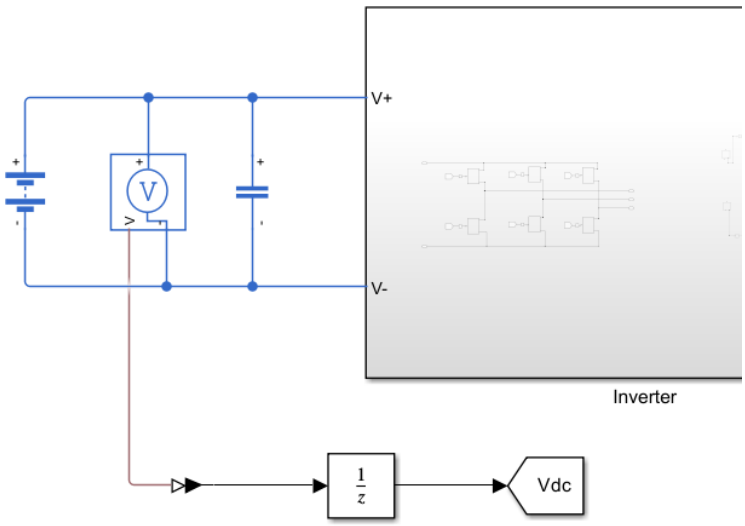


Figure 4.16: DC Bus model on *Simulink*

Because the converter is supplied from the battery model rather than from an ideal source, the bus voltage is itself a result of the test. Figure 4.17 shows it alongside the torque, so that cause and effect are read on the same time axis.

The torque step draws current from the pack, and the bus voltage falls with it, from 516.7 to 512.5 V. The drop of 4.2 V is the voltage across the internal resistance of the battery: at 14.67 N·m and 3000 rpm the mechanical power is 4.6 kW, which corresponds to some 8.9 A drawn from the pack and therefore about 5.5 V across the 614.4 mΩ of internal resistance, the measured value being slightly lower because part of that power is supplied while the bus capacitor discharges.

The drop is not large at this operating point, but it scales with power, and this is the reason the voltage limiter is referred to the measured bus voltage rather than to a fixed value. At the peak power of the drive the drop would be several times greater, and a limiter set to a constant derived from the nominal pack voltage would authorise a voltage that the pack, under load, cannot actually supply.

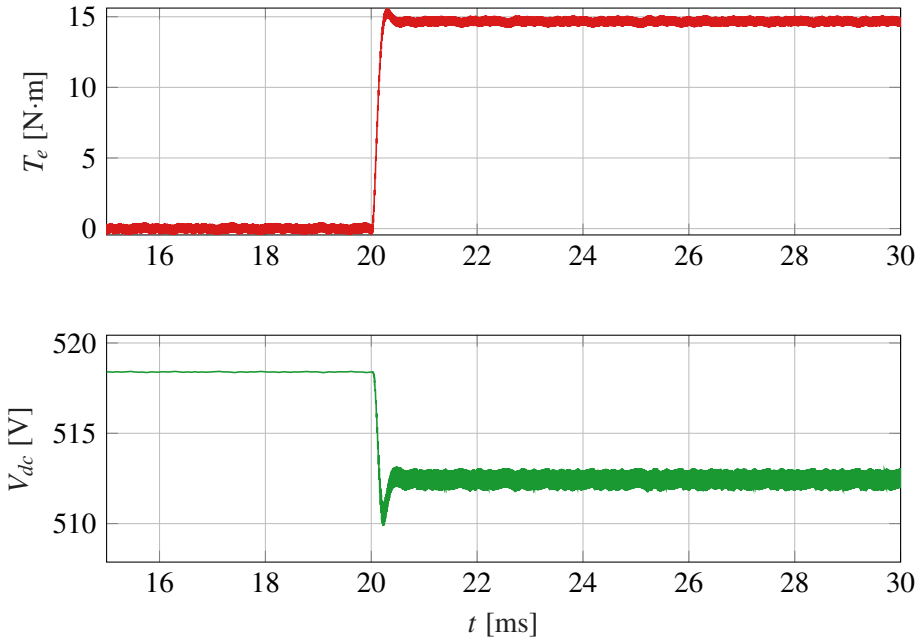


Figure 4.17: Torque step and the resulting drop in DC bus voltage.

4.3.3.4 Frequency Content

Figure 4.18 shows the spectrum of the phase current. The fundamental appears at 200 Hz, and the ripple is concentrated in families of components grouped around discrete frequencies rather than spread across the spectrum. The dominant family is not at the switching frequency but at twice it, at 80 kHz, where the largest component reaches 0.40 A against 0.09 A at 40 kHz. This follows from the modulator using a triangular carrier: each inverter leg switches twice per carrier period, so the ripple repeats at $2f_{sw}$, and the family at the carrier frequency itself is largely suppressed. The effect is favourable, since it places the dominant disturbance twice as far from the control bandwidth as the switching frequency alone would suggest. This grouping is characteristic of carrier-based modulation with a fixed switching frequency, and it is a practical advantage: the disturbance is confined to a narrow and known band, far above the 2 kHz bandwidth of the current loop, so it is neither amplified by the regulator nor confused with the signals it is meant to track. It is also the reason why the feedforward path required a filter once the inverter was introduced, since the flux maps are indexed by the measured currents and would otherwise be driven by the ripple; the cut-off of that filter, at 2 kHz, sits between the loop bandwidth and the switching families.

The absence of low-frequency components is equally significant. Modulation strategies with variable switching frequency, such as hysteresis control, spread their spectral content and can excite mechanical resonances unpredictably; the fixed-frequency space-vector modulation adopted here does not.

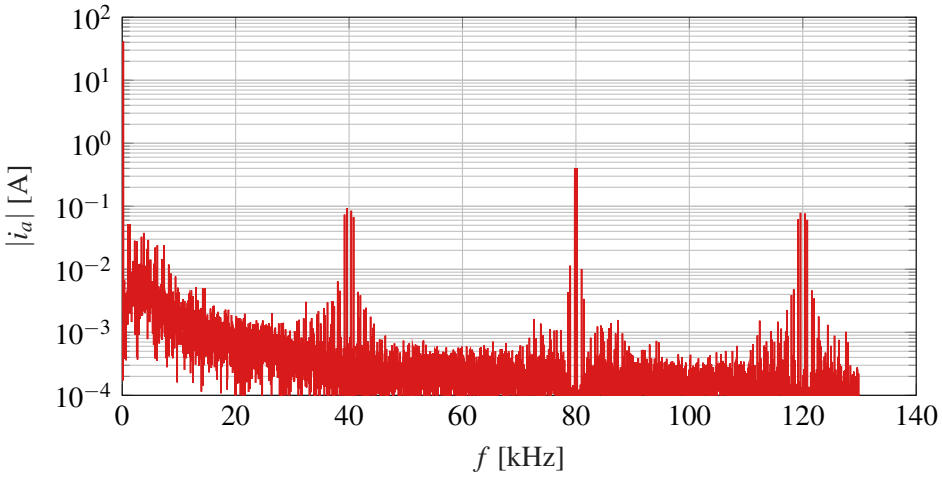


Figure 4.18: Spectrum of the phase current, showing the fundamental and the switching families.

4.3.3.5 Torque Tracking

Beyond the ripple, this test answers the question of whether the control chain survives the introduction of the real modulator. The mean torque delivered in steady state is 14.67 N·m against a command of 15 N·m, an error of $-2.21\,\%$. The corresponding figure obtained in Section 4.3.1.3 with ideal voltage sources, at the same torque, was $-2.17\,\%$.

The difference of four hundredths of a percentage point is the significant result. The shortfall is the one already identified and explained in Section 4.3.1.3, caused by evaluating the MTPA trajectory and the torque constant with unsaturated parameters, and the modulator adds essentially nothing to it. Space-vector modulation, the switching inverter and the battery with its internal resistance therefore do not degrade the accuracy of the torque produced: the current regulators reject the switching disturbance within their bandwidth, and the mean value of the current is unaffected. Every element of the drive designed in this chapter continues to behave as designed when the idealisation of the voltage source is removed.

4.3.4 Simulation with Vehicle

The final test places the drive against the mechanical load of the vehicle, so that the torque control developed in this chapter is exercised in the conditions it was designed for: accelerating the car, coasting, and recovering energy under braking. The purpose is no longer to isolate a property of the regulator but to observe the complete drive operating over a realistic manoeuvre.

The mechanical model consists of the motor driving a fixed gear ratio of 12.6, a wheel of radius 0.203 m, and the translating mass of the vehicle. The mass moved by this drive is taken as one quarter of the total vehicle mass of 270 kg, that is 67.5 kg, because the car is driven by four motors, one at each wheel, and each drives its own share of the mass. The model includes no aerodynamic drag and no rolling resistance; the only forces on the shaft are the motor torque and

the inertia of the vehicle. The consequences of this simplification are examined below and are, in one respect, the most instructive feature of the test.

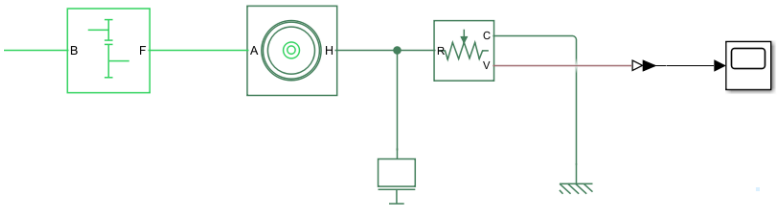


Figure 4.19: Mechanical Load on *Simulink*

Unlike the previous test, this one is run with ideal controlled voltage sources rather than the switching inverter. The manoeuvre lasts more than two seconds, which at the time step required to resolve switching would be prohibitively long to simulate, and the switching behaviour has already been characterised in Section 4.3.3. What is under examination here is the mechanical response, for which the ideal source is an adequate and much faster representation. The DC bus is accordingly held constant.

The torque reference follows a four-phase profile: a demand of $+20$ N·m from 0.05 to 1.0 s to accelerate the car, zero torque from 1.05 to 1.55 s to coast, a demand of -20 N·m from 1.60 to 2.15 s to brake regeneratively, and zero thereafter. Figure 4.20 shows the commanded and delivered torque, the resulting vehicle speed, and the dq currents.

4.3.4.1 Acceleration

During the first phase the car accelerates from rest to 60.1 km/h, which corresponds to 9891 rpm at the motor, at an average of 16.0 m/s^2. The torque is tracked with a mean error of $-4.9\,\%$ over the phase. This is the same shortfall identified in Section 4.3.1.3, arising from the use of unsaturated parameters in the MTPA trajectory, and it is somewhat larger here than the $-3.4\,\%$ measured for the same torque at low speed, because the higher speed of this test drives the machine deeper into saturation, where the real flux linkage departs further from the constant value assumed.

The average acceleration of 16.0 m/s^2 is 1.6 g, which is at or above the limit that tyre grip could sustain on a real car. The *Tire (Simple)* block used here transmits whatever longitudinal force the torque demands without modelling the slip that would occur once that limit is exceeded, so the simulated acceleration is not constrained by adhesion. The figure should therefore be read as the acceleration the drive is capable of commanding, not as one the vehicle could realise on track.

4.3.4.2 Coasting

The coasting phase is the clearest result of the test. With the torque commanded to zero, the speed remains constant: it varies by 0.0002 m/s over the half-second of the phase, a drift of 0.001 %. This

is the direct consequence of the model containing no aerodynamic drag and no rolling resistance. With no torque applied and no resistive force present, there is nothing to decelerate the vehicle, and its speed is held by its own inertia.

This is exactly the expected behaviour, and it serves as a check on the mechanical model in both directions. Had the speed fallen during coasting, an undeclared friction would have been present in the model; had it risen, there would have been a sign error in the torque path. That it remains flat confirms that the mechanical connection is correct and that the only stored quantity is the kinetic energy of the mass. It must be stressed, however, that a real car does decelerate when coasting, and often appreciably; the model is representative of the drive, not of the road behaviour of the vehicle, and its speed and range predictions should not be read as such.

4.3.4.3 Regenerative Braking

In the final phase a negative torque is commanded and the car decelerates from 66.2 to 23.9 km/h. The torque reverses cleanly and is tracked with the same accuracy as during acceleration. The quadrature current follows the torque and reverses with it, while the direct-axis current remains negative throughout, at -3.7 A, essentially the value it held during acceleration at the same current magnitude. This is the practical confirmation, now against the vehicle load, of the symmetry established in Section 3.5.3: because the MTPA trajectory depends on i_q^2, reversing the torque reverses i_q and leaves i_d unchanged, so the same reference generator serves traction and braking without modification.

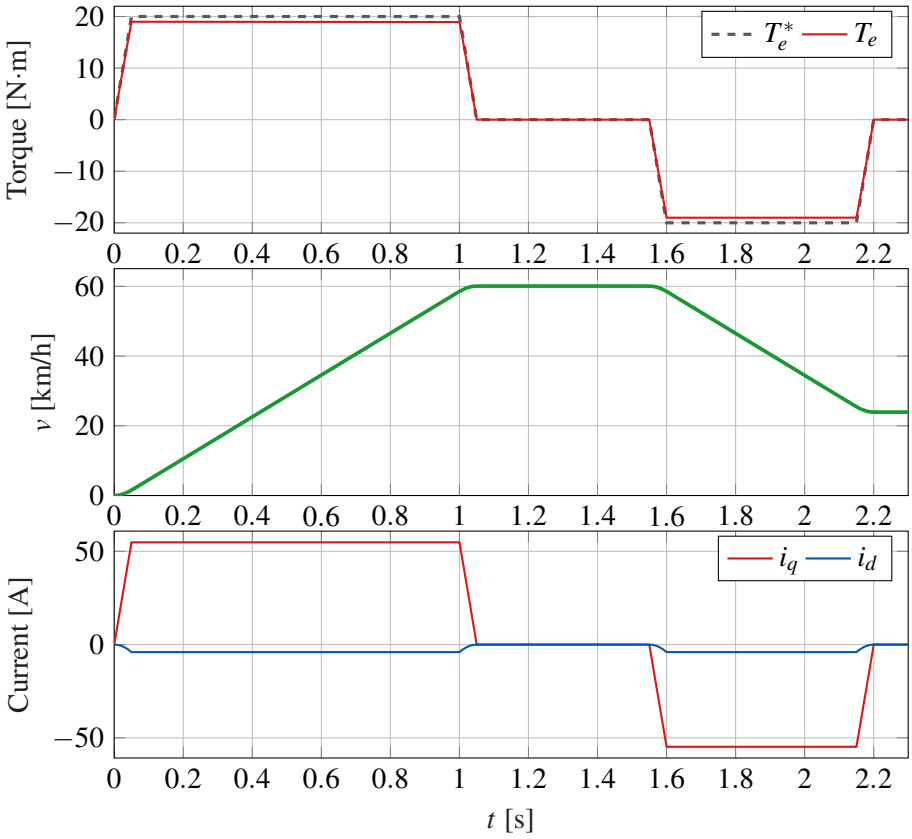


Figure 4.20: Torque, vehicle speed and dq currents over an acceleration, coasting and regenerative braking manoeuvre.

Taken together, the three phases show the drive delivering, holding and reversing torque against a realistic inertial load, with the current references behaving throughout as the design intends. The limitations of the test are those of the mechanical model rather than of the drive: the absence of resistive forces makes the coasting and the achievable acceleration unrealistic, and the fixed bus and ideal sources omit the switching effects quantified separately in Section 4.3.3. Within those bounds, the drive performs the complete traction and braking cycle as designed.

Chapter 5

Conclusions and Future Work

This thesis detailed the development of a motor controller for a Formula Student prototype. The research established the interior Permanent Magnet Synchronous Motor (PMSM) and a two-level three-phase inverter as the selected architecture for the vehicle. To guarantee precise torque delivery and efficiency, a Field-Oriented Control (FOC) scheme was combined with Space-Vector Modulation (SVM) and a Maximum Torque Per Ampere (MTPA) strategy.

The control system was modeled and evaluated in MATLAB/Simulink. By parameterizing the PMSM with Finite Element Analysis (FEA) data, the model successfully captured nonlinear behaviors, including magnetic saturation and cross-coupling. Closed-loop tests demonstrated that the current regulators, decoupling feedforward, and anti-windup mechanisms function correctly together.

5.1 Achievement of Objectives

Overall, the primary objectives set for this project were partially achieved:

- **Theoretical Research and Modeling (Achieved):** The research into power electronics and the Simulink modeling of the powertrain subsystems were successfully completed.
- **Controller Simulation and Tuning (Achieved):** A functional motor controller was implemented and tuned in Simulink, maintaining stable and precise torque delivery across dynamic tests.
- **System Limitations (Needs Improvement):** The simulation environment and the current control scheme have limitations that should be addressed in future iterations:
  - The vehicle's mechanical model currently neglects resistive forces, such as aerodynamic drag and rolling resistance.
  - The MTPA strategy was implemented using unsaturated machine parameters, causing a slight torque shortfall at high currents.

  - Field-weakening control was not implemented; therefore, the operating speed range of this controller is limited compared to the requirements and desired performance of a Formula Student vehicle.

- **Experimental Validation (Not Achieved):** The controller was validated exclusively through simulation and has not yet been experimentally validated on the physical inverter and PMSM. This transition from simulation to the embedded physical prototype remains a critical next step for future developments.

5.2 Future Work

To continue the development of the powertrain and address the current limitations, the following steps are proposed for future work:

- **Field-Weakening Implementation:** Develop and integrate a field-weakening control loop in Simulink to extend the motor's operating speed range, maximizing the vehicle's top speed capabilities.

- **Advanced Control Functionalities:** Implement system-level features, such as vehicle speed limiters and DC power limits, to ensure compliance with Formula Student competition rules and improve overall safety.

- **Mechanical Model Enhancement:** Upgrade the mechanical and vehicle dynamics model in Simulink by incorporating aerodynamic drag, rolling resistance, and tire slip behavior to obtain more accurate and realistic system-level simulations.

- **MTPA Trajectory Optimization:** Improve the MTPA strategy by deriving the trajectory numerically from the saturated finite element data (e.g., using look-up tables or signal injection methods) to eliminate the torque shortfall at high currents.

- **Hardware Transition and Testing:** Transfer the control logic to the embedded microcontroller software. Subsequently, validate the system using Hardware-in-the-Loop (HIL) simulations, followed by physical testbench validation with the real DC-AC converter and PMSM.

References

[1] F. S. Germany, *Formula student rules 2025*, 2025.

[2] P. G. S. Cristóvão, "Development of a DC/AC three-phase power converter optimizing energy efficiency," Dissertação de Mestrado, Faculdade de Engenharia da Universidade do Porto, Porto, Portugal, 2024.

[3] S.-H. Kim, *Electric Motor Control: DC, AC, and BLDC Motors*. Elsevier, 2017.

[4] N. Hashemnia and B. Asaei, "Comparative study of using different electric motors in the electric vehicles," in *18th International Conference on Electrical Machines (ICEM)*, 2008, pp. 1–5. DOI: 10.1109/ICELMACH.2008.4800157

[5] Z. Cao, A. Mahmoudi, S. Kahourzade, and W. L. Soong, "An overview of electric motors for electric vehicles," in *31st Australasian Universities Power Engineering Conference (AUPEC)*, 2021, pp. 1–6. DOI: 10.1109/AUPEC52110.2021.9597739

[6] M. Zeraoulia, M. E. H. Benbouzid, and D. Diallo, "Electric motor drive selection issues for HEV propulsion systems: A comparative study," *IEEE Transactions on Vehicular Technology*, vol. 55, no. 6, pp. 1756–1764, 2006. DOI: 10.1109/TVT.2006.878719

[7] C. C. Chan, "The state of the art of electric, hybrid, and fuel cell vehicles," *Proceedings of the IEEE*, vol. 95, no. 4, pp. 704–718, 2007. DOI: 10.1109/JPROC.2007.892489

[8] R. Carlson, M. Lajoie-Mazenc, and J. C. dos Santos Fagundes, "Analysis of torque ripple due to phase commutation in brushless DC machines," *IEEE Transactions on Industry Applications*, vol. 28, no. 3, pp. 632–638, 1992. DOI: 10.1109/28.137450

[9] T. J. E. Miller, *Switched Reluctance Motors and Their Control*. Magna Physics Publishing and Oxford University Press, 1993.

[10] A. Vagati, "The synchronous reluctance solution: A new alternative in AC drives," in *Proceedings of IECON'94 - 20th Annual Conference of IEEE Industrial Electronics*, vol. 1, 1994, pp. 1–13. DOI: 10.1109/IECON.1994.397741

[11] J. de Santiago et al., "Electrical motor drivelines in commercial all-electric vehicles: A review," *IEEE Transactions on Vehicular Technology*, vol. 61, no. 2, pp. 475–484, 2012. DOI: 10.1109/TVT.2011.2177873

[12] F. Blaabjerg, Ed., *Control of Power Electronic Converters and Systems*. Academic Press, 2018.

[13] S. Kouro, J. I. León, L. G. Franquelo, J. Rodríguez, and B. Wu, "Dc–ac converters," in *The Industrial Electronics Handbook Power Electronics and Motor Drives*. CRC Press, 2011, ch. 14.

[14] J. Millán, P. Godignon, X. Perpiñà, A. Pérez-Tomás, and J. Rebollo, "A survey of wide bandgap power semiconductor devices," *IEEE Transactions on Power Electronics*, vol. 29, no. 5, pp. 2155–2163, 2014. DOI: 10.1109/TPEL.2013.2268900

[15] Z. Dong, Y. Liu, R. Huang, and C. Liu, "Comparative study on topologies of three-level and open-winding converters for PMSM drives," in *1st International Power Electronics and Application Symposium (PEAS)*, 2021.

[16] M. Schweizer and J. W. Kolar, "Design and implementation of a highly efficient three-level T-type converter for low-voltage applications," *IEEE Transactions on Power Electronics*, vol. 28, no. 2, pp. 899–907, 2013. DOI: 10.1109/TPEL.2012.2203151

[17] Typhoon HIL, *Three-phase three-level t-type inverter/rectifier*, https://www.typhoon-hil.com/documentation/typhoon-hil-software-manual/References/three-phase_three-level_t-type_inverter_rectifier.html, Accessed in 2026.

[18] V. Agarwal, "Dc-ac converters (inverters)," in *Encyclopedia of Electrical and Electronic Power Engineering*, Elsevier, 2023, pp. 36–45.

[19] H. W. van der Broeck, H.-C. Skudelny, and G. V. Stanke, "Analysis and realization of a pulsewidth modulator based on voltage space vectors," *IEEE Transactions on Industry Applications*, vol. 24, no. 1, pp. 142–150, 1988. DOI: 10.1109/28.87265

[20] The MathWorks, Inc., *Space vector modulation (SVM) for motor control*, https://www.mathworks.com/discovery/space-vector-modulation.html, Accessed in 2026.

[21] V. M. Bida, D. V. Samokhvalov, and F. S. Al-Mahturi, "PMSM vector control techniques — a survey," in *IEEE Conference of Russian Young Researchers in Electrical and Electronic Engineering (EIConRus)*, 2018, pp. 577–581. DOI: 10.1109/EIConRus.2018.8317164

[22] P. D. C. Perera, F. Blaabjerg, J. K. Pedersen, and P. Thøgersen, "A sensorless, stable V/f control method for permanent-magnet synchronous motor drives," *IEEE Transactions on Industry Applications*, vol. 39, no. 3, pp. 783–791, 2003. DOI: 10.1109/TIA.2003.810631

[23] The MathWorks, Inc., *Open-loop and closed-loop control*, https://www.mathworks.com/help/mcb/gs/open-loop-and-closed-loop-control.html, Accessed in 2026.

[24] I. Takahashi and T. Noguchi, "A new quick-response and high-efficiency control strategy of an induction motor," *IEEE Transactions on Industry Applications*, vol. IA-22, no. 5, pp. 820–827, 1986. DOI: 10.1109/TIA.1986.4504799

[25] M. Abassi, A. Khlaief, O. Saadaoui, A. Chaari, and M. Boussak, "Performance analysis of FOC and DTC for PMSM drives using SVPWM technique," in *16th International Conference on Sciences and Techniques of Automatic Control and Computer Engineering (STA)*, 2015, pp. 228–233.

[26] E. Levi, "Foc: Field-oriented control," in *The Industrial Electronics Handbook Power Electronics and Motor Drives*. CRC Press, 2011, ch. 24.

[27] S. Morimoto, Y. Takeda, T. Hirasa, and K. Taniguchi, "Expansion of operating limits for permanent magnet motor by current vector control considering inverter capacity," *IEEE Transactions on Industry Applications*, vol. 26, no. 5, pp. 866–871, 1990. DOI: 10.1109/28.60058

[28] The MathWorks, Inc., *FEM-parameterized PMSM*, https://www.mathworks.com/help/sps/ref/femparameterizedpmsm.html, Accessed in 2026.

[29] Infineon Technologies AG, *IMZA120R014M1H coolSiC™ 1200 v sic trench mosfet datasheet*, https://www.infineon.com/assets/row/public/documents/60/49/infineon-imza120r014m1h-datasheet-en.pdf, Revision 1.40, 2024.

[30] J. M. d. S. Pinto, "Formula student — sistema de tração," Dissertação de Mestrado, Faculdade de Engenharia da Universidade do Porto, Porto, Portugal, 2014.

[31] S.-K. Sul, *Control of Electric Machine Drive Systems* (IEEE Press Series on Power Engineering). Hoboken, New Jersey: John Wiley & Sons, 2011, vol. 55, ISBN: 9780470590799.

[32] L. Harnefors and H.-P. Nee, "Model-based current control of ac machines using the internal model control method," *IEEE Transactions on Industry Applications*, vol. 34, no. 1, pp. 133–141, 1998. DOI: 10.1109/28.658735

[33] K. Ogata, *Modern Control Engineering*, 5th. Upper Saddle River, New Jersey: Prentice Hall, 2010, ISBN: 9780136156734.

[34] D. G. Holmes, T. A. Lipo, B. P. McGrath, and W. Y. Kong, "Optimized design of stationary frame three phase AC current regulators," *IEEE Transactions on Power Electronics*, vol. 24, no. 11, pp. 2417–2426, 2009. DOI: 10.1109/TPEL.2009.2029548

[35] T. M. Jahns, G. B. Kliman, and T. W. Neumann, "Interior permanent-magnet synchronous motors for adjustable-speed drives," *IEEE Transactions on Industry Applications*, vol. IA-22, no. 4, pp. 738–747, 1986. DOI: 10.1109/TIA.1986.4504786

[36] S. Morimoto, Y. Takeda, and T. Hirasa, "Current phase control methods for permanent magnet synchronous motors," *IEEE Transactions on Power Electronics*, vol. 5, no. 2, pp. 133–139, 1990. DOI: 10.1109/63.53150

[37] S. Morimoto, M. Sanada, and Y. Takeda, "Effects and compensation of magnetic saturation in flux-weakening controlled permanent magnet synchronous motor drives," *IEEE Transactions on Industry Applications*, vol. 30, no. 6, pp. 1632–1637, 1994.

[38] *Motor datasheet TI085-052-070-04B7S-07S04BE2 with fieldweakening*, Project no. 2018-079-2; preliminary, calculated values, Fischer Elektromotoren GmbH, Billigheim-Allfeld, Germany, Sep. 2018.

[39] K. J. Åström and B. Wittenmark, *Computer-Controlled Systems: Theory and Design*, 3rd. Upper Saddle River, New Jersey: Prentice Hall, 1997, ISBN: 9780133148992.

[40] K. J. Åström and T. Hägglund, *PID Controllers: Theory, Design, and Tuning*, 2nd. Research Triangle Park, North Carolina: Instrument Society of America, 1995, ISBN: 9781556175169.

[41] T. Sun, J. Wang, and X. Chen, "Maximum torque per ampere (MTPA) control for interior permanent magnet synchronous machine drives based on virtual signal injection," *IEEE Transactions on Power Electronics*, vol. 30, no. 9, pp. 5036–5045, 2015. DOI: 10.1109/TPEL.2014.2365814

Appendix A

Motor datasheet
TI085-052-070-04B7S-07S04BE2

FISCHER Elektromotoren

Motor Datasheet [calculated Values]

TI085-052-070-04B7S-07S04BE2

with Fieldweaking

High-Speed & Power-Systems

Project-No.: 2018-079-2

	Symbol	Unit	Value	
Rated Data Water cooled (φ = 0°)				
Nominal Torque	T_{NomWC}	Nm	11,1	
Nominal Current	I_{NomWC}	A_{rms}	22,6	
Nominal Speed	n_{NomWC}	rpm	13250	
Nominal Power	P_{NomWC}	W	15404	
Winding Losses[1] / *Total Losses* [1,2]	P_{DWC}	W	254	*617*
Holding Torque	T_{HWC}	Nm	7,9	
Holding Current	I_{HWC}	A_{rms}	16	
Rated Data S6 duty (φ = -10°)				
Torque	T_{S6}	Nm	24,6	
Current	I_{S6}	A_{rms}	51	
Speed	n_{S6}	rpm	12100	
Power	P_{S6}	W	31199	
Winding Losses[1] / *Total Losses* [1,2]	P_{DS6}	W	1282	*1619*
Holding Torque	T_{HS6}	Nm	17,4	
Holding Current	I_{HS6}	A_{rms}	35,9	
Peak Data (φ = -10°)				
Peak Torque	T_{Peak}	Nm	29,1	
Peak Current	I_{Peak}	A_{rms}	61	
Speed at Peak Torque	n_{Peak}	rpm	11600	
Peak Power	P_{Peak}	W	35366	
Winding Losses[1] / *Total Losses* [1,2]	P_{DPeak}	W	1843	*2167*
Data				
Torque Constant	k_t	Nm/A_{rms}	0,492	
BEMF Constant (Phase - Phase)	k_e	V_{rms}/(rad/s) V_{rms}/rpm	0,296 0,031	
Motor Constant	k_m	Nm/√W	0,447	
Idle Speed	n_{idle}	rpm	13650	
max. Speed (Fieldweaking)	n_{max}	rpm	20000	
max. Frequency (Idle/*Fieldweaking*)	f_{max}	Hz	910	*1333*
DC Bus Voltage	U_{DC}	V_{DC}	600	
Ø Resistance per Phase (Winding only)	R_{Ph20}	Ω	0,126	
Ø Inductance per Phase (Winding only)	L_{Ph}	mH	0,393	
electr. Time Constant τ=L/R	τ_{el}	ms	3,11	
Number of Polepairs	n		4	
Winding Connection			Star	

preliminary Data

Motor Datasheet [calculated Values]

TI085-052-070-04B7S-07S04BE2

with Fieldweaking

High-Speed & Power-Systems

	Symbol	Unit	Value
Data Watercooling			
Inlet Temperature of Coolant	T_{in}	°C	10 … 40
Max. Temperature rise of Coolant	T_{max}	K	5
Min. required Coolant flow	Q_{min}	l/min	---
Volume of cooling channel	V_{cool}	l	---
thermal Time Constant	τ_{th}	min	---

Data Mechanics			
Rotor Inertia (assembly set)	J	kgm²	0,33*10^-3
Weight of Motor w/o Housing	m	kg	2,8
Outer Stator Diameter w/o Housing	dA	mm	85
Inner Stator Diameter	dAg	mm	51,6
Length of Stator	l	mm	70

Annotations - Losses

[1] Winding Losses are referred to a Coil Temperature of 100°C.

[2] The total Losses are made up of: Winding Losses; Stator Iron Losses; Rotor Losses;

Calculation of total Losses:

Winding Losses + Stator Iron Losses (at speed X) + Rotor Losses (at speed X)

Annotations - general

Ensure that your servo drive can handle the Nominal- and Peakcurrent of the Motor.
An adjustment of the Speed and DC Bus Voltage can be done after consultation.
The nominal data in this datasheet are based on an ambient/coolant temperature of 20°C
The stated nominal Torques are without consideration of friction losses through Bearings or Sealings.

Annotations - thermal monitoring system

Because the exact duty type depends also on the thermal connection of the motor, the embedded thermal monitoring system has to be analysed and attented. However, attention has to be payed that the temperature sensors do not show the exact temperature of the winding and this could be up to 20 K higher due to thermal capacities. Despite an electrical insulation towards the winding, you are only allowed to connect the sensors to your controller by using a galvanic separation in between.

preliminary Data

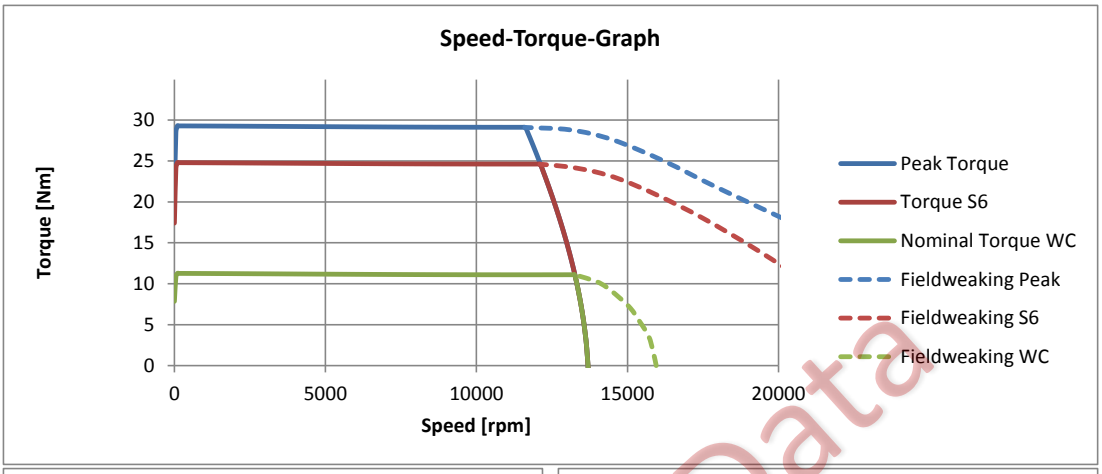
FISCHER Elektromotoren

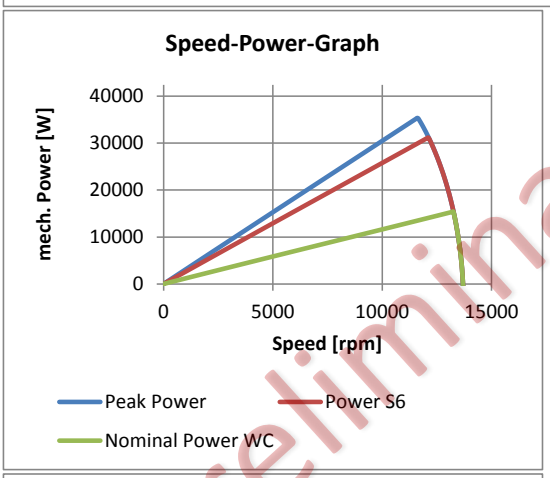
Motor Datasheet [calculated Values]

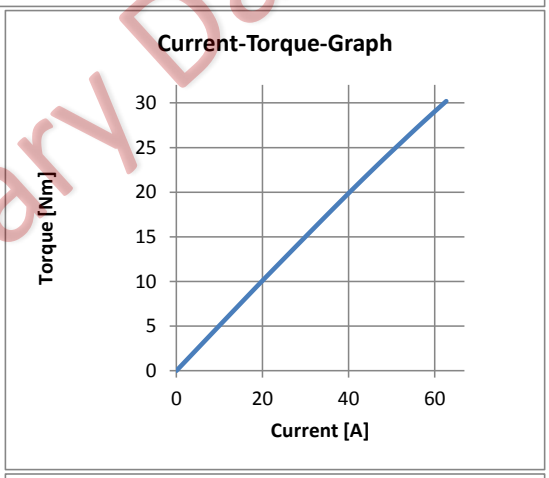
TI085-052-070-04B7S-07S04BE2

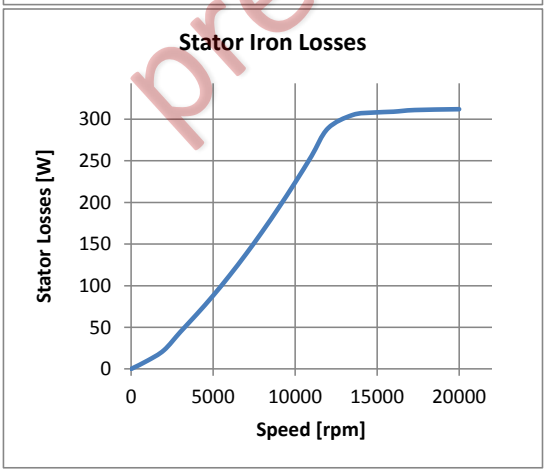
with Fieldweaking

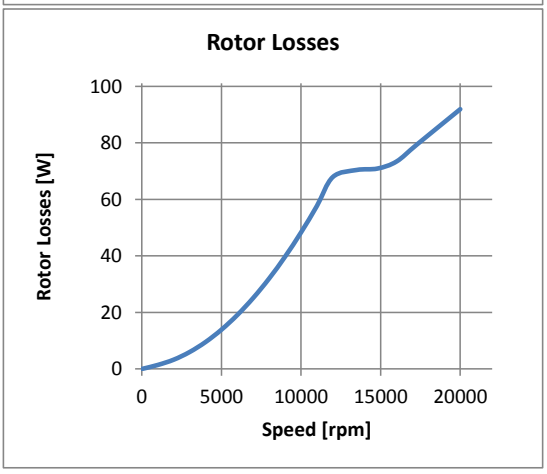
High-Speed & Power-Systems











preliminary Data

Fischer Elektromotoren GmbH
Schützenstraße 19
D-74842 Billigheim-Allfeld
Geschäftsführer: Peter Fischer
Tel. (0049) 6265/9222-0
Fax: (0049) 6265/9222-22
info@fischer-elektromotoren.de
www.fischer-elektromotoren.de

Appendix B

Excerpt TI085-052-070-04B7S-07S04BE2 FEA spreadsheet

TI085-052-070-04B7S-07S04BE2

Ld-Lq-Tabelle (berechnete Daten!)

Is_eff [A]	θ[°el]	Id_eff [A]	Iq_eff [A]	Ld [H]	Lq [H]	FluxM_eff [Vs]	FluxD_eff [Vs]
5	-80	-4.92E+00	8.68E-01	3.17E-04	3.77E-04	4.28E-02	4.13E-02
5	-70	-4.70E+00	1.71E+00	3.06E-04	3.77E-04	4.28E-02	4.14E-02
5	-60	-4.33E+00	2.50E+00	3.06E-04	3.77E-04	4.28E-02	4.15E-02
5	-50	-3.83E+00	3.21E+00	3.07E-04	3.76E-04	4.28E-02	4.16E-02
5	-40	-3.21E+00	3.83E+00	2.99E-04	3.76E-04	4.28E-02	4.18E-02
5	-30	-2.50E+00	4.33E+00	2.97E-04	3.75E-04	4.28E-02	4.21E-02
5	-20	-1.71E+00	4.70E+00	2.94E-04	3.74E-04	4.28E-02	4.23E-02
5	-10	-8.68E-01	4.92E+00	2.80E-04	3.73E-04	4.28E-02	4.25E-02
5	0	0.00E+00	5.00E+00	2.90E-04	3.72E-04	4.28E-02	4.28E-02
11	-80	-1.08E+01	1.91E+00	3.05E-04	3.86E-04	4.28E-02	3.95E-02
11	-70	-1.03E+01	3.76E+00	3.03E-04	3.85E-04	4.28E-02	3.97E-02
11	-60	-9.53E+00	5.50E+00	3.03E-04	3.84E-04	4.28E-02	3.99E-02
11	-50	-8.43E+00	7.07E+00	3.01E-04	3.83E-04	4.28E-02	4.02E-02
11	-40	-7.07E+00	8.43E+00	2.94E-04	3.81E-04	4.27E-02	4.06E-02
11	-30	-5.50E+00	9.53E+00	2.88E-04	3.80E-04	4.27E-02	4.11E-02
11	-20	-3.76E+00	1.03E+01	2.82E-04	3.77E-04	4.27E-02	4.16E-02
11	-10	-1.91E+00	1.08E+01	2.77E-04	3.74E-04	4.27E-02	4.21E-02
11	0	0.00E+00	1.10E+01	2.67E-04	3.72E-04	4.27E-02	4.27E-02
16	-80	-1.58E+01	2.78E+00	3.04E-04	3.93E-04	4.28E-02	3.80E-02
16	-70	-1.50E+01	5.47E+00	3.04E-04	3.93E-04	4.28E-02	3.82E-02
16	-60	-1.39E+01	8.00E+00	3.01E-04	3.91E-04	4.27E-02	3.86E-02
16	-50	-1.23E+01	1.03E+01	2.96E-04	3.90E-04	4.27E-02	3.90E-02
16	-40	-1.03E+01	1.23E+01	2.90E-04	3.87E-04	4.26E-02	3.96E-02
16	-30	-8.00E+00	1.39E+01	2.83E-04	3.84E-04	4.26E-02	4.03E-02
16	-20	-5.47E+00	1.50E+01	2.77E-04	3.80E-04	4.25E-02	4.10E-02
16	-10	-2.78E+00	1.58E+01	2.71E-04	3.76E-04	4.25E-02	4.17E-02
16	0	0.00E+00	1.60E+01	2.61E-04	3.72E-04	4.25E-02	4.25E-02
21	-80	-2.07E+01	3.65E+00	3.07E-04	4.02E-04	4.28E-02	3.65E-02
21	-70	-1.97E+01	7.18E+00	3.05E-04	4.01E-04	4.28E-02	3.67E-02
21	-60	-1.82E+01	1.05E+01	3.00E-04	4.00E-04	4.27E-02	3.72E-02
21	-50	-1.61E+01	1.35E+01	2.94E-04	3.98E-04	4.26E-02	3.78E-02
21	-40	-1.35E+01	1.61E+01	2.88E-04	3.94E-04	4.25E-02	3.86E-02
21	-30	-1.05E+01	1.82E+01	2.80E-04	3.90E-04	4.24E-02	3.95E-02
21	-20	-7.18E+00	1.97E+01	2.71E-04	3.84E-04	4.23E-02	4.04E-02
21	-10	-3.65E+00	2.07E+01	2.57E-04	3.78E-04	4.23E-02	4.13E-02
21	0	0.00E+00	2.10E+01	2.47E-04	3.71E-04	4.22E-02	4.22E-02
26	-80	-2.56E+01	4.51E+00	3.07E-04	4.14E-04	4.28E-02	3.49E-02
26	-70	-2.44E+01	8.89E+00	3.05E-04	4.13E-04	4.27E-02	3.53E-02
26	-60	-2.25E+01	1.30E+01	2.99E-04	4.11E-04	4.26E-02	3.58E-02
26	-50	-1.99E+01	1.67E+01	2.93E-04	4.07E-04	4.25E-02	3.66E-02
26	-40	-1.67E+01	1.99E+01	2.85E-04	4.02E-04	4.23E-02	3.75E-02
26	-30	-1.30E+01	2.25E+01	2.74E-04	3.95E-04	4.22E-02	3.86E-02
26	-20	-8.89E+00	2.44E+01	2.64E-04	3.87E-04	4.20E-02	3.97E-02
26	-10	-4.51E+00	2.56E+01	2.55E-04	3.79E-04	4.20E-02	4.08E-02
26	0	0.00E+00	2.60E+01	2.44E-04	3.71E-04	4.20E-02	4.20E-02
31	-80	-3.05E+01	5.38E+00	3.14E-04	4.24E-04	4.28E-02	3.32E-02
31	-70	-2.91E+01	1.06E+01	3.09E-04	4.24E-04	4.27E-02	3.37E-02