POLITECNICO DI TORINO

Corso di Laurea
in Mechatronics Engineering

Tesi di Laurea

Advanced Electric Drive Control Systems for Silent Electric Powertrains

Relatori
prof. Paolo Pescetto

Candidato
Giovanni De Martino

Anno Accademico 2025-2026

Sommario

Objectives

This thesis addresses the gap between simulation models and real hardware behavior in permanent magnet synchronous motors for electric vehicle applications. The main objectives were: developing a reliable test-bench infrastructure for motor characterization, implementing stable field-oriented control algorithms, and building a workflow to update simulation models from FEA using experimental data. The goal was to provide feedback from real testing to improve both motor design and control strategies.

Methodology

The test-bench (Figure 1) consisted of a surface-mounted PMSM (800W, 8.2Nm, 5 pole pairs) coupled to a loading machine, driven by a dSPACE MicroLabBox controller and a custom six-leg inverter.

Critical electromagnetic interference (EMI) issues corrupting rotor position measurements were resolved by installing ferrite cores on encoder cables, separating power supplies for control and power electronics, and implementing proper shielding. A rigorous alignment procedure was developed using both open-loop voltage control and closed-loop current control in the stator reference frame, with the latter proving to be more robust. A braking chopper was implemented by exploiting the fourth inverter leg to allow for four-quadrant operation and protect the DC supply during regenerative braking. Current fuses were added for thermal protection.

Field-oriented control (FOC) was implemented in Simulink with two PI current controllers in the dq rotating frame, converting the current errors into the corresponding voltage references V_d and V_q. PI parameters were initially calculated theoretically, then fine-tuned experimentally with the add-on of anti-windup mechanisms. Automated workflows for systematic testing were created, and CAN communication between controller and loading machine was established for both velocity and torque control.

A systematic procedure was developed to characterize the motor across its operating range. Permanent magnet flux linkage was estimated through no-load run-up tests measuring back-EMF. Torque maps were generated by testing multiple current operating points (Id, Iq) with direct torque measurement. Initial flux map attempts using steady-state voltage equations failed due to temperature rise in the stator resistance. The Motor-Generator-Motor (MGM) technique was then implemented, which eliminates resistance dependency by exploiting flux symmetry towards the direct axis and assuming linear resistance increase over time during a single current setpoint. This approach proved significantly more reliable under varying temperature conditions.

Results

The test-bench improvements successfully eliminated EMI, enabling stable rotor position measurements and precise field-oriented control. Current control achieved fast transient response (few milliseconds settling time) with zero steady-state error for both d and q axes, demonstrating the expected symmetry of a low-saliency motor.

Comprehensive torque and flux maps (Figure 2) were generated in the admissible current range. A study on limit conditions was also performed, translating the loading machine's torque limit into the machine under test's current bound. Permanent magnet flux linkage was estimated at approximately 9.2mWb. The updated Amesim model incorporated these experimental characteristics, significantly improving simulation accuracy compared to the original FEA-based model using approximate manufacturer's information. MTPA validation confirmed the optimal operating trajectory lies close to the Iq axis, as expected for low-saliency motors such as the machine under test.

Individual Contribution

EMI issues were independently diagnosed and resolved through systematic troubleshooting and testing of different mitigation strategies. The braking chopper circuit and control logic were designed and implemented, along with several additional software and hardware safety measures. The complete Simulink control architecture was developed, including current loops, coordinate transformations, space vector modulation, and alignment procedures. The PI tuning process was conducted from theoretical calculations to experimental trial and error. The loading machine velocity control was integrated through CANopen, allowing the complete test-bench to be operated from the same real-time environment.

The most significant contribution was recognizing temperature effects in flux estimation and developing the Motor-Generator-Motor implementation. The experimental protocol was designed, measurement and data processing algorithms were implemented, and results were validated. Automated testing workflows were created that systematically swept the motor's operating range and generated maps for model updating. All data analysis, including FFT harmonic evaluation and filter design, was performed independently.

Keywords: PMSM characterization, Field-Oriented Control, Test-bench development, Model updating, EMI mitigation, Motor-Generator-Motor technique, Amesim simulation, Hardware-in-the-loop

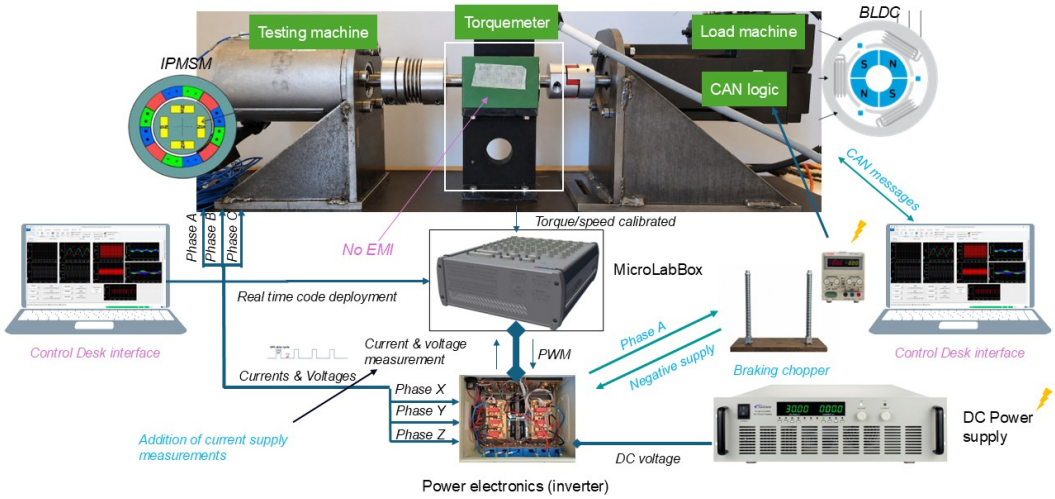


Figura 1: Final test-bench setup developed

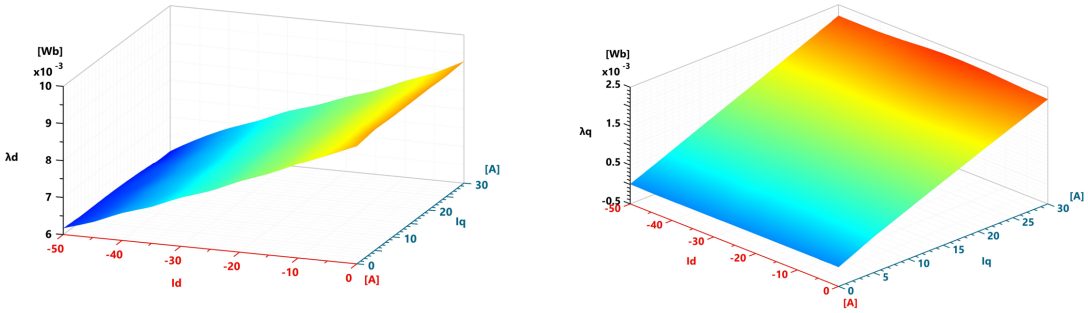


Figura 2: Flux maps obtained through Motor-Generator-Motor technique: (left) d-axis flux linkage, (right) q-axis flux linkage

Indice

Elenco delle tabelle 7

Elenco delle figure 8

1 Introduction 11
- 1.1 Motivation: electric drives in modern powertrains . . . 11
- 1.2 AC motors as electrical machines . . . 12
- 1.3 Permanent Magnet Synchronous Motors . . . 14
- 1.4 Thesis objectives and structure . . . 17

2 A walk through the open-loop control 19
- 2.1 Overview of the initial test-bench . . . 19
- 2.2 DC power supply . . . 20
- 2.3 Inverter . . . 22
  - 2.3.1 Initial stage of the Drivers . . . 23
  - 2.3.2 Main driver . . . 24
  - 2.3.3 Current measurements . . . 25
- 2.4 Machine under test . . . 28
- 2.5 Loading machine . . . 31
- 2.6 First loading test . . . 33
- 2.7 Phase determination procedure . . . 37
- 2.8 Open-loop control . . . 39
  - 2.8.1 Open loop control initial state-of-art . . . 39
  - 2.8.2 EMC EMI . . . 42
  - 2.8.3 Open-loop control after shielding . . . 46

3 Current to torque control 49
- 3.1 Current control and parameter tuning . . . 49
- 3.2 DC loading motor and CAN protocol . . . 56
  - 3.2.1 Brief introduction to CAN and CANopen . . . 57
  - 3.2.2 Objects used for the loading machine . . . 58
  - 3.2.3 Implementation in Simulink and Stateflow . . . 59
- 3.3 Braking chopper . . . 62
- 3.4 Experimental torque and flux maps . . . 67

3.4.1 Aim of the characterization 67
3.4.2 Experimental workflow 67
3.4.3 Torque map measurement 68
3.4.4 Initial flux map estimation from voltage equations 69
3.4.5 Problem of the stator resistance 70
3.4.6 DC flux maps and limits of harmonic analysis 72
3.4.7 Construction of the maps 72
3.4.8 Use of the maps for model updating 74

A Main Inverter Driver Schematic 75

Elenco delle tabelle

3.1 Main CANopen objects used for the control of the PD6 loading machine. . 59

Elenco delle figure

1 Final test-bench setup developed . . . 4
2 Flux maps obtained through Motor-Generator-Motor technique: (left) d-axis flux linkage, (right) q-axis flux linkage . . . 4
1.1 General structure of an electric drive system. . . . 11
1.2 Transformer working principle. . . . 13
1.3 Generic electric motor. . . . 13
1.4 Simplified structure of a Permanent Magnet Synchronous Motor. . . . 15
1.5 Park transformation and rotating reference frame. . . . 16
1.6 MTPA trajectory of a surface-mounted PMSM. . . . 17
2.1 Schematic representation of the initial test-bench and data exchange . . . 19
2.2 Twintex PCL9000-60 . . . 20
2.3 DC supply with no four-quadrant operation connected to an inductive load. 21
2.4 Distorted DC voltage supply from Twintex . . . 21
2.5 PeakTech 1580 DC supply . . . 22
2.6 Internal view of the inverter . . . 22
2.7 Front view of the inverter. . . . 23
2.8 Rear view of the inverter . . . 23
2.9 Initial stage of the drivers' scheme . . . 23
2.10 Main driver . . . 24
2.11 Main driver's schematic . . . 24
2.12 Distorted signal at the optocoupler input . . . 24
2.13 PCB after resistor integration . . . 25
2.14 Inverter output voltage when asked 30V, 10%DC . . . 26
2.15 Inverter output voltage when asked 30V, 50%DC . . . 26
2.16 Inverter output voltage when asked 30V, 90%DC . . . 26
2.17 Inverter output voltage when asked 30V, 100%DC . . . 26
2.18 Generation of a PWM signal with frequency 100Hz . . . 27
2.19 Generation of a PWM signal with frequency 1kHz . . . 27
2.20 Generation of a PWM signal with frequency 10kHz . . . 27
2.21 Generation of a PWM signal with frequency 100kHz . . . 27
2.22 Hall sensors' appearance . . . 27
2.23 Hall-effect current measurement principle . . . 27
2.24 Measurement circuit PCB . . . 28
2.25 Current measurement calibration setup . . . 28

2.26 Ideal measured current over reconstructed current . . . 29
2.27 Current measurement system characteristics including calibration . . . 29
2.28 Machine under test . . . 29
2.29 Machine under test's rotor . . . 29
2.30 $\lambda_d(I_d, I_q)$ . . . 30
2.31 $\lambda_d(I_d)$ . . . 30
2.32 $\lambda_d(I_q)$ . . . 30
2.33 $\lambda_q(I_d, I_q)$ . . . 30
2.34 $\lambda_q(I_d)$ . . . 30
2.35 $\lambda_q(I_q)$ . . . 30
2.36 $torque(I_d, I_q)$ . . . 31
2.37 $torque(I_d)$ . . . 31
2.38 $torque(I_q)$ . . . 31
2.39 PD6-EB80CD-E-65-2 . . . 31
2.40 Former configuration . . . 32
2.41 Updated configuration . . . 32
2.42 First run-up Amesim model . . . 33
2.43 Line-to-line voltage simulation at 150rpm . . . 34
2.44 Line-to-line voltage simulation at 900rpm . . . 34
2.45 Tektronix P5200A voltage probe . . . 34
2.46 Line-to-line voltage measurement with and without probes . . . 35
2.47 Line-to-line voltage measurement with and without probes close up . . . 35
2.48 Line-to-line voltage measurement with and without probes at 900rpm . . . 35
2.49 All three line-to-line voltage measurements with probes and with a ramp velocity . . . 35
2.50 FFT applied to one line-to-line voltage, close up to the main electrical frequency . . . 36
2.51 FFT comparison of the line-to-line voltage with and without probes . . . 36
2.52 FFT comparison after filtering the signal measured with probes . . . 36
2.53 Electrical odd harmonics at 150rpm . . . 37
2.54 Electrical odd harmonics at 300rpm . . . 37
2.55 Electrical odd harmonics at 600rpm . . . 37
2.56 Electrical odd harmonics at 900rpm . . . 37
2.57 Electrical harmonics analysis while swapping the probes . . . 38
2.58 Mechanical angle during one electrical rotation . . . 38
2.59 Voltage open-loop control on Simulink . . . 39
2.60 SPM modeling on Amesim . . . 39
2.61 Predicted torque . . . 40
2.62 Predicted velocity . . . 40
2.63 Vq command during the failed open-loop test . . . 41
2.64 EMC EMI effect on velocity . . . 41
2.65 Rotor position during a failed open-loop cycle leading to reverse rotation . . . 42
2.66 EMC EMI possible sources . . . 42
2.67 Single-point ground scheme . . . 43
2.68 Multi-point ground scheme . . . 43

2.69 Ferrite core addition . . . 45
2.70 Shielded cable . . . 46
2.71 V_q command . . . 47
2.72 Resultant speed of the rotor . . . 47
3.1 Generic current control . . . 50
3.2 Bode plot used for the first PI tuning . . . 52
3.3 P only control when asked $I_d = 10A$ . . . 52
3.4 P only control when asked $I_q = 10A$ . . . 52
3.5 I only control when asked $I_d = 10A$ . . . 53
3.6 I only control when asked $I_q = 10A$ . . . 53
3.7 PI control with P part slowed down by 100 times . . . 53
3.8 PI control with nominal values . . . 53
3.9 PI control oscillations close up at steady state when increasing k_p by one hundred times . . . 54
3.10 PI control with too large proportional gain . . . 54
3.11 PI control instability with too large proportional gain . . . 54
3.12 PI control with I part slowed down by 100 times . . . 55
3.13 PI control with nominal values . . . 55
3.14 PI control with a faster integral part . . . 55
3.15 Steady state current control when speed is increased . . . 56
3.16 Current control at steady state high speed . . . 56
3.17 Current control with combined inputs of (a) I_d, (b) I_q, and (c) speed. . . . 56
3.18 Overview of the implemented CANopen sequence. The upper Stateflow chart manages the communication with the PD6 loading machine through CANopen, while the lower chart controls the electrical states of the Machine Under Test. The two state machines are synchronized so that the loading machine starts rotating only after the alignment and initialization of the current controller have been completed. . . . 62
3.19 Motoring and generating modes . . . 63
3.20 Braking chopper integration in the final test-bench configuration . . . 65
3.21 Experimental DC torque map obtained from the torquemeter measurement. 69
3.22 Estimated DC flux linkage maps: (a) d-axis flux linkage over I_d, (b) d-axis flux linkage overview, (c) q-axis flux linkage over I_q, and (d) q-axis flux linkage overview. . . . 74
A.1 Complete schematic of the main inverter driver. . . . 76

Capitolo 1

Introduction

1.1 Motivation: electric drives in modern powertrains

Hybrid and electric vehicles rely on electric drives to produce motion. In this kind of application the motor is not working alone, but it is part of a complete chain made by the DC supply, the inverter, the control unit, the sensors as well as the mechanical load. If one of these elements does not behave correctly, the whole drive is affected.

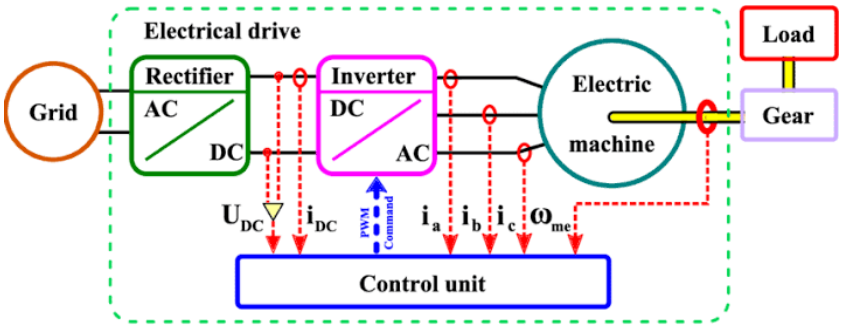


Figura 1.1: General structure of an electric drive system.

This thesis starts from this practical problem. The test bench was already present, but it was not ready to be used directly for motor characterization (update) and torque control. Before doing any advanced control, we had to first understand how the setup was behaving, which measurements were reliable, and which parts of the bench were creating problems as we will see later on.

In simulation, the motor and the inverter can be described in an easy and straightforward way. In a real test bench, instead, several non-idealities appear. The rotor position measurement can be corrupted by electromagnetic interference, the current sensors need calibration, the inverter does not behave as an ideal voltage source, and the motor parameters are not exactly the ones expected from the initial model. These aspects are

not secondary, because current control and torque control depend directly on how these measurements are.

Among the different electrical machines used in automotive applications, Permanent Magnet Synchronous Motors are largely used because they offer high efficiency, compact size and high torque density. The machine considered in this work is a surface-mounted PMSM or simply SPM. This kind of machine is convenient from a control point of view, because the torque is mainly linked to the quadrature current. However, this is true only if the rotor position and the phase convention are known.

For this reason, the work was not only about writing a control algorithm but an extensive part of the activity was about making the test bench usable: checking the power supply, testing the inverter, understanding the phase order, solving encoder issues, decreasing EMIs, and only after that moving to current control. The final goal was to have a setup that could be used to collect reliable experimental data and update the motor model.

1.2 AC motors as electrical machines

An electrical machine is a device that exchanges energy through electromagnetic interaction. In this thesis we are mainly interested in rotating electrical machines, but it is useful to briefly place them inside the wider electrical machine family.

The whole electrical machine world can be divided into three main groups:

- Transformers: they exchange electrical energy between two circuits without moving parts. A primary winding creates a varying magnetic flux in the core, and this flux induces a voltage in the secondary winding. Since the working principle is based on a time-varying magnetic field, transformers are naturally associated with AC quantities.

  In ideal conditions, the power balance is:

  $$P_1 = V_1 I_1 = V_2 I_2 = P_2, \qquad V_2 > V_1 \Rightarrow I_2 < I_1$$

  The voltage ratio depends on the ratio between the number of turns of the two windings.

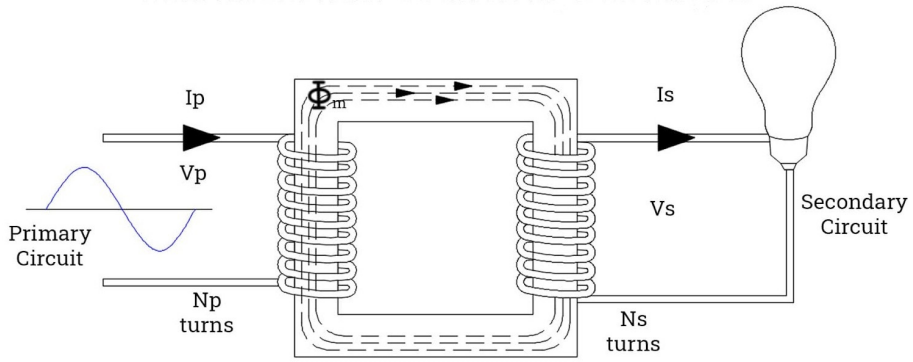


Figura 1.2: Transformer working principle.

- Motors: they convert electrical energy into mechanical energy. This is the case that mainly interests us. On the electrical side, voltage and current are supplied to the machine. On the mechanical side, the useful quantities are torque and angular speed. In a simplified way:

  $$P_1 = VI = T\omega = P_2$$

  and torque is often related to current through:

  $$T = k_t I$$

  where k_t is the torque constant.



Figura 1.3: Generic electric motor.

- Generators: they convert mechanical energy into electrical energy. The same machine can often work both as a motor and as a generator, depending on the sign of

torque and speed. This point is important for our test bench, because during some tests the machine can enter regenerative operation.

In this document, motors and generators will sometimes be treated together as terminology, since the same electrical machine can exchange energy in both directions. This was also relevant for the braking chopper activity, because the DC supply could not withstand regenerated energy.

The main categorization for rotating electrical machines is between DC machines and AC machines.

DC machines are supplied by direct current. The most classical example is the brushed DC machine. It is composed of a stator field winding, which creates the excitation field, and a rotor armature winding supplied through brushes and a mechanical commutator. The field current creates the magnetic field, while the armature current is mainly responsible for torque production.

The torque can be written as:

$$T_m = k_t I_a I_f$$

where I_a is the armature current, I_f is the field current, and k_t is the torque constant. If the field current is kept constant, the torque becomes mainly proportional to the armature current. This is the reason why brushed DC machines are simple to understand and control.

However, brushes and mechanical commutation bring limitations. They introduce wear, maintenance, friction, electrical noise and speed limits. For compact and efficient electric drives, these drawbacks are not ideal.

Brushless solutions remove the mechanical commutator and use electronic commutation instead. In this case the rotor does not need to be supplied through brushes, and the stator currents are generated by a power electronic converter. This also makes the distinction between DC and AC less immediate: the drive can start from a DC link, but the machine itself receives phase currents that vary in time.

AC machines are supplied by alternating currents that are phase shifted. In rotating AC machines, the stator currents generate a rotating magnetic field, and torque is produced by the interaction between this field and the rotor. In modern drives, the AC supply is usually generated by an inverter starting from a DC voltage source. For this reason, the inverter is not just an addition but a fundamental part of the drive [Krishnan, 2010].

This is exactly the situation of the test bench considered in this thesis. The motor behaviour cannot be separated from the inverter, the sensors and the control unit. The following section therefore focuses on the PMSM model and on the quantities needed for Field-Oriented Control.

1.3 Permanent Magnet Synchronous Motors

Permanent Magnet Synchronous Motors belong to the family of AC machines. The stator is supplied by three-phase currents, while the rotor contains permanent magnets. These magnets create the rotor magnetic field without the need for an external excitation winding.

The rotor follows the rotating stator field at synchronous speed. Therefore, the electrical angle used by the controller is linked to the mechanical rotor position through:

$$\theta_e = p\theta_m$$

where θ_e is the electrical angle, θ_m is the mechanical angle, and p is the number of pole pairs. The same relation is valid for the angular speed:

$$\omega_e = p\omega_m$$

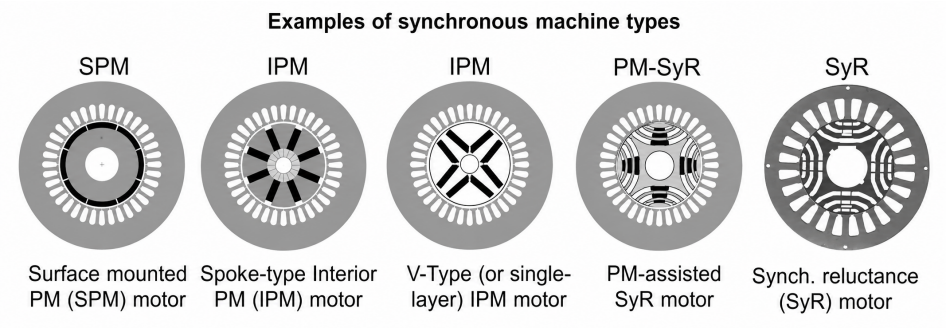


Figura 1.4: Simplified structure of a Permanent Magnet Synchronous Motor.

This relation is not only theoretical. In the test-bench, the rotor position is one of the most important measured quantities. If the measured position is incorrect, the controller applies the coordinate transformation using a wrong electrical angle. Consequently, even if the numerical command appears correct, the physical voltage or current applied to the machine is not the intended one.

To control the PMSM, the three-phase quantities are transformed into the rotating dq reference frame. The d-axis is aligned with the rotor magnetic flux, while the q-axis is orthogonal to it. This transformation is known as the Park transformation [Krishnan, 2010, Sul, 2011].

In this work, the power-invariant form of the Park transformation is considered. For a generic three-phase quantity x_{abc}, it can be written as:

$$\begin{bmatrix} x_d \\ x_q \\ x_0 \end{bmatrix} = \sqrt{\frac{2}{3}} \begin{bmatrix} \cos\theta_e & \cos\left(\theta_e - \frac{2\pi}{3}\right) & \cos\left(\theta_e + \frac{2\pi}{3}\right) \\ -\sin\theta_e & -\sin\left(\theta_e - \frac{2\pi}{3}\right) & -\sin\left(\theta_e + \frac{2\pi}{3}\right) \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} x_a \\ x_b \\ x_c \end{bmatrix}.$$

The zero-sequence component x_0 is zero for an ideally balanced three-phase system. Therefore, the quantities of main interest for the motor control are x_d and x_q.

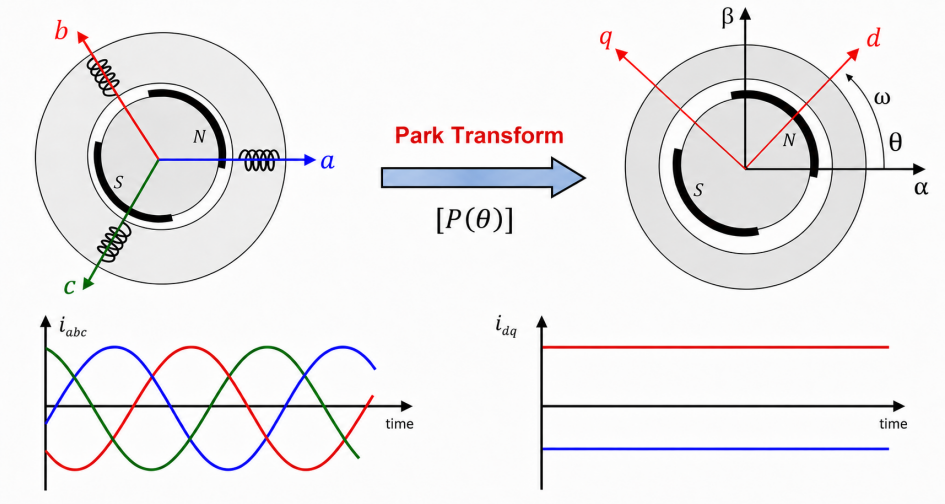


Figura 1.5: Park transformation and rotating reference frame.

In the rotating dq reference frame, the flux linkages of a PMSM can be written as:

$$\lambda_d = \lambda_{pm} + L_d I_d$$

$$\lambda_q = L_q I_q$$

where λ_{pm} is the permanent-magnet flux linkage, L_d and L_q are the direct- and quadrature-axis inductances, and I_d and I_q are the corresponding current components.

With the power-invariant transformation adopted above, the electromagnetic torque is expressed as:

$$T_m = p\left(\lambda_d I_q - \lambda_q I_d\right).$$

By substituting the flux-linkage expressions:

$$T_m = p\left[\lambda_{pm} I_q + (L_d - L_q) I_d I_q\right].$$

This expression shows that the electromagnetic torque has two main contributions. The first one originates from the interaction between the permanent-magnet flux and the quadrature current. The second one depends on the difference between L_d and L_q and represents the reluctance torque.

In a surface-mounted PMSM, the magnets are placed on the rotor surface. The rotor therefore presents relatively low magnetic saliency, and it is generally possible to assume:

$$L_d \simeq L_q.$$

As a consequence, the reluctance contribution becomes small and the torque mainly depends on the quadrature current:

$$T_m \simeq p\lambda_{pm} I_q.$$

This behaviour is particularly relevant for the machine under test. Since it behaves as a low-saliency machine, most of the electromagnetic torque is expected to be produced through I_q. This simplification was useful during the first control implementation, since the quadrature current could be considered the main torque-producing current.

The direct-axis current I_d is instead mainly associated with the magnetic flux condition. Therefore, for a surface-mounted PMSM operating below the voltage limit, a common strategy is to keep I_d close to zero and regulate I_q according to the required torque.

This is also one of the reasons why Field-Oriented Control is particularly suitable for PMSM drives. By expressing the electrical quantities in the rotating dq frame, the controller can regulate the direct and quadrature current components independently, separating, at first approximation, the flux-related and torque-producing effects[Sul, 2011].

The initial FEA maps of the machine under test show the expected low-saliency behaviour, with a relatively small influence of I_d on torque and an MTPA trajectory located close to the I_q-axis.

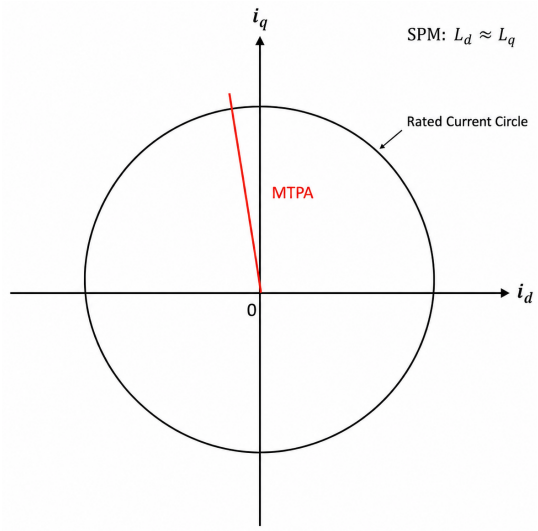


Figura 1.6: MTPA trajectory of a surface-mounted PMSM.

1.4 Thesis objectives and structure

The work described in this thesis was performed on an already existing electric motor test bench. The objective was not to build a theoretical PMSM controller only, but to make the real setup work correctly enough to perform control and characterization tests.

The work followed a practical order. First, the main hardware components had to be understood: the DC supply, inverter, MicroLabBox, current sensors, encoder, machine

under test and loading machine. Several problems appeared, especially related to inverter output, current measurements, phase convention and electromagnetic interference.

After that, open-loop voltage control was implemented. This was useful to check if the inverter commands, the phase order and the rotor alignment were consistent. It also showed how dangerous a wrong position measurement can be, because the controller could reconstruct apparently correct physical quantities while the motor receives the wrong physical command.

Once the setup became more reliable, closed-loop current control was implemented in the *dq* frame. PI controllers were first computed from theoretical values and then tuned experimentally. This step was necessary because the real machine and inverter did not behave exactly like the simplified model.

The final objective was to move towards torque control and motor characterization. The available motor data were not accurate enough for a reliable torque model, so experimental maps were needed to describe the real machine better.

Chapter 2 presents the test bench, the main hardware components, the first tests, the phase determination procedure and the open-loop control activity. Chapter 3 focuses on closed-loop current control, PI tuning, loading machine communication, torque measurement and the steps needed to move toward torque control and motor characterization.

Capitolo 2

A walk through the open-loop control

2.1 Overview of the initial test-bench

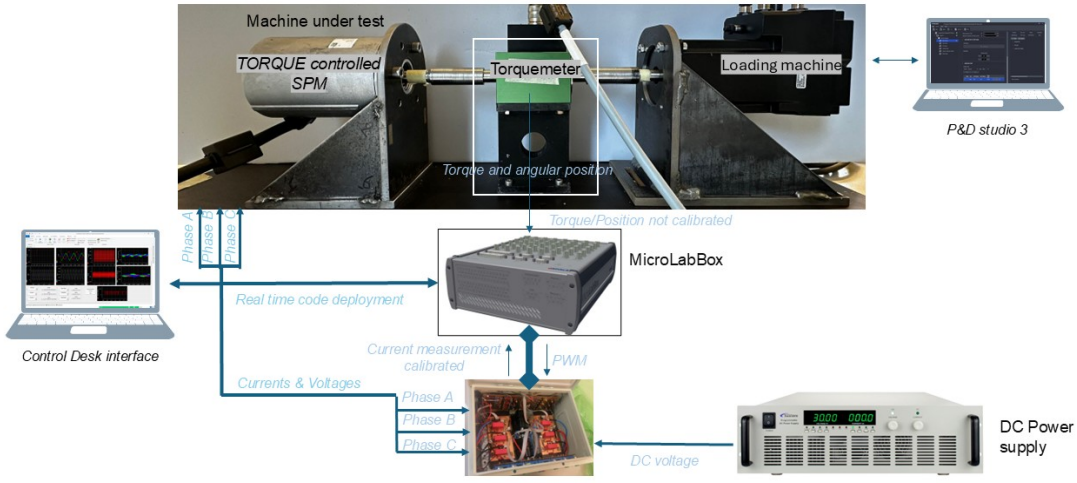


Figura 2.1: Schematic representation of the initial test-bench and data exchange

Figure 2.1 shows the initial state of the test-bench considered in this work. The machine under test and the loading machine are both mounted on metal supports bolted on a wooden plate. Their shafts are connected through a torquemeter, which also provides the angular displacement measurement [burster]. Two softwares are used in this first configuration. P&D Studio 3 is used to command the loading machine through a USB-CAN cable connected directly to the laptop. ControlDesk is instead used as interface for the inputs and outputs of the dSPACE MicroLabBox DS1202. This connection is physically realized through Ethernet. The control logic is developed in Simulink with a

model-based approach, then the corresponding C code is generated and transferred to the MicroLabBox to be executed in real time [dSPACE, 2023]. Once the code is running, the control stage interacts with the power electronics, represented by a custom-made six-leg inverter for medium-power applications. The inverter is supplied by a high-power, high-current medium-voltage Twintex DC supply. From the MicroLabBox it receives three non-inverted PWM commands; the complementary signals are then generated internally by the driver stage, before the power stage produces the corresponding square-wave voltages. On the machine-under-test side, the three motor phases receive the voltages generated by the inverter, as in a standard PMSM drive. The reaction of the machine is both electrical and mechanical. On the electrical side, the imposed voltages produce phase currents that are measured and sent back to the control system. On the mechanical side, the loading machine can impose either constant velocity or constant torque, depending on the test to be performed.

2.2 DC power supply

The power supply is the component through which the whole power stage is fed. It provides a constant DC voltage to the inverter, therefore creating the DC link. The inverter arms then switch this voltage and generate the phase voltages applied to the machine under test. Depending on the reference point, the inverter output can be seen as switching between approximately $\pm V_{dc}/2$, or between 0 V and V_{dc} with respect to the negative DC terminal.

In terms of current, the supply had to match the requirements of the planned tests. The test-bench was intended to work with a DC link voltage up to 60 V and phase currents up to 50 Arms for at least five minutes. For this reason, the selected supply was the Twintex PCL9000-60 [Twintex, 2022].

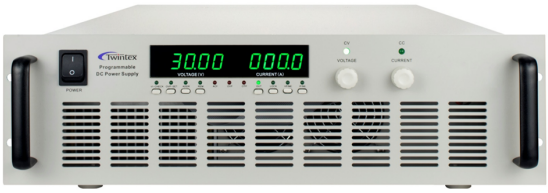


Figura 2.2: Twintex PCL9000-60

The main specifications are:

- 0~60 V output voltage range, matching our 30~60 V working range, with 0.01 V resolution.
- 0~150 A output current range, much larger than what was really needed, ensuring this would not represent a bottleneck in our control. Current resolution is 0.1 A.
- Rated power equal to 9 kW, meaning that voltage and current can be exploited in their full range simultaneously.
- Tunable over-voltage protection up to 63 V.
- Voltage ripple $\leq 100 mV_{pp}$.

- No four-quadrant operation support.

A simple schematic of the connection between the power source and the three-phase motor, represented as a generic inductive load, is reported below.

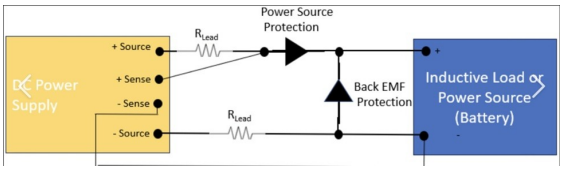


Figura 2.3: DC supply with no four-quadrant operation connected to an inductive load.

A form of power source protection was then implemented later, to avoid accidentally injecting current back into the supply and damaging it. This is a common situation when working with electric motors, since they can enter generator mode when torque and speed have opposite signs.

The voltage ripple specification was also tested, and the result is reported in the following figure.

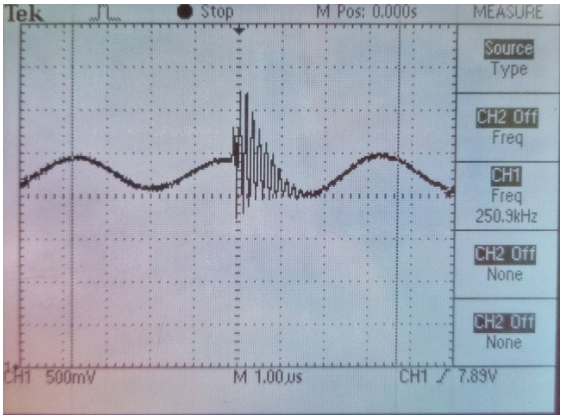


Figura 2.4: Distorted DC voltage supply from Twintex

During this test, a simple DC output of 8 V was requested to the source. According to the nominal specifications, the supply was expected to have no trouble performing such a basic task, since the voltage ripple lower than $100mV_{pp}$ was referred to the rated voltage. The outcome was instead the superimposition of the desired DC signal with a ripple of about 0.5 Vpp at $\simeq 250kHz$, plus another signal at much higher frequency. No spectral analysis was performed on this last contribution due to its high-frequency nature.

This fault in the DC voltage generation was believed to be the cause of the EMC/EMI problem on the encoder measurement, or at least one of the possible causes. For this reason, the test-bench was also equipped with a second DC supply, the PeakTech 1580.

This second DC supply was much smaller than the Twintex one, but apparently it improved the EMC/EMI problem due to the higher stability of its signal. It was therefore a very suitable additional supply for the first testing phase, where the two machines had



Figura 2.5: PeakTech 1580 DC supply

no common DC link. The front output is limited to 5 A, while the higher-current output is available from the rear terminals [PeakTech, 2024].

2.3 Inverter

Inverter's features:

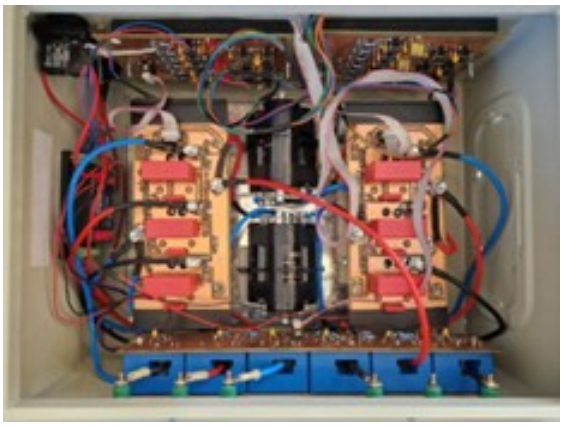


Figura 2.6: Internal view of the inverter

- Nominal DC link supply voltage: 30V-60V
- 2 three-phase converters
- Rated phase current: 30Arms
- Internal DC supply for the drivers and sensors
- Internal deadtime set to 700ns
- Input: 3+3 PWM signals, referenced to ground
- Two separate DC links, with common negative reference in the considered configuration
- Cooling is performed via continuously operating fans

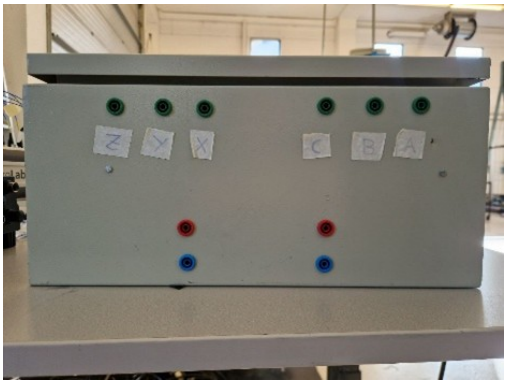


Figura 2.7: Front view of the inverter.



Figura 2.8: Rear view of the inverter

Such an inverter is capable of supplying two different three-phase motors at the same time, with no additional cables needed to make them share the negative reference. Our test-bench was in fact configured to eventually test two engaged models of the same three-phase motor, and the power features of the inverter are suitable for such an application.

The rear part is equipped with an ON/OFF button, the power supply cable, both not included in the picture, and a 25-pin port as well. The aim of the power supply is to power all the drivers and internal circuitry, but also to compensate for the losses. The I/O port accepts digital inputs for MOSFET control from the MicroLabBox, and analog outputs coming from the Hall sensors [Nemeş, 2021, Holmes and Lipo, 2003, Iannuzzo, 2025].

2.3.1 Initial stage of the Drivers

In an inverter, drivers have the duty to process the signal coming from the PWM generator referred to the high-side only, propagate it to command the high-side MOSFET and negate it for the low-side one.



Figura 2.9: Initial stage of the drivers' scheme

For obvious reasons, two different drivers are needed to supply the two different branches of the same inverter. Drivers' specifications:

- optically separated from the control board
- include 700ns deadtime between upper and lower ON/OFF signals
- require 3 input PWM signals, one for each phase
- automatically generate the negated signal
- all 3 input signals must be either TOP or BOTTOM from the controller output

- do not allow deadtime variation or modification
- no over-current/voltage protection is embedded

2.3.2 Main driver

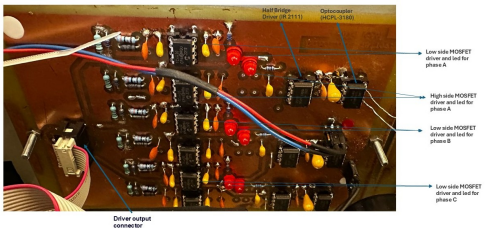


Figura 2.10: Main driver

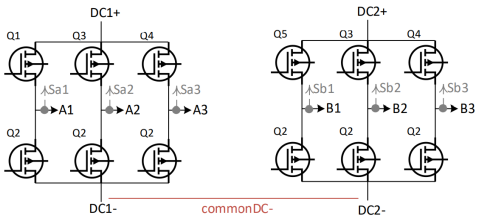


Figura 2.11: Main driver's schematic

The main driver's full scheme with all its components and connections can be found in Appendix A. It is composed of:

- 3 optocouplers (HCPL-3180)
- 3 half-bridge inverter circuits (IR2111)
- 6 MOSFET drivers (TC4420)
- 6 power MOSFETs

At first, the inverter did not produce the commanded output regardless of the supplied voltage, in the allowed range, or the PWM period. It appeared that no matter what the command was, the output was close to a disrupted square wave with $DC \approx 90\%$. After checking the correctness of the digital signal coming out of the MLBX and the signal coming out of the driver, it was clear the error had to reside in the optocouplers.

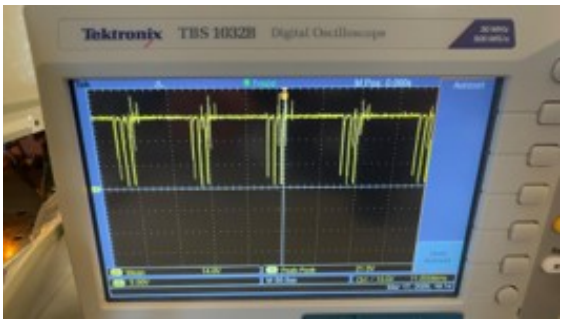


Figura 2.12: Distorted signal at the optocoupler input

It was observed that the input of the optocoupler did not have current limiting on the LED side. The input current exceeded the permissible forward current (I_F) for the LED. The current was higher than what the LED could handle, which caused the MicroLabBox to enter a current limitation mode. This was the reason for the shape of the output signal waveform. The solution we opted for was the addition of a series resistor to limit the

input current to the LED of the optocoupler, ensuring it does not exceed the maximum forward current rating.

Given the parameters:

- $I_{F(min)} = 0.01A$
- $I_{F(max)} = 0.016A$
- $V_{IN} = 5V$
- $V_F = 1.5V$ typical

For 5V input we get a minimum resistor value of:

$$R_{min} = \frac{V_{IN} - V_F}{I_{F(max)}}$$

With the same input voltage we get a maximum resistor value of:

$$R_{max} = \frac{V_{IN} - V_F}{I_{F(min)}}$$



Figura 2.13: PCB after resistor integration

After mounting the current limiting resistors, normal behavior of each leg was achieved.

The pictures above show the square waves generated by the inverter at 100ms PWM period with a supply of 15V, outside the nominal range. Notice that it works properly until $DC = 99\%$, while at $DC = 100\%$ the signal decreases by roughly 10%. For this reason, we always included saturation at 99% duty cycle in our algorithms. Acceptable overshoots are also present.

We also reported the square waves generated for different periods. The signal is still regular at 1ms, while 10ms or 10μs periods bring about the degradation of the signal.

2.3.3 Current measurements

All six inverter legs benefit from current measurements. Such measurements were realized through LA 125-P Hall sensors paired with an integrated PCB circuit to have a voltage in the range of ~10V [LEM, 2013].

Hall sensors' specifications and features:

- Primary nominal RMS current: 125A
- Measuring range: 0 to $\pm 200A$
- Measurement resistor approximately 33Ohm
- Accuracy at max I_{PN}, at $\pm 15V$ supplied equal to $\pm 0.6\%$

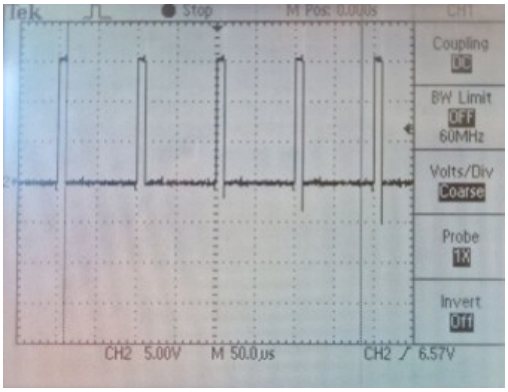


Figura 2.14: Inverter output voltage when asked 30V, 10%DC

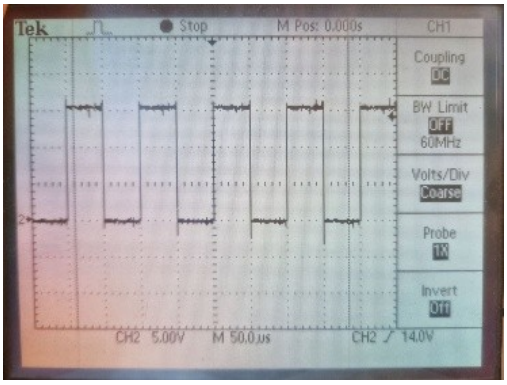


Figura 2.15: Inverter output voltage when asked 30V, 50%DC

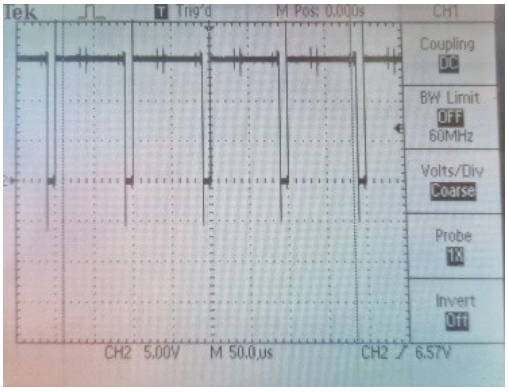


Figura 2.16: Inverter output voltage when asked 30V, 90%DC

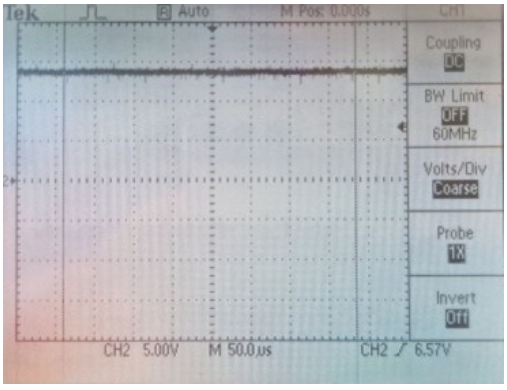


Figura 2.17: Inverter output voltage when asked 30V, 100%DC

- Linearity error $< 0.15\%$
- Closed loop compensated current transducer using the Hall effect
- Insulating plastic case recognized according to UL 94-V0
- Non-disruptive measurement

These inaccuracies were later taken into account during the calibration phase.

Since the PCB circuit outputs an analog voltage, while we were interested in tracing back the phase currents, we proceeded to compute the transresistance from the primary current to such output. This way, we would directly have the current measurement in Ampere. This simply corresponds to computing the gain and, in case, the offset to apply to our measurement system.

The test-bench for the static calibration is reported below.

One inverter leg and the negative point are connected to a series resistor of 2.2Ohm each. Current measurements of the Hall sensor are then compared to the one of the clamp meter model Extech EX655 in DC mode and after proper demagnetization.

Since the resistors were just stuck on a plastic table, which is an insulated support, we

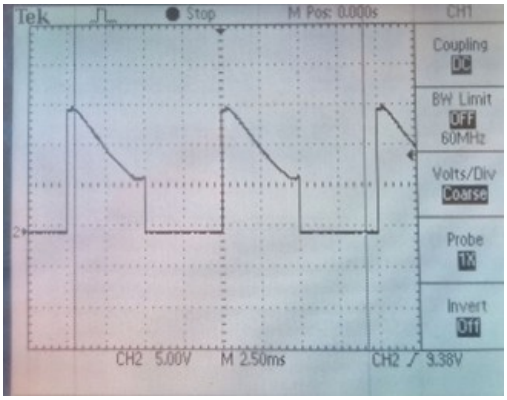


Figura 2.18: Generation of a PWM signal with frequency 100Hz

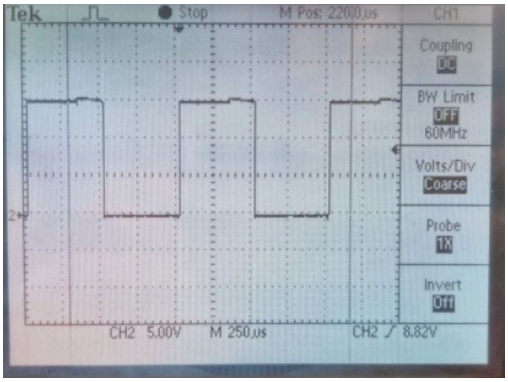


Figura 2.19: Generation of a PWM signal with frequency 1kHz

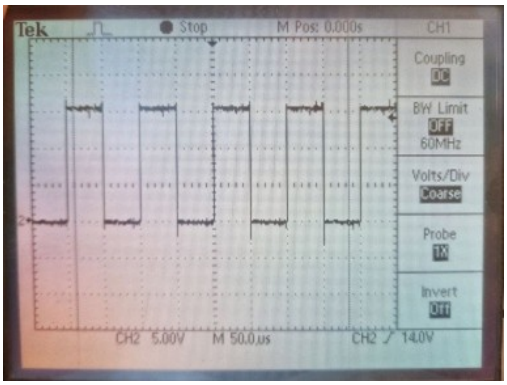


Figura 2.20: Generation of a PWM signal with frequency 10kHz

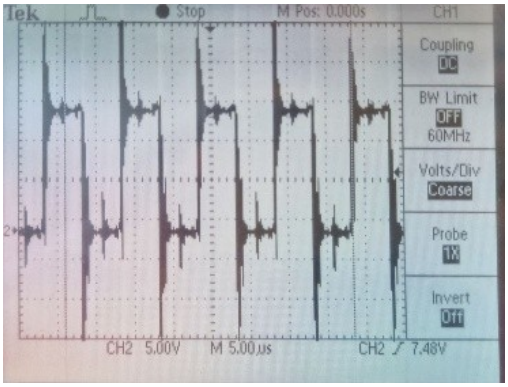


Figura 2.21: Generation of a PWM signal with frequency 100kHz

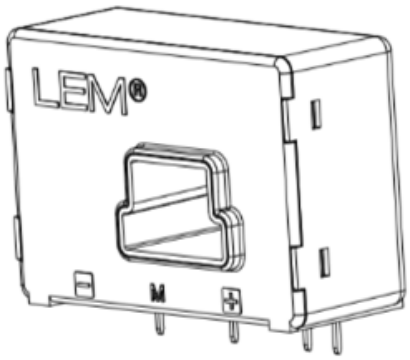


Figura 2.22: Hall sensors' appearance

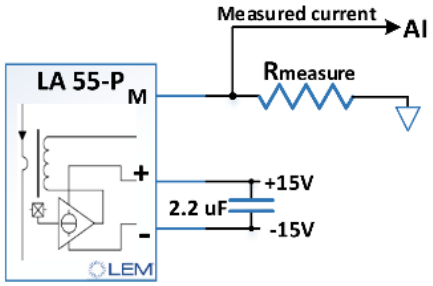


Figura 2.23: Hall-effect current measurement principle

could not push further than 6A for safety reasons. A correct configuration would include

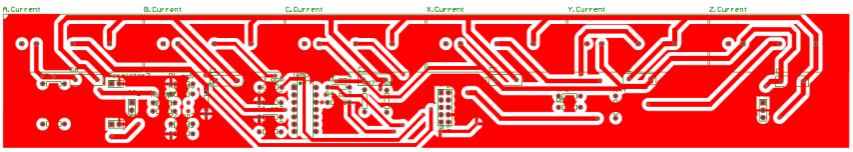


Figura 2.24: Measurement circuit PCB

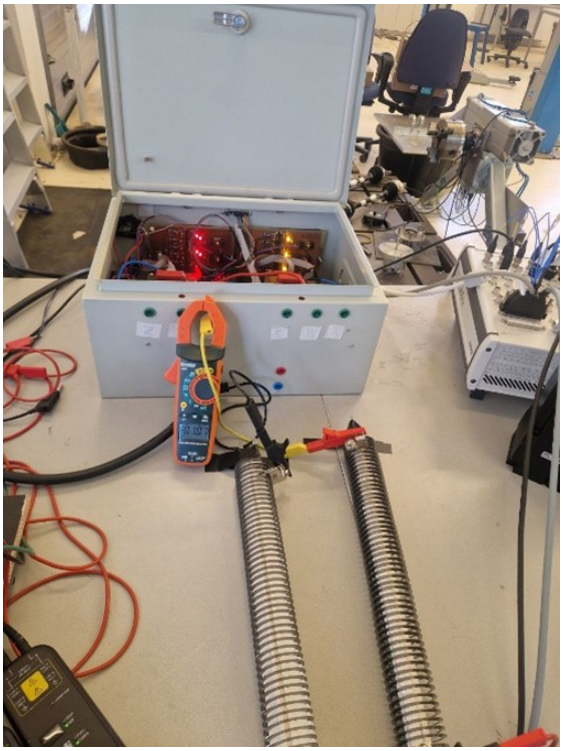


Figura 2.25: Current measurement calibration setup

the clamping of such resistors on a metal cage, but it was not possible.

This makes our comparison weaker, since our test conditions allow for currents up to at least 30Arms.

Offset and gain were computed through least square estimate:

2.4 Machine under test

We dealt with a surface mounted permanent magnets motor

- Base speed 1000rpm
- Max speed 3600rpm
- Power 800W
- Rms voltage 4V

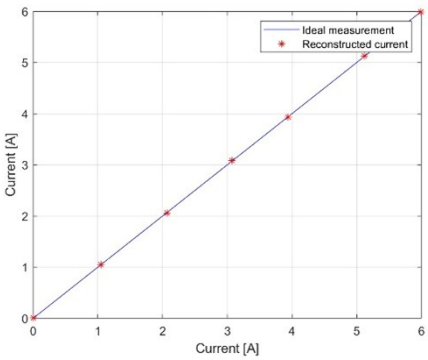


Figura 2.26: Ideal measured current over reconstructed current

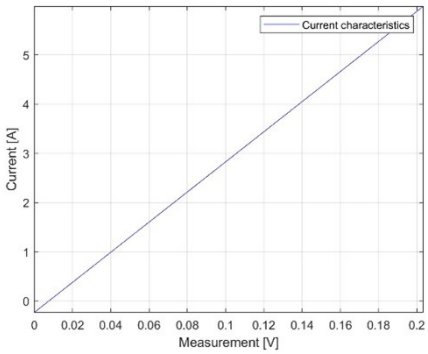


Figura 2.27: Current measurement system characteristics including calibration

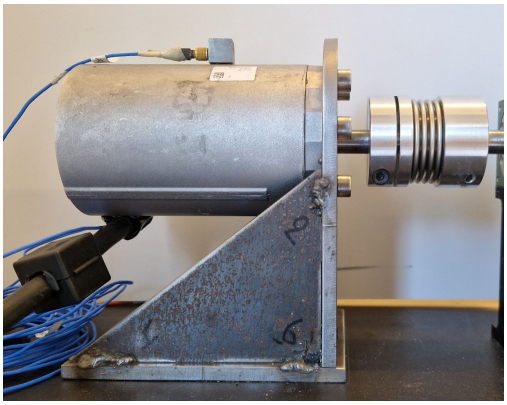


Figura 2.28: Machine under test



Figura 2.29: Machine under test's rotor

- RMS current 110A
- Concentrated windings, star connection
- Base torque 8.2Nm
- Max torque at max speed 2.3Nm
- 5 pole pairs
- Very low stator resistance 1.6 $\simeq \Omega$, not measurable with a standard resistor

From the nominal characteristics we can see that this is a low-voltage-high-current

motor, suitable for power steering applications from Brose.

FEA analysis was previously performed through E-machine Design, providing the following characteristics:

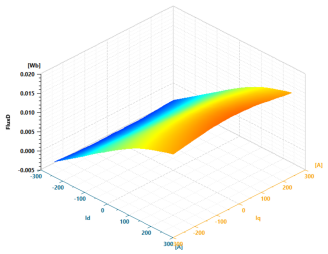


Figura 2.30: $\lambda_d(I_d, I_q)$

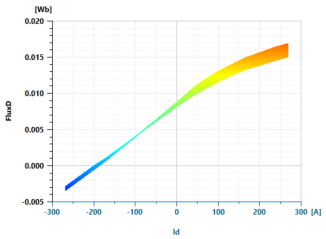


Figura 2.31: $\lambda_d(I_d)$

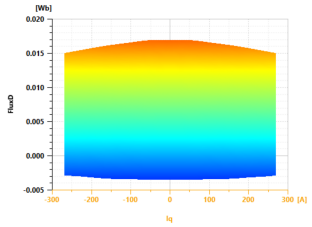


Figura 2.32: $\lambda_d(I_q)$

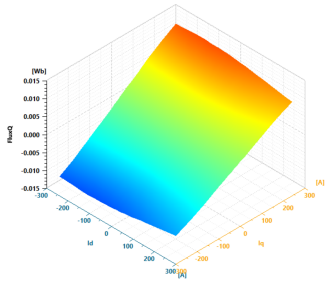


Figura 2.33: $\lambda_q(I_d, I_q)$

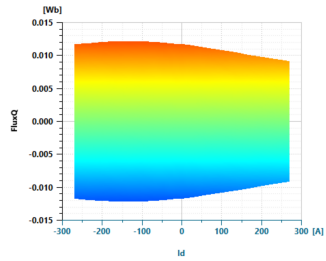


Figura 2.34: $\lambda_q(I_d)$

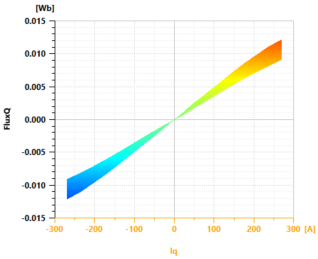


Figura 2.35: $\lambda_q(I_q)$

The motor's type can be easily spotted from that: $L_d \simeq L_q$, which suggests a low-saliency behaviour, very close to the one expected for an SPM motor. Moreover, at $I_d = I_q = 0$ we can see flux along d axis, roughly 0.0085Wb. We were then using the standard convention.

From torque maps instead, we can see how I_d has little to no influence on torque, again as we would expect. The MTPA is located in the second quadrant but very close to I_q axis.

Saturation can be witnessed in the same graphs, but very close to the rated currents thus neglectable. Both saturation and cross saturations were taken into account though.

Notice how these characteristics don't exactly match the specifications provided by the manufacturer; this is going to be the reason why our torque control might be inaccurate: steady state currents will be perfectly tracked due to the presence of PI controllers, whereas the torque does not represent a feedback in our control. This problem motivated the later characterization activity, since more accurate motor maps were required before moving toward reliable torque control.

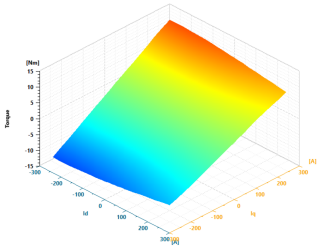


Figura 2.36: $torque(I_d, I_q)$

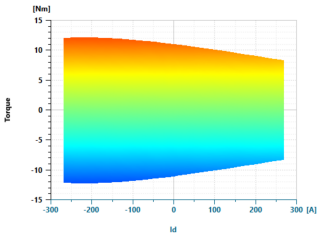


Figura 2.37: $torque(I_d)$

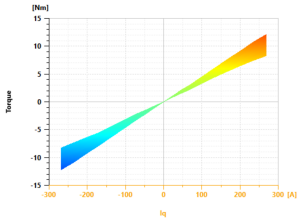


Figura 2.38: $torque(I_q)$

2.5 Loading machine

PD6 motor represented our loading machine throughout the whole testing period [Nanotec, 2025]. It was a ready-to-use and flexible tool, allowing us to work in different operating conditions.

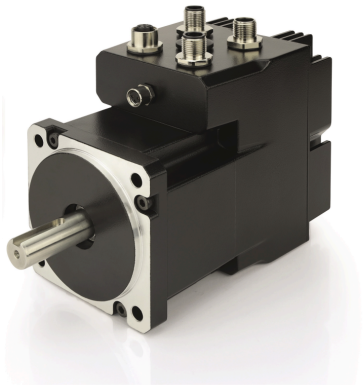


Figura 2.39: PD6-EB80CD-E-65-2

- Operating voltage (rated value) 48 V
- Permissible range of the operating voltage 12-57.6 V
- Logic voltage 24 V DC
- Voltage range of the logic supply 12-30V
- Rated current 25 A rms
- Peak current: 75 Arms
- Rated torque 3Nm, unfortunately not matching our testing conditions
- Rated speed 3000 rpm
- Rated power 940W
- Single-turn absolute encoder
- Support of voltage control, USB-CAN control, CAN control.

Its main limitation was the rated torque being lower than the machine under test, together with the complete absence of safety functions such as STO.

PD6 came with a USB-CAN cable to be connected from the motor to the laptop, and with "Plug & Drive Studio 3", a user-friendly program to directly access all the functionalities of the loading motor. Despite this, we wanted to get rid of such cable but especially of every dependency on third-party software. The goal was to control the whole test-bench in a unique environment, which was ControlDesk. Through ControlDesk we would not only monitor the useful data, but also feed them to our controller, so that everything could be integrated with each other.

The final aim was to create an RTI plug-and-play task, allowing the loading machine operating condition and the machine-under-test current references to be commanded directly from the real-time environment. Such possibility was provided by the loading machine thanks to its support of the CANopen protocol, that will be later discussed in Section 3.2

The change in the configuration is reported below.

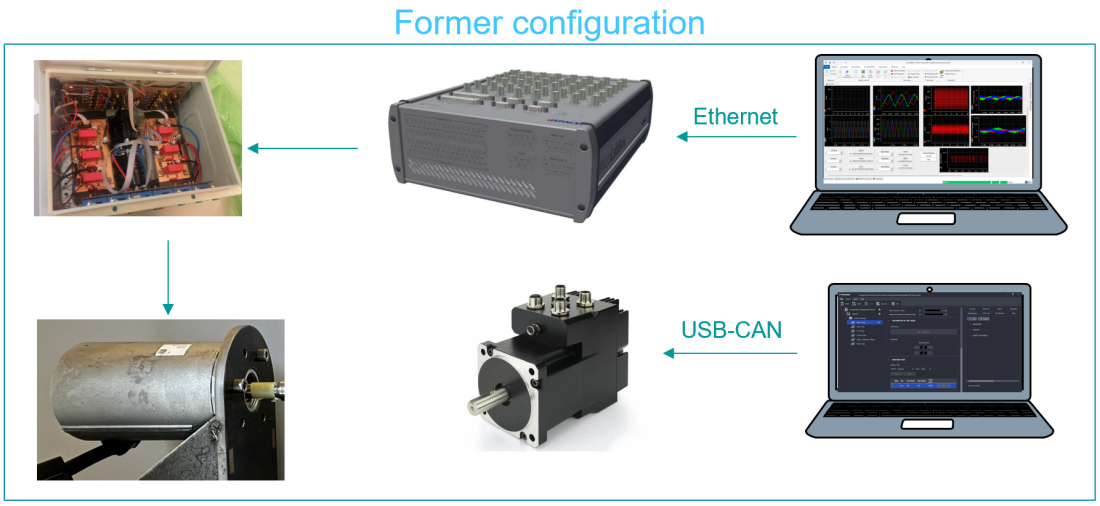


Figura 2.40: Former configuration

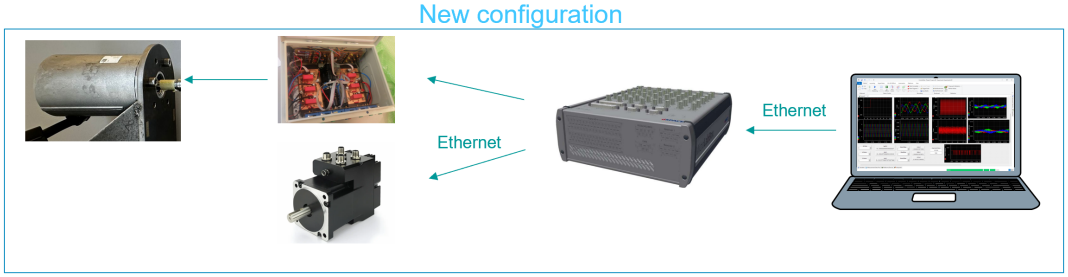


Figura 2.41: Updated configuration

2.6 First loading test

Our first loading test on the machine under test consisted of spinning it at different increasing speeds while measuring the line-to-line voltages. Before carrying out the experiment, an Amesim model was built in order to predict the behavior of the considered signals.

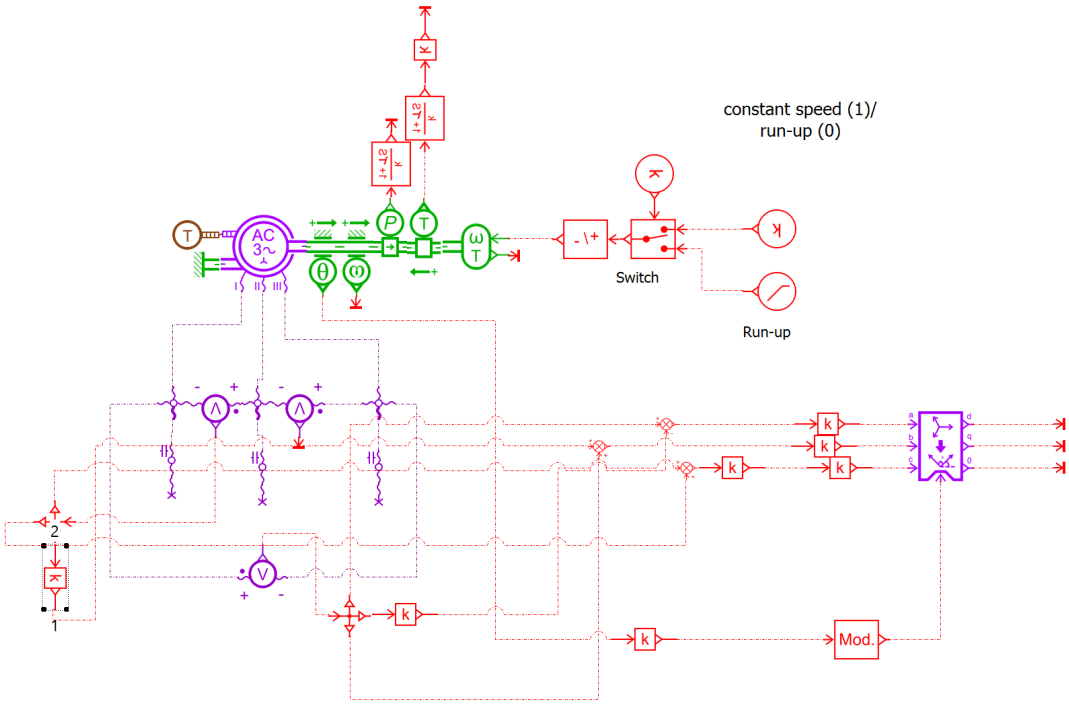


Figura 2.42: First run-up Amesim model

The top part represents the two machines, while the bottom part represents the measurements fed to dSPACE: the aforementioned line-to-line voltages, the reconstruction of the phase voltages, the Park transformation used to obtain V_{dq}, and finally the current imposition, which corresponds to an open circuit in all the phases, since the inverter is not connected.

Since the phase currents are imposed to be null, from the generic voltage equations of a PMSM we get:

$$V_{an} = R_s i_a + \frac{d\lambda_a}{dt}, \qquad V_{bn} = R_s i_b + \frac{d\lambda_b}{dt}, \qquad V_{cn} = R_s i_c + \frac{d\lambda_c}{dt}$$

but, since $i_a = i_b = i_c = 0$, the line-to-line voltages allow us to trace back to the back-EMF, or better to their instantaneous difference.

For every 150rpm, our motor's characteristics lead to an increase in the differences of back-EMF of about 1V. Since our measurement system ranges from -10V to 10V, we had to add three voltage probes.

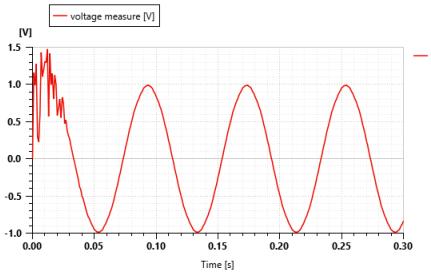


Figura 2.43: Line-to-line voltage simulation at 150rpm

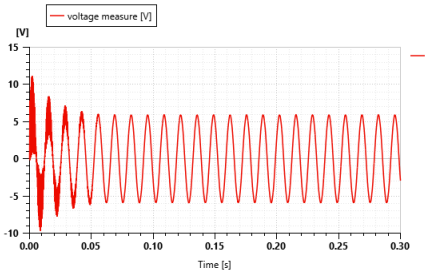


Figura 2.44: Line-to-line voltage simulation at 900rpm



Figura 2.45: Tektronix P5200A voltage probe

Such probes provide sufficient protection against over voltage, by dividing the voltage measurement by 50 times, therefore allowing a voltage measurement through MicroLabBox up to 500V, much more than needed. The input to our measurement system is then digitally multiplied by 50.
The main drawbacks of this procedure are:

- Necessity of a different power source, for all three of them
- Need for calibration
- Mandatory filtering action due to the presence of noise.

Here below we reported the outcome of the experiment for one of the three legs at 150rpm, therefore expecting an AC signal of about 1V peak.

Results of our first loading test at different speeds are reported below. The three voltage probes helped us to appreciate the 120 degrees displacement of the line-to-line voltages, even in non-stationary conditions, which are the ones we are interested in.

Predictions were confirmed by this first testing stage. Since we are exploiting a shrunk fraction of the probes measurement range, a clear difference can be spotted between different velocities: at low velocity, and thus with smaller generated voltages, the noise-to-signal ratio is larger, while at 6V the noise is smaller with respect to the fundamental frequency signal.

A graph with all the three line-to-line voltages was also added, highlighting both their correct phase displacement and the correctness of the convention, since V_{ac} succeeds V_{cb} that succeeds V_{ba} with a positive velocity; if the convention was wrong, we would have

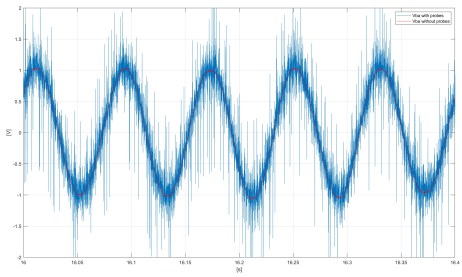


Figura 2.46: Line-to-line voltage measurement with and without probes

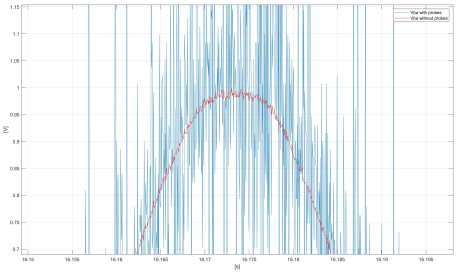


Figura 2.47: Line-to-line voltage measurement with and without probes close up

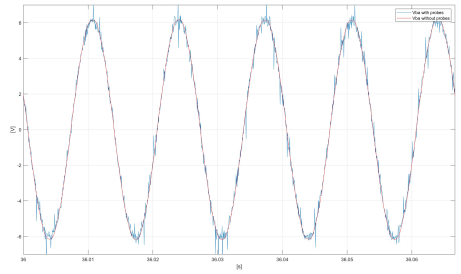


Figura 2.48: Line-to-line voltage measurement with and without probes at 900rpm

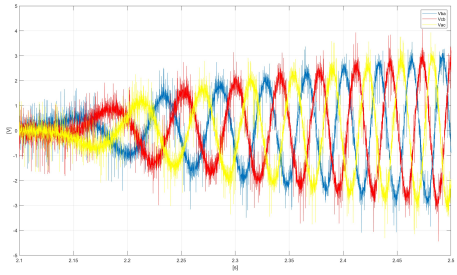


Figura 2.49: All three line-to-line voltage measurements with probes and with a ramp velocity

the opposite succession order.

To the same signals an FFT was applied in order to delve into the harmonics study, and at the same time to retrieve additional information about the distinction between machine under test's harmonics and noise coming from the acquisition setup.

The test was conducted at 150rpm with variations of $\pm 5rpm$. Since the machine has 5 pole pairs, the corresponding electrical frequency is around 12.5Hz.
The three main peaks are located around the main electrical frequency. We have three of them due to fluctuations in the loaded speed, but we can also witness the presence of additional peaks around multiples of it.
Reprocessing the same tests by focusing only on the expected harmonic frequencies, we selected the bins placed at the electrical frequency and at its multiples, where we expect the motor contribution to appear.

The investigation gave a negative result, i.e. mechanical harmonics can be claimed to be irrelevant.
In the same graphs, though, we also reported the signal coming from a connector without probes, to go deeper into the study of the measurement system we adopted.

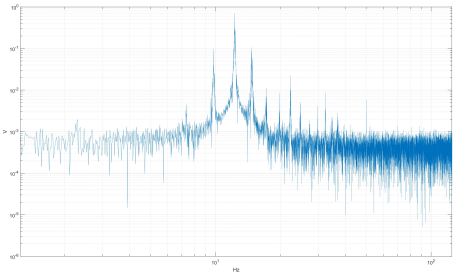


Figura 2.50: FFT applied to one line-to-line voltage, close up to the main electrical frequency

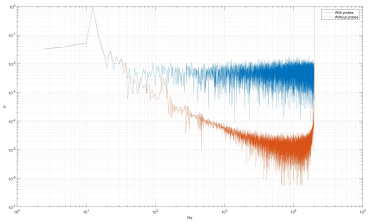


Figura 2.51: FFT comparison of the line-to-line voltage with and without probes

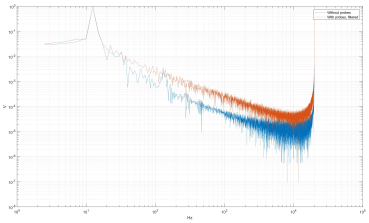


Figura 2.52: FFT comparison after filtering the signal measured with probes

The most trivial effect was the introduction of a high-frequency noise having a quite constant trend over frequency. This noise was visible in figure 2.47, making the fundamental frequency invisible to the naked eye.
The introduction of a low pass filter, as shown in picture 2.52, allowed us to fully reconstruct the original signal with a high signal-to-noise ratio.

Eventually, we applied the same methodology, but focusing only on the electrical frequencies of interest, which are the odd multiples. When we filter, in fact, we want to make sure to only filter the addition given by our measurement system, and not mistake the motor's harmonics for noise. Catching fast varying signals is the base we should be equipped with in order to develop more complex control algorithms against them.

Here below I reported such study at several different velocities.

Very low harmonic content is shown at electrical frequency too. Such phenomenon is even more relevant than its mechanical correspondent, because electrical harmonics in the motor voltage waveforms translate into harmonic components in the magnetic field. These harmonic magnetic fields create unwanted torque ripple, additional losses and heating, reducing the smoothness of the operation.

It is also important to apply the same study to the remaining legs and then swapping the probes, to be able to tell whether imbalances were present and if they were due to the probes or to the three-phase motor's little imbalances.

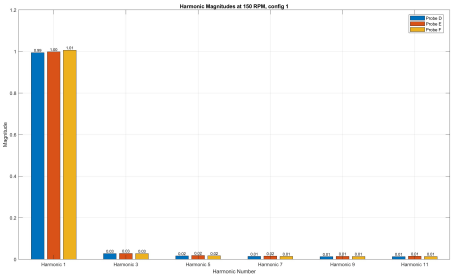


Figura 2.53: Electrical odd harmonics at 150rpm

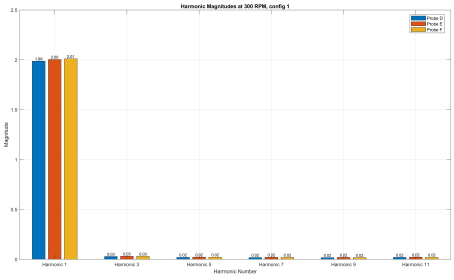


Figura 2.54: Electrical odd harmonics at 300rpm

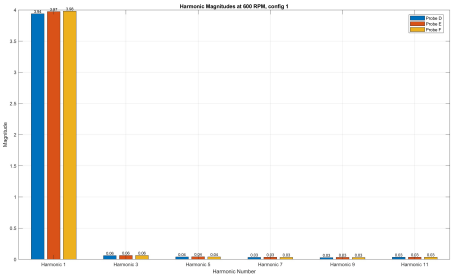


Figura 2.55: Electrical odd harmonics at 600rpm

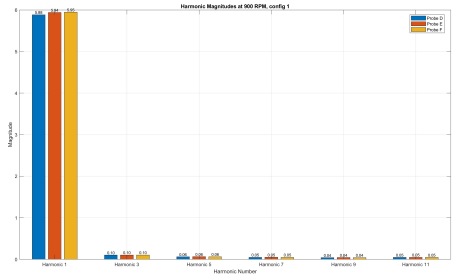


Figura 2.56: Electrical odd harmonics at 900rpm

No imbalance was detected, or just a very subtle one that brought about a smaller voltage on phase a and a larger one in phase c, proving our voltage probes needed little-to-no calibration. After many run ups, we could state such relative error settles at less than 1%, therefore our stator's windings are well balanced.

2.7 Phase determination procedure

While the inverter has a clear reference for the legs, a three-phase motor has, in general, none for its phases. If the load represented by the interaction between the stator currents and the magnetic flux inside the rotor that generates the back-EMF is well balanced as well as the three stator resistances and windings, there is no such thing as phase a, b or c before defining a convention.
However, it is fundamental for every future test to always have the same convention. The procedure entails sending a steady voltage signal from the inverter's leg we want to be phase a (in the inverter's figure, leg z). That is to say, $V_z = 1V$ (for instance) but also $V_y = V_x = -0.5V$ since voltages should always be balanced. This way we will get our rotor's magnet's north pole to align along phase z. Then, we proceed by resetting the encoder angular measurement to 0deg. This way we cleared the procedure needed to have

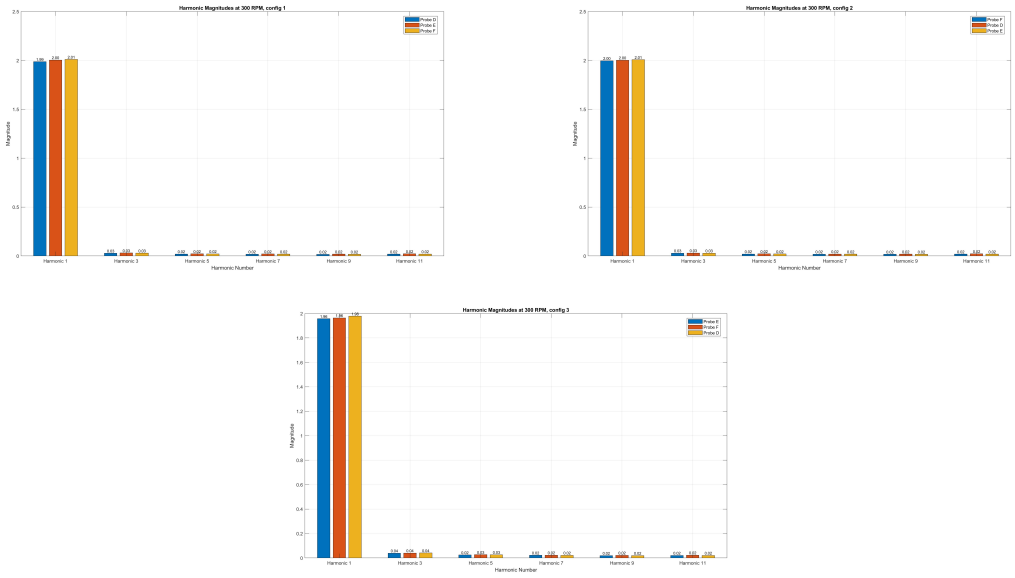


Figura 2.57: Electrical harmonics analysis while swapping the probes

$V_z = V_a$.
Second, we impose the same voltage across one of the other two arms to get a counter clock-wise rotation of the machine under test, then check that the encoder recorded a positive rotation. If not, just reverse the angle's convention.
In our case, we tested and checked y to be phase b, i.e.

$$V_{zyx} = [-0.5 \quad 1 \quad -0.5]$$

that gave us a positive rotation of $120/n_{pp} = 24deg$.

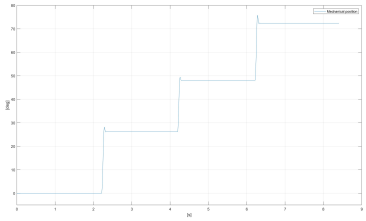


Figura 2.58: Mechanical angle during one electrical rotation

2.8 Open-loop control

The meaning of open-loop control is generally an algorithm that provides a constant signal, which is an input, to the plant model and then measures the output response. The main difference with a closed-loop control is how we manage such acquisitions; while in closed loop control we do have a system that checks instantaneously the speed (for instance) adherence to the desired one and acts upon the input, the same happens in open-loop control but with no adjustment of the input signal.
It is clear that, while in the closed loop control the command is controlled by a controller usually acting on the state error, in open-loop control our input is straightforward a command for the system.
In CVC FOC open-loop control, this principle is still respected by stating the wanted voltages to the inverter, but we are generally interested in the direct and quadrature voltages more than the phase voltages. In the light of this, the implementation of an open-loop algorithm is open-loop with respect to current and speed, but still dependent on the measured rotor position. Rotor position is used to switch from V_{dq} to V_{abc} for the command obtainment, and then in the measurements system from V_{abc} to V_{dq}. This process is mathematically realized by the Park transformation matrix.
This adds more space to a possible failure, which is likely the wrong measurement of the rotor position. In the next Chapter we will discuss this, how such inaccurate variable could make the whole mechanism inceppare.

2.8.1 Open loop control initial state-of-art

For the several tentatives of open-loop control, we employed the following scheme

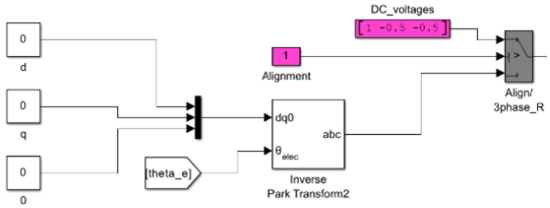


Figura 2.59: Voltage open-loop control on Simulink

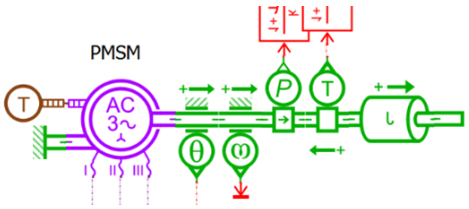


Figura 2.60: SPM modeling on Amesim

On the left, the command generation system is inclusive of the alignment. On the right, the test-bench: the machine under test is connected to the loading machine that is not acting as the load; its only influence is on the inertia of the system.
This simple test aimed at feeding V_d and V_q, reconstructing them from the measured signals (here after, line-to-line voltages). Here below you can see the expected outcome of a supply $V_{dq} = [0\ 2]$:
At first, V_q supply to the plant brings about torque that makes the machine under test

start spinning. Currents follow the same course as the torque since as we know they are strictly related.

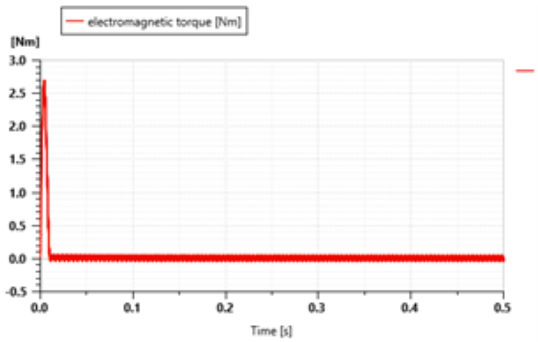


Figura 2.61: Predicted torque

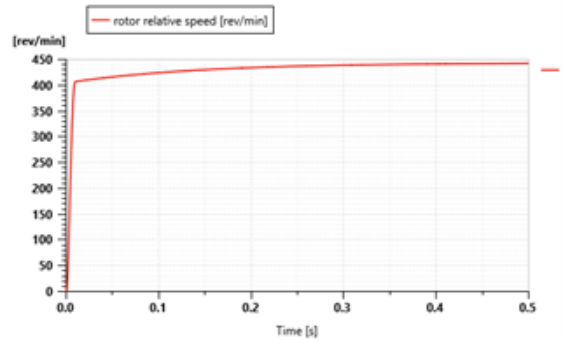


Figura 2.62: Predicted velocity

At steady state the rotor spins at constant rpm, currents are null and so is the torque: the back-EMF has completely saturated the voltage supply, and no further velocity could be reached unless a larger supply of voltage is achieved.
This phenomenon can be explained by the voltage equations, were current terms can be canceled at steady state and the voltage is fully exploited on the motional term. This simple correlation is the main reason why we command quadrature and direct quantities instead of the phase ones.

$$V_d \simeq -\omega_e \lambda_q$$
$$V_q \simeq \omega_e \lambda_d$$

Results were not encouraging at first. After providing the same amount of voltage, the rotor would start spinning in the right direction which suggests a correct alignment procedure; such speed would slowly drop and eventually after no more than twenty seconds the rotation would stop.

From V_q, we could state that everything works correctly: the command is correctly obtained at the terminals of the inverter, with the proper transformations, and with the inverse transformations reconstructed back. But dynamics of the system are not satisfactory.

During the whole test, we are able to measure a torque signal that suggests the presence of friction in the system. This friction does not explain by itself the time course of the velocity, since dynamic friction is proportional to the speed, thus it should affect transient dynamics and limit the velocity at steady state, surely not the reason why rotation is interrupted.
The unfiltered speed (which is already filtered) explains it all: something is affecting the measurement of the rotor position, thereby losing the angle.
Since the angle is lost, a command on V_q is not properly converted to phase voltages V_{abc}, and the plant response if far from the predicted one. But, when measuring V_{abc} and applying the inverse transformation, we are still employing the wrong angle to transform,

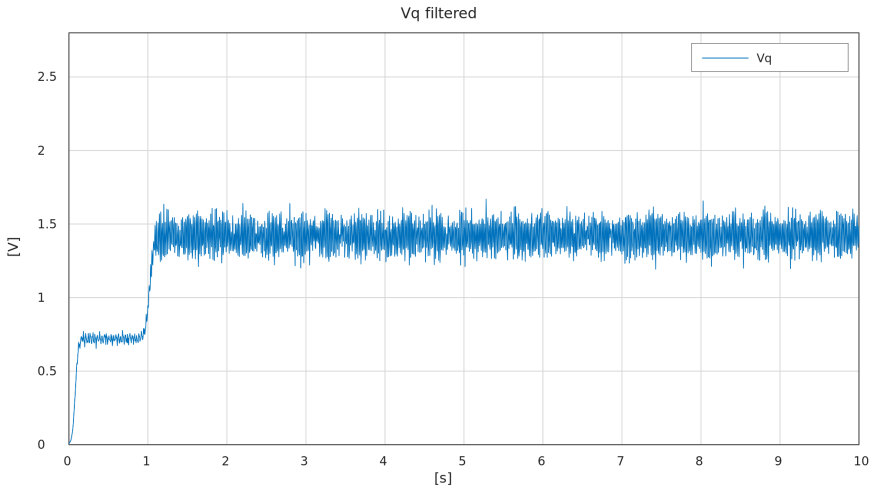


Figura 2.63: Vq command during the failed open-loop test

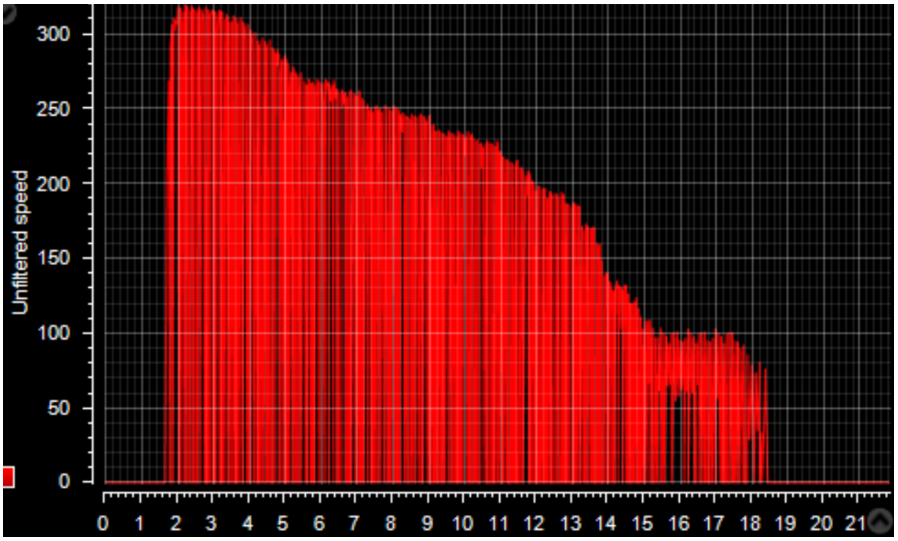


Figura 2.64: EMC EMI effect on velocity

so on V_q reconstruction side everything looks smooth and to be working.
Instead, since the angle is wrong, instead of supplying $V_d = 0$, $V_q = 2V$ we are slowly supplying a larger and positive V_d, and a smaller V_q. Since the machine under test is a surface mounted permanent magnet, a different V_d does not affect the speed, and decreasing V_q does by decreasing it and this is exactly the effect we get. In current terms, we are moving out of the MTPA. In encoder terms, we are using the alignment, which is the first fundamental procedure.
Other experiments took place highlighting how dangerous an angle loss could be: if, at times, the angle loss were fast enough and voltages were sufficiently high, an overly angle

loss could take place in few sampling instants, falling into a situation where angle loss is >90 degrees. This leads to the machine under test to start spinning backwards!!!!

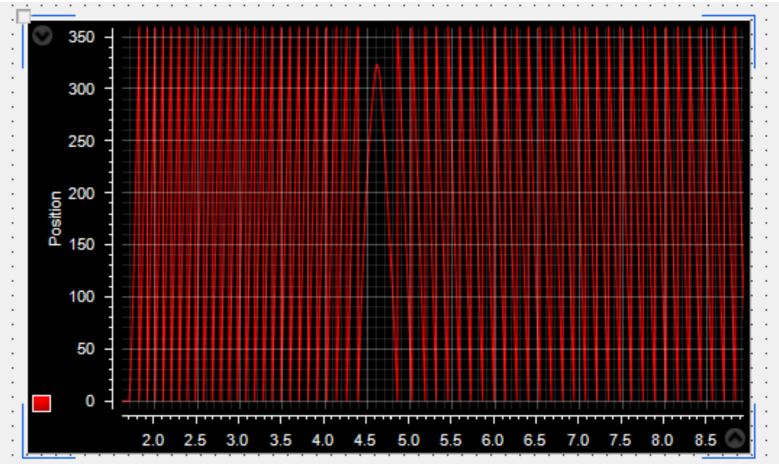


Figura 2.65: Rotor position during a failed open-loop cycle leading to reverse rotation

2.8.2 EMC EMI

Electro magnetic interference, also known as EMI, is a general term to describe the electric disturbance caused by one system considered external to another considered of interest. Since in motor applications the general trend is to pack the several components closer and closer,thus o increase the power density, this topic must be taken on in this paper.
In our case, the simultaneous presence and function of the power supplies, inverter, and three-phase machine and DC one brought about many possible sources of interference that compromised the correct working of the encoder, mainly.

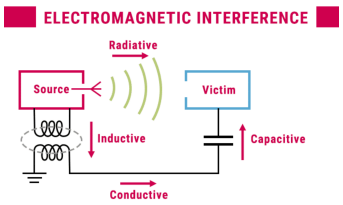


Figura 2.66: EMC EMI possible sources

There could be four types of interference: inductive, capacitive, conductive and radiative [Analog Devices, 2009].

- Radiative: one circuit acts as an antenna, sending electromagnetic waves to another one that catches them thereby ruining its correct working principle.
- Conductive: common impedance coupling occurs when two or more circuits share a common electrical path, and the current from one circuit causes a voltage drop

across that common path, which then affects the other circuits. This is the only kind of interference that takes place by direct contact between two circuits.
- Capacitive: electric field coupling occurs when an electric field from one circuit induces a voltage in another nearby circuit. Two parallel plates start forming a capacitor, but an unintended one.
- Inductive: also known as magnetic coupling, occurs when a changing magnetic field from one circuit induces a voltage in another nearby circuit thereby acting as an inductor.

Stumbling upon all these kinds of interference was unavoidable due to our test bench configuration.

Causes of EMI coming from the inverter: The main troublesome mechanism inside the inverter is by far the PWM switching one. Switches have the duty to rapidly turn the power on and off.
The switching action can cause conductive interference: its rapidity can create current to back propagate into the power supply, disrupting its output signal. This disruption results in voltage ripples and spikes.
Parasitic inductances of the motor cables and again the rapid change of currents, more voltage spikes will be induced. Because of the high frequency of these signals, electromagnetic fields will be radiated.

While the perfect configuration for attenuating EMIs would see the presence of a single point grounding, the configuration we could deal with would include a common ground for all the devices in our system.

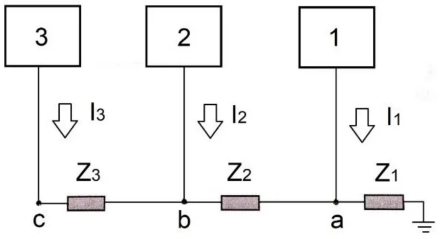


Figura 2.67: Single-point ground scheme

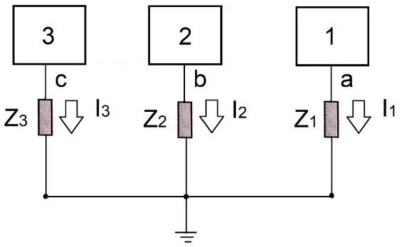


Figura 2.68: Multi-point ground scheme

In an ideal test-bench, during every test we would expect currents to flow from our system to the ground and vice versa, with no voltage drop across the connections.
This does not happen in reality, where voltage losses happen in between. In figure, such distributed loss are represented by impedances Z1,Z2,Z3 that for sake of simplicity we can think of them as resistive only. Given the currents I1, I2, I3 is clear that for the common ground configuration we have:

$$V_a = (I_1 + I_2 + I_3)Z_1$$

$$V_b = (I_1 + I_2 + I_3)Z_1 + (I_2 + I_3)Z_2$$
$$V_c = (I_1 + I_2 + I_3)Z_1 + (I_2 + I_3)Z_2 + I_3Z_3$$

This means that the ground in a real application is not shared equally among all the devices, instead it is varying across the ground line. When the ground currents flowing in the common paths are low or absent, the reference potential is essentially the same in all subsystems, but this is not the standard situation we will fall in. The presence of a shared ground path brings about interaction between the circuits, because the impedance of the common return path is not null. This is particularly dangerous when operating with very different energy levels such as power and control: while power supply, inverter and motors are in the same wavelength, MicroLabBox is not. In this configuration power equipment's high energy impulses may couple to the control signals.
For the single ground configuration:

$$V_a = I_1Z_1$$
$$V_b = I_2Z_2$$
$$V_c = I_3Z_3$$

i.e., we switch from series to parallel impedance. This configuration mitigates common-mode noise, but it is more costly and bulky. It is not compact enough for the final test-bench requirements.
Pic of the first configuration.

Since we couldn't rely on the second option, we made sure to properly attenuate the interaction between:

- Power stage, represented by the inverter arms; for this reason, we added ferrite cores on inverter's legs
- Control stage, represented by the PWM signals to the driver and the input from the encoder and torque sensor. We added ferrite cores here too
- Control stage, represented by the PWM signals to the driver and the input from the encoder and torque sensor. We added ferrite cores here too [Murata Manufacturing Co., Ltd.].

Common vs separate DC supply: a measure we adopted in the early stages of test-bench design was the possibility to exploit two different power supplies.
In case of the implementation of a motor control power stage with a shared line, just like before there is no voltage drop (thus power losses) along the return path for the current if everything works ideally. In real applications instead, a distributed impedance will be also present thereby the voltage drop acts as an unwanted signal or noise for the other circuits sharing that same path. This is known as "common Impedance coupling", leading the two circuits to interact with each other. This is a type of conductive electromagnetic interference, since the two circuits are physically connected.
Finally, a common DC supply was used since the final application required it

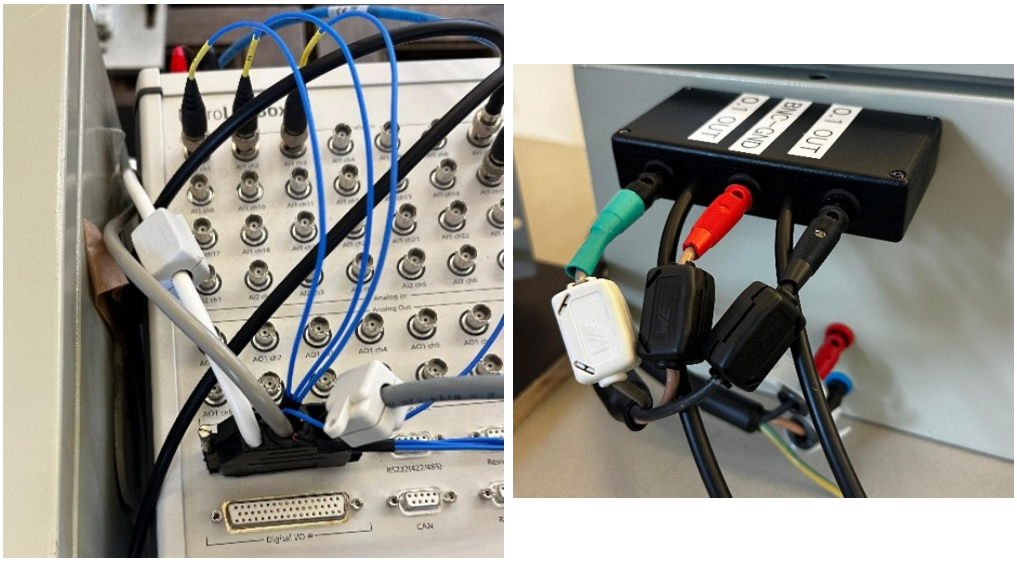


Figura 2.69: Ferrite core addition

Capacitor The presence of a capacitor as a DC link for the loading machine represents both a drawback and an advantage while talking about EMIs. Other than its primary role in storing charges as we already discussed, it has secondary functions:

- It acts as a short circuit to high-frequency signals while blocking DC, which is fundamental in applications of PMSM control like ours.
- It also acts as a filter to noise by shunting high-frequency noise to ground or providing a low-impedance path for ripple currents.

Other than that, there are a couple of phenomena worth mentioning. First and foremost, we mustn't forget the inductive nature of Electric motors; the co-presence of a capacitor and one (or multiple in this case) inductance could be dangerous, due to the possibility of forming an RL circuit with the possibility of the system noise excite this resonant frequency, leading to very high circulating currents or overvoltage. This phenomenon will instead certainly happen in case of employment of a variable frequency drive, not our case. Second, since the capacitor is intended to smooth voltage spikes coming from the loading machine, it can sometimes create sharper current edges.
For this reason, the capacitor was stably placed far from the power devices.

Encoder shielding
Other than the aforementioned solutions to mitigate electromagnetic interference between devices, we adopted one that turned out to be a game changer.
Insulation protects a cable mechanically from scraps and abrasion and environmentally from moisture, but it is transparent to electromagnetic energy and offers no protection. Shielding is needed to combat the effects of EMI. The working principle is the same as the Faraday cage.

The shield surrounds the power-carrying conductors. The shield can act on EMI in two ways. First, it can reflect energy. Second, it can pick up the noise and conduct it to ground. In either case, the EMI does not reach the conductors. Some energy still passes through the shield, but it is so highly attenuated that it doesn't cause interference.

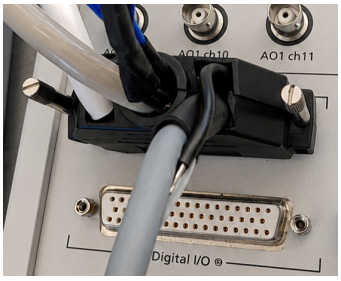


Figura 2.70: Shielded cable

The cable we employ uses foil to realize the shielding action. Foil shielding includes a thin layer of aluminum, typically attached to a carrier such as polyester to add strength and ruggedness. It provides full coverage of the conductors it surrounds.
Braid should also be added. It provides a low-resistance path to ground and is much easier to terminate by crimping or soldering when attaching a connector [Belden Inc.].
In our case, the aluminum foil solution was employed only, but the possibility of adding a braid was considered too.

2.8.3 Open-loop control after shielding

By reperforming the same experiment as before, thereby imposing V_q at different levels, we could finally experience the correct functioning of the encoder. Unlike the previous tests, the position measurement remained stable and the machine behaved as expected.

The Machine Under Test was mechanically connected to the Loading Machine, the encoder, the torquemeter and the couplings. Since the Loading Machine was not actively imposing a load during this test, these components mainly contributed to the total rotary inertia of the system.

When a positive V_q is applied, a quadrature current initially flows through the machine and produces electromagnetic torque. The rotor therefore accelerates and the electrical speed increases. From the q-axis voltage equation:

$$V_q = R_s I_q + \frac{d\lambda_q}{dt} + \omega_e \lambda_d$$

we can see that, as the speed increases, the motional term $\omega_e \lambda_d$ also increases. Such torque leads to an increase of the speed, thus also increasing the motional term in the voltage equation. This term keeps growing until most of the provided voltage is employed to compensate the back-EMF, thereby ending the accelerating action.

Under the ideal no-load steady-state approximation, the current becomes very small and the flux-linkage variation can be neglected. Since the d-axis flux is then mainly represented by the permanent-magnet flux linkage, the voltage equation can be approximated as:

$$V_q \approx \omega_e \lambda_{pm}$$

Therefore, for a constant applied V_q, the rotor tends towards a corresponding steady-state speed.

The experimental steady-state values did not exactly match the ones predicted by the Amesim simulation. This was not unexpected, since the initial motor characteristics used in the simulation were obtained from approximate FEA data. At this stage, however, the main objective was not to obtain an exact correspondence between simulation and experiment, but to verify that the real system reproduced the expected physical dynamics.

In this case, the transient lasts a few seconds before reaching the final speed. After that, the open-loop control can keep such speed without the progressive drop experienced before. More importantly, the position measurement remains correct throughout the test, which was the result we were mainly looking for after introducing the shielding.

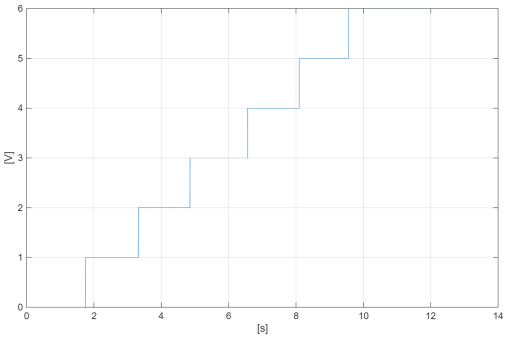


Figura 2.71: V_q command

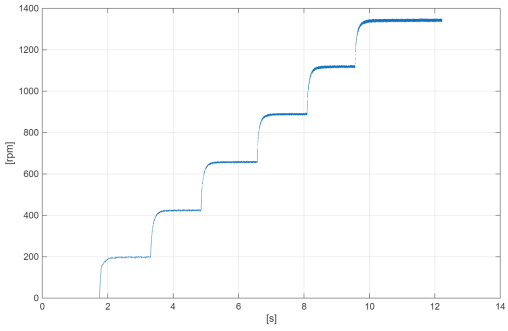


Figura 2.72: Resultant speed of the rotor

A further check was performed by imposing V_d, both with positive and negative values. In this case, almost no mechanical reaction was observed. This behaviour is coherent with the low-saliency SPM model discussed in Chapter 1.

With the power-invariant Park transformation adopted in this work, the electromagnetic torque of the machine can be approximated as:

$$T_e \simeq p \lambda_{pm} I_q$$

because $L_d \simeq L_q$ and the reluctance torque contribution is therefore small.

Consequently, when no quadrature current is produced, the electromagnetic torque remains approximately zero:

$$I_q \simeq 0 \qquad \Rightarrow \qquad T_e \simeq 0$$

This is exactly what we experienced during the test: acting on V_d alone produced almost no rotation, while V_q was the component actually responsible for the acceleration of the machine.

Capitolo 3

Current to torque control

After implementing open-loop voltage control, the next step was to move to closed-loop current control. In this configuration, the controlled quantities are the d- and q-axis currents, while the outputs of the current controllers are the corresponding voltage references V_d and V_q.

3.1 Current control and parameter tuning

Like in the most common applications, we relied on the use of two PI controllers. A PI controller is a tool that takes as input a desired value of a current, its real measured one as well, computes the difference between them, known as tracking error, and provides as output the corresponding voltage needed to close such gap. A PI controller employed for current control has the following time domain expression [Bojoi and Pescetto, 2025, Sul, 2011]:

$$u(t) = K_p e(t) + K_i \int_0^t e(t')\, dt'$$

therefore, applied to our specific case:

$$V_d(t) = K_p e_d(t) + K_i \int_0^t e_d(t')\, dt' \qquad V_q(t) = K_p e_q(t) + K_i \int_0^t e_q(t')\, dt'$$

Thus, the voltage command is given by the sum of a term that is proportional to the tracking error, and one that is proportional to the integral of it. It is important to include both because they have different effects on the response of the system. A P-only controller can be employed, but we have no guarantee that the steady-state error will drop to zero. Actually, that is not the case in general. The integral part is what makes sure to reach the desired value, unless instability or saturation problems arise.

In frequency domain, the same PI controller shows the following expression:

$$PI(s) = K_p + \frac{K_i}{s} = \frac{K_p s + K_i}{s}$$

This transfer function contains one pole at the origin, introduced by the integral action, and one zero located at:

$$s = -\frac{K_i}{K_p}$$

The role of the PI controller is to combine two effects: the proportional term reacts immediately to the current tracking error, while the integral term accumulates the error over time and removes the steady-state error.

In Figure 3.1 I reported the general scheme of a PI controller used to close-loop control a generic plant. The final goal is to have a stable closed-loop system, but in automatic control it is very common to start from the study of the open-loop transfer function. More precisely, what is studied is the loop transfer function:

$$L(s) = C(s)P(s)$$

where $C(s)$ is the controller and $P(s)$ is the plant.

This is useful because the behaviour of the loop transfer function gives information on what will happen after closing the feedback loop. From its Bode plot we can understand the bandwidth of the system, and also have information on stability margins. For this reason, instead of choosing the PI parameters only by trial and error, we first try to shape the open-loop transfer function in a reasonable way. After that, the values are checked and corrected experimentally [Bojoi and Pescetto, 2025].

In order to calibrate the current control, we generally need to measure: phase currents, to be converted in rotary frame currents and fed again to the controller as feedback, rotor position for the Park transformation, and the apparent inductance l_s. The permanent magnet flux λ_m is instead more useful when back-EMF compensation or torque estimation are considered.

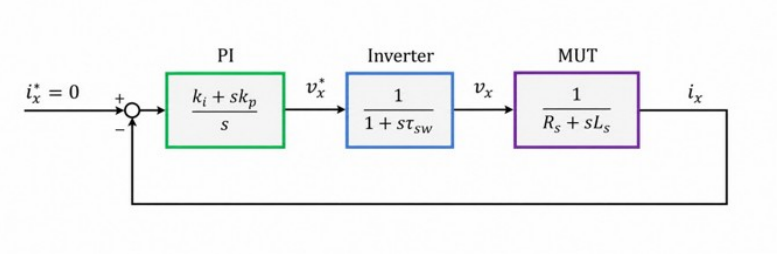


Figura 3.1: Generic current control

As reported in the picture above, the plant is represented by the system inverter + motor. The inverter turns desired voltages into actual ones, therefore it is modeled as a first order lag. In a very simplified fashion, the PMSM is instead modeled as a first order lag between voltage in one direction and the corresponding current. That is to say that the motor resembles a first order RL system with time constant:

$$\tau_s = \frac{l_s}{R_s}$$

This is obviously an approximation, because the real PMSM dynamics depend on the operating point and, when the rotor is spinning, also on the cross-coupling and back-EMF terms. Nevertheless, it is still useful as a first starting point for the tuning of the current controller.

For the calibration phase, we want the open-loop transfer function to have a good behaviour around the crossover frequency. The crossover frequency must be much lower than the sampling and PWM frequency, otherwise the real-time control would try to react to dynamics that are too fast for the implementation.

Around the crossover frequency, a slope close to $-20\,\mathrm{dB/dec}$ is usually preferred, because it gives a safer phase margin than a steeper crossing. For this reason, the zero of the PI controller,

$$\omega_z = \frac{K_i}{K_p}$$

has to be placed coherently with the motor and inverter dynamics. In simple words, the PI zero should help increasing the gain at low frequency, so that the current error is removed, without making the system too oscillatory or unstable.

At first approximation we can say that:

$$\omega_b \approx \frac{k_p}{l_s}$$

therefore the only constraint is:

$$k_i \ll k_p \omega_b$$

so we have one degree of freedom.

The parameter tuning is executed this way:

$$\omega_b = \frac{2\pi f_s}{10} \qquad k_p = \omega_b l_s \qquad k_i = \frac{k_p \omega_b}{20}$$

Where 20 and 10 at the denominators are safety factors used to keep the controller far enough from the sampling frequency and to avoid an excessively aggressive integral action.

With this procedure in mind, three problems show up:

- Only an approximate model was employed for the SPM motor. This model is useful for a first tuning, but it cannot describe all the real dynamics of the system.
- The apparent inductance l_s is not a constant value given only by the system geometry, but instead it is a differential quantity that depends on the operating point. For such reason, a good norm is to pick a conservative value inside the current range of interest.
- Wide inaccuracy on the flux maps, as discussed before, can make torque estimation and torque-to-current conversion not very reliable.

This method can only provide a starting value for our algorithm, that then is tuned by trial and error.

Final approximate transfer function will be:

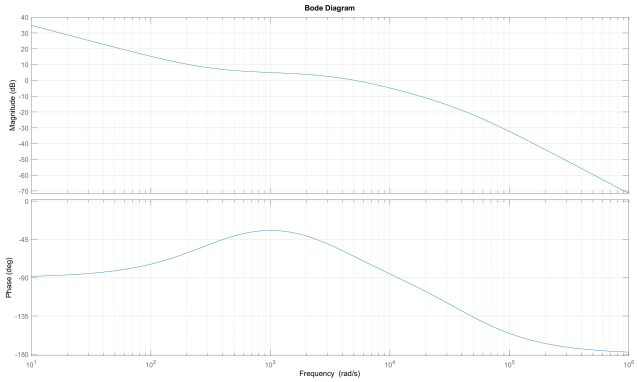


Figura 3.2: Bode plot used for the first PI tuning

We can start by evaluating the single contribution:

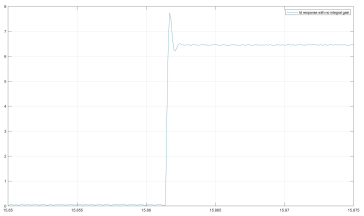


Figura 3.3: P only control when asked $I_d = 10A$

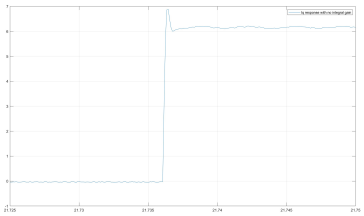


Figura 3.4: P only control when asked $I_q = 10A$

In the pictures above it is shown the step response of our system with the computed values of k_p, but with $K_i = 0$, to highlight the importance of the integral gain. Without that, the proportional part can provide an almost instantaneous response to the current demand, but it does not reach the desired value. In fact, with only the proportional gain, a non-zero steady-state error is needed to generate the voltage required by the plant.

In these last pictures instead, an I-only control proves that the integral action is enough to remove the steady-state error. The addition of a proportional part helps to improve the performance indexes, especially the transient response.

Moreover, we can highlight the symmetry of the system: it looks like the dynamics of the system in response to a command along d axis is similar to the one along q axis. This is coherent with the behaviour expected from a SPM motor, where the apparent inductances along the two axes are close to each other.

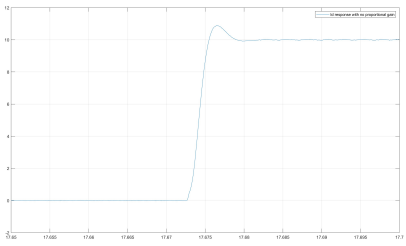


Figura 3.5: I only control when asked $I_d = 10A$

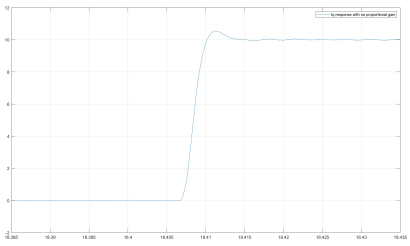


Figura 3.6: I only control when asked $I_q = 10A$

We can then perform tests with both parts, starting from the evaluation of P contribution by scaling it at different ratios, while keeping the computed k_i constant.

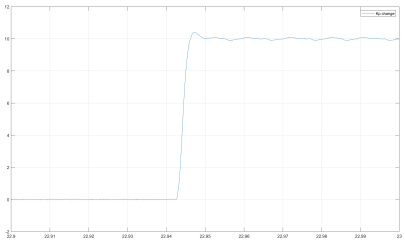


Figura 3.7: PI control with P part slowed down by 100 times

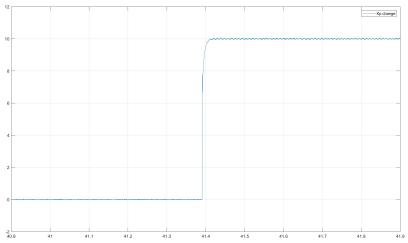


Figura 3.8: PI control with nominal values

Proportional gain influences both the rising time during transient and the oscillations at steady state, nonetheless the general shape of the response. When this factor is damped, rising time $90\% \approx 30ms$, whereas with the nominal value for the proportional gain it dropped to only $10ms$. This influence on the speed of the system dynamics is exactly what we would expect: the bigger k_p, the faster the controller.

Moreover, it is clear that at steady state k_p also influences the amplitude of the oscillations, but their frequency does not change. Oscillations happen around the current setpoint, and this gap is closed faster by the proportional side.

The last try to obtain the maximum performance out of our controller is trying to impose a proportional gain larger than the theoretical one. Here we fell into instability: a too large gain makes the command for V_{dq} skyrocket, such that we end up with a current alternatively pulled up and pulled down.

In light of this, it is better to stick to the nominal values.

For the integral side we performed the same experiments while focusing on I_q, just to change.

Having a slower integral part makes our control still effective, but with very large rising time. In Figure 3.12 we still recognize the presence of an instantaneous proportional

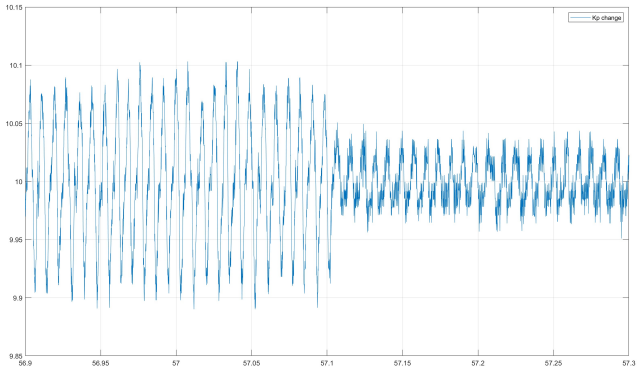


Figura 3.9: PI control oscillations close up at steady state when increasing k_p by one hundred times

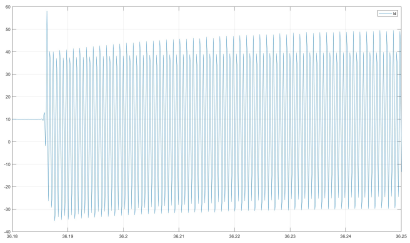


Figura 3.10: PI control with too large proportional gain

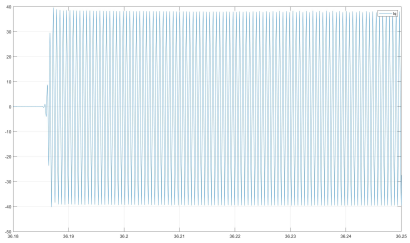


Figura 3.11: PI control instability with too large proportional gain

response, and then the integral one takes over to complete the exponential growth till the desired value.

Just like before, we can add one more level to this control and try to speed it up.

A too fast integral contribution makes the tracking faster, but it also brings about oscillations. These oscillations then make the settling time larger despite having a lower rising time. Moreover, the oscillations are slowly damped over time, therefore they do not seem to actually affect the steady state response. Pushing further would also arise wind-up issues, which is not convenient.

To sum up, the just computed values are good to achieve high performances both in terms of quickness and stability in a broader way.

Last but not least, we need to evaluate the impact of the speed on our control, since so far current control was performed with a mechanically blocked rotor test.

As we can see in picture, velocity does not destroy the transient or the steady state response of the current controller in the tested range. The current is still tracked, but a noise appears in I_{dq}, whose amplitude and frequency depend on the rpm. This is expected,

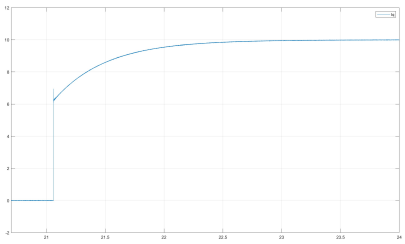


Figura 3.12: PI control with I part slowed down by 100 times

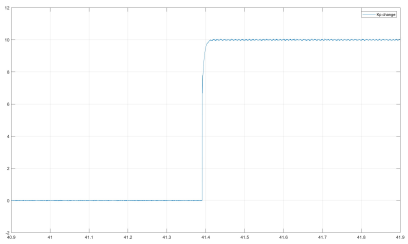


Figura 3.13: PI control with nominal values

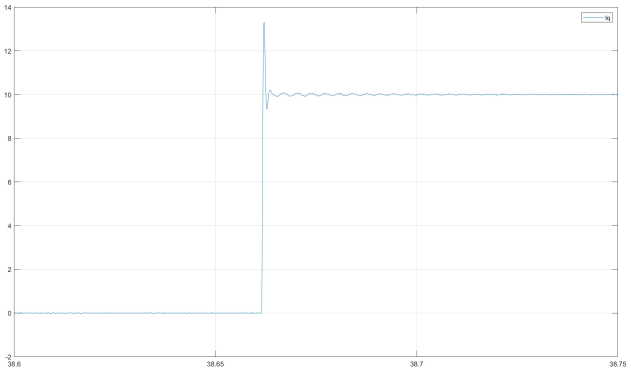


Figura 3.14: PI control with a faster integral part

because when the motor is spinning the back-EMF and the position measurement have a larger influence on the reconstructed dq quantities.

In the end, it is possible to evaluate the co-effect of all these inputs, how they affect each other.

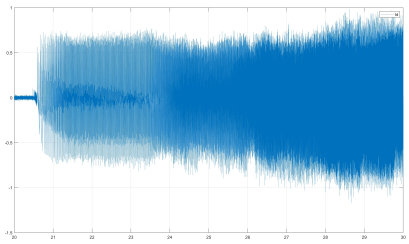


Figura 3.15: Steady state current control when speed is increased

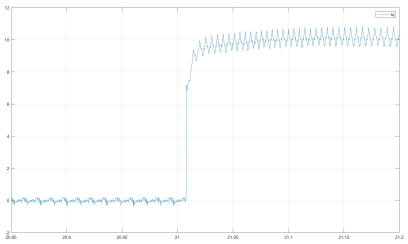
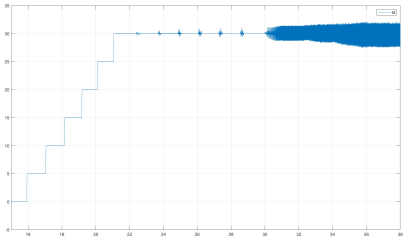
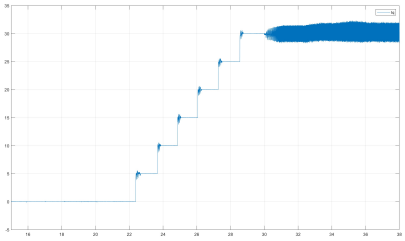


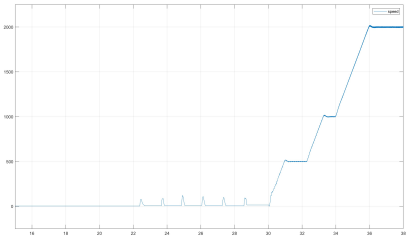
Figura 3.16: Current control at steady state high speed



(a) I_d



(b) I_q



(c) Speed

Figura 3.17: Current control with combined inputs of (a) I_d, (b) I_q, and (c) speed.

3.2 DC loading motor and CAN protocol

As already introduced in Section 2.5, the loading machine was initially controlled through P&D Studio 3 and a USB-CAN cable connected directly to the laptop. This configuration was useful during the first testing phase, because it allowed us to move the loading machine in a simple way and check if the mechanical connection with the machine under test was correct. However, it was not suitable for the final test-bench configuration. The goal was to control both machines from the same real-time environment, namely ControlDesk and the MicroLabBox, without depending on an external software running on the laptop.

For this reason, CAN communication was implemented between the dSPACE MicroLabBox and the PD6 loading motor. In this way, the loading machine could be commanded directly from the Simulink model running in real time. This was necessary because the current control of the machine under test and the velocity control of the loading machine had to work together during the same experiment. The machine under test was controlled electrically through the inverter, while the loading machine imposed the mechanical speed through its internal controller.

At first, this part was not very intuitive, because sending a velocity command to the loading machine was not equivalent to simply writing one value on the CAN bus. The PD6 had its own internal state machine, and it had to be brought to the correct state before accepting the speed reference. Also, each brushless DC motor might share the same "language" but not the same "words", that is to say that the firmware version decides what is implemented and what is not; for this reason, controlling the PD6 motor required an initial study of its object dictionary.

3.2.1 Brief introduction to CAN and CANopen

CAN, which stands for Controller Area Network, is a communication protocol largely used in automotive and industrial applications. It is based on a two-wire differential bus, usually called CAN-H and CAN-L. The use of two wires makes the communication more robust with respect to disturbances, because the useful information is not represented by the absolute voltage of one wire, but by the voltage difference between the two lines.

One important aspect of CAN is that it is not based on the direct addressing of a node. In other words, a message is not sent to a specific receiver in the same way as in other communication protocols. Instead, every node connected to the bus receives the message, and then decides if the message is useful or not. The identifier of the CAN frame defines the meaning of the message and also its priority. Lower identifiers have higher priority [Rapelli, 2024a].

This is also the reason why CAN uses a non-destructive arbitration system. If two nodes start transmitting at the same time, both nodes keep reading the bus while transmitting. Since the dominant bit overwrites the recessive one, the node sending the message with lower priority detects that the bus is not following its transmitted bit and stops transmitting. The higher priority message is therefore not destroyed, and continues on the bus. This is very useful in real-time systems, because important messages can be sent first without completely corrupting the communication [Rapelli, 2024b].

CAN itself defines the lower communication layers, so basically how the message is transmitted on the bus. However, for controlling a motor, we also need a higher level structure that gives a meaning to the data. This is where CANopen is used. CANopen is an application layer built on top of CAN. It defines standard communication objects and standard device behaviour. In our case, this was useful because the PD6 motor supports CANopen and follows the CiA 402 profile for drives and motion control [CiA, 2011, 2024].

In CANopen, the main idea is the object dictionary. Each parameter or variable of the drive is stored in an object, identified by an index and a subindex. For example, the controlword is object $6040h$, while the modes of operation object is $6060h$. By writing or

reading these objects, it is possible to configure the drive, enable it, select the operating mode and send the speed command.

The PD6 motor behaves as a CANopen slave, while the MicroLabBox behaves as the master in our setup. The master has to bring the node into the correct communication state and then send the correct objects to the drive. CANopen uses different services for this purpose. The most relevant ones for this work are:

- NMT, used for network management, for example to move the drive to the Operational state.
- SDO, used to write or read objects from the object dictionary.
- PDO, used for faster process data exchange during operation.
- Heartbeat, used by the node to signal that it is alive on the network.

In the Simulink implementation, the messages were built by separating a header and a data value. The header contained the CANopen command, the index and the subindex of the object, while the data value contained the value to be written. Even if in the model the signals were called `PDO_header` and `PDO_data`, the logic behind them is the access to the CANopen object dictionary of the PD6 motor.

3.2.2 Objects used for the loading machine

The loading machine was used mainly in velocity mode. This means that the PD6 internal drive controller was responsible for imposing the requested speed, while the machine under test was controlled in current through the inverter. In this way, the loading machine was not directly controlled by our current controller, but it was used to create a repeatable mechanical condition for the machine under test.

The main CANopen objects used in the implementation are reported below [Nanotec, 2025, CiA, 2024].

Object	Name	Use in the implementation
$6040h$	Controlword	Used to move the drive through the CiA 402 power state machine. The values $0006h$, $0007h$ and $000Fh$ were used to shutdown, switch on and enable operation.
$6060h$	Modes of operation	Used to select the velocity mode of the PD6 motor. In our case the value 2 was written.
$6042h$	Target velocity	Used to impose the speed reference to the loading machine. This value was updated during the test according to the desired speed.
$6048h$	Velocity acceleration	Used to define the acceleration ramp of the velocity controller.
$6049h$	Velocity deceleration	Used to define the deceleration ramp of the velocity controller, when required.
$6041h$	Statusword	Used as feedback from the drive to check the state of the controller, even if the main implementation was based on a predefined sequence.

Tabella 3.1: Main CANopen objects used for the control of the PD6 loading machine.

The controlword is the most important object during the start-up phase. The drive cannot directly receive a speed command and start moving. First, it has to pass through the correct internal states. For this reason, the start-up sequence was realized by sending three main controlword values. First, $0006h$ is sent to perform the shutdown command and bring the drive in a ready-to-switch-on condition. Then, $0007h$ is sent to switch on the drive. Finally, $000Fh$ is sent to enable the operation. Only after this sequence, the motor can actually react to the velocity command.

The mode of operation was selected by writing the value 2 in object $6060h$. This corresponds to velocity mode for the PD6. Once this mode is selected, object $6042h$ can be used to send the target velocity. The acceleration ramp was also set, in order to avoid applying an abrupt speed change to the mechanical system. This was important because the two machines were connected through the shaft and the torquemeter, and sudden speed changes could create useless mechanical stress or oscillations.

3.2.3 Implementation in Simulink and Stateflow

The CANopen communication was implemented directly inside the Simulink model executed by the MicroLabBox. The main difficulty was not only sending one CAN message, but coordinating the loading machine with the electrical control of the machine under test. For this reason, the implementation was made through a Stateflow chart.

The Stateflow chart manages two actions at the same time. On one side, it sends the CANopen messages to the PD6 motor. On the other side, it also controls the state of the machine under test controller. This was necessary because the loading machine should not start spinning before the machine under test has been correctly aligned and before the current control is ready.

The first part of the chart is used to activate the communication. At the beginning, the system remains in a silent state, with the operational trigger equal to zero. After a fixed time, the trigger is activated. This allows the CANopen communication to move the PD6 motor towards the Operational state. This delay was introduced to avoid sending messages immediately at the beginning of the real-time execution, when all the signals and blocks may not yet be in a stable condition.

After this first activation, the chart starts the sequence of CANopen messages. The first object written is the controlword $6040h$. The first value sent is:

$$6040h = 0006h$$

which corresponds to the shutdown command. After that, the value:

$$6040h = 0007h$$

is sent to switch on the drive. Then:

$$6040h = 000Fh$$

is sent to enable operation. These values follow the standard CiA 402 power state machine. In the Simulink implementation, each command was sent after a small time delay, so that the drive had enough time to move from one state to the next one.

In parallel, the machine under test follows its own control sequence. First, the alignment is performed. This is necessary because the current control is based on the Park transformation, and therefore on the knowledge of the rotor position. If the alignment is wrong, the d and q axes are not correctly oriented and the current control would not have the expected physical meaning.

After the alignment, the current controller is activated with:

$$I_{dq} = [0 \quad 0]$$

This step is important because it allows the current control to be already closed before the loading machine starts spinning. In this way, the machine under test is not producing torque, but the controller is ready and the measured currents are already being regulated around zero. This avoids starting the mechanical test while the electrical controller is still in a transient or uninitialized condition.

Only after this phase, the velocity mode of the PD6 is selected by writing:

$$6060h = 2$$

At this point, the PD6 internal controller is configured to work in velocity mode. Then the ramp parameters are accessed and the acceleration value is set. In the implemented sequence, object $6048h$ was used for the velocity acceleration ramp. The goal was to avoid a step-like speed command and to make the loading machine reach the desired velocity in a controlled way.

Finally, the target velocity is sent to the loading machine through object:

$$6042h$$

The value written in this object is the desired speed. In the Stateflow chart, this message is not sent only once, but is refreshed cyclically. This makes the speed command adjustable during the experiment. Therefore, by changing the speed variable in ControlDesk, the target velocity of the PD6 can be changed while the test is running.

The full procedure can therefore be summarized as follows:

1. The real-time application starts and the CAN communication remains silent for a short initial time.
2. The operational trigger is activated and the PD6 is moved to the CANopen Operational state.
3. The machine under test performs the alignment procedure.
4. The current controller of the machine under test is activated with $I_d = 0$ and $I_q = 0$.
5. The PD6 drive is enabled through the controlword sequence $0006h$, $0007h$, $000Fh$.
6. Velocity mode is selected on the PD6 by writing $6060h = 2$.
7. The ramp parameters are set, so that the loading machine does not receive an abrupt speed command.
8. The desired speed is sent through object $6042h$.
9. Once the speed condition is reached, the current control of the machine under test can be tested with the desired I_{dq} values.

This order is quite important. If the loading machine starts spinning before the alignment, then the electrical angle used by the controller may not be initialized correctly. If the desired current is applied before the speed control is ready, the mechanical condition is not controlled. Instead, by first aligning the machine, then closing the current loop at zero current, then activating the velocity control of the PD6, the two controls are started in a clean order.

The role of the loading machine in this test is therefore to impose the mechanical operating point. The role of the machine under test controller is instead to impose the electrical operating point. This separation is useful because it allows us to test the current controller at different speeds and with different current references. For example, once the PD6 is spinning at the desired speed, we can ask the machine under test to track a certain I_d and I_q. The torque response and the current tracking can then be evaluated under a controlled speed condition.

Another practical aspect is that all of this happens inside the same real-time task. This is a big improvement with respect to the first configuration, where the loading machine was controlled through P&D Studio 3 from the laptop. In that case, the loading machine and the machine under test were controlled from two different environments. With the CANopen implementation, the MicroLabBox became responsible for both the electrical control of the machine under test and the command of the loading machine. This made the workflow more repeatable and easier to automate.

The implementation was also useful for future tests. Once the CANopen communication was working, it became possible to create automatic tests where the speed of the loading machine and the current references of the machine under test are changed according to a predefined sequence. This is exactly the kind of workflow needed for motor characterization, because the machine has to be tested in many operating points. Without CANopen integration, every speed change would have required a manual action on the external software, making the test slower and less repeatable.

To conclude, the CANopen implementation was not only a communication exercise, but a necessary step to make the test-bench more integrated. The PD6 loading motor could be controlled directly by the MicroLabBox, while the machine under test was controlled through the inverter current loop. The Stateflow logic ensured that the two actions happened in the correct order: first alignment, then zero-current control, then velocity control of the loading machine, and finally current control with the desired I_{dq} values. This made the test-bench ready for systematic current-control tests and, later, for torque and flux characterization.

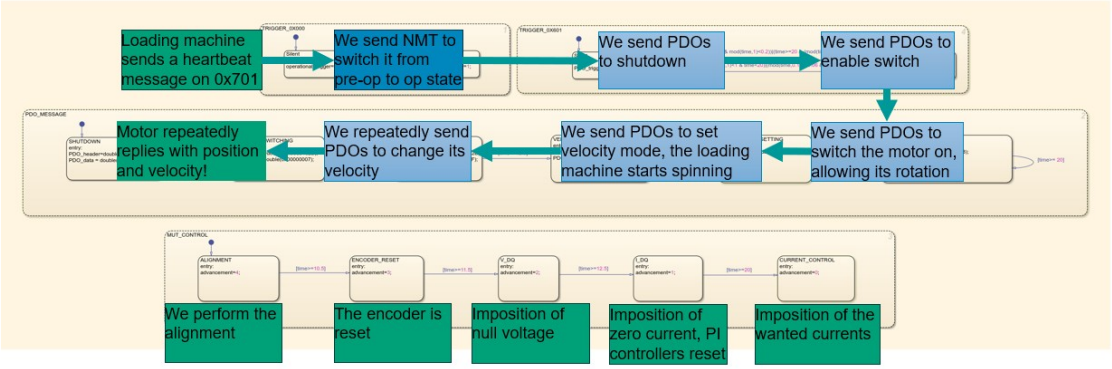


Figura 3.18: Overview of the implemented CANopen sequence. The upper Stateflow chart manages the communication with the PD6 loading machine through CANopen, while the lower chart controls the electrical states of the Machine Under Test. The two state machines are synchronized so that the loading machine starts rotating only after the alignment and initialization of the current controller have been completed.

3.3 Braking chopper

After implementing the communication with the loading machine and the current control of the machine under test, another important problem had to be faced. The test-bench was not only working in motoring conditions but also in generating one. Depending on the sign of torque and speed, both the machine under test and the loading machine could also work as generators. This is a normal behaviour for an electric machine, but it becomes a problem if the power supply is not able to absorb current.

In fact, whatever the applied control strategy is, an electrical machine can work both in motor mode and in generator mode. In motor mode, electrical power is taken from the DC supply, passes through the inverter, and is converted into mechanical power. In generator mode, the opposite happens: mechanical power is converted into electrical power and injected back into the DC link.

This can be summarized in a simple way by looking at the DC power:

$$P_{dc} = V_{dc} I_{dc}$$

If the current has the standard positive sign, the power supply is providing power to the system. Instead, if the DC current becomes negative, the system is trying to send power back to the power supply. This is the regenerative condition.

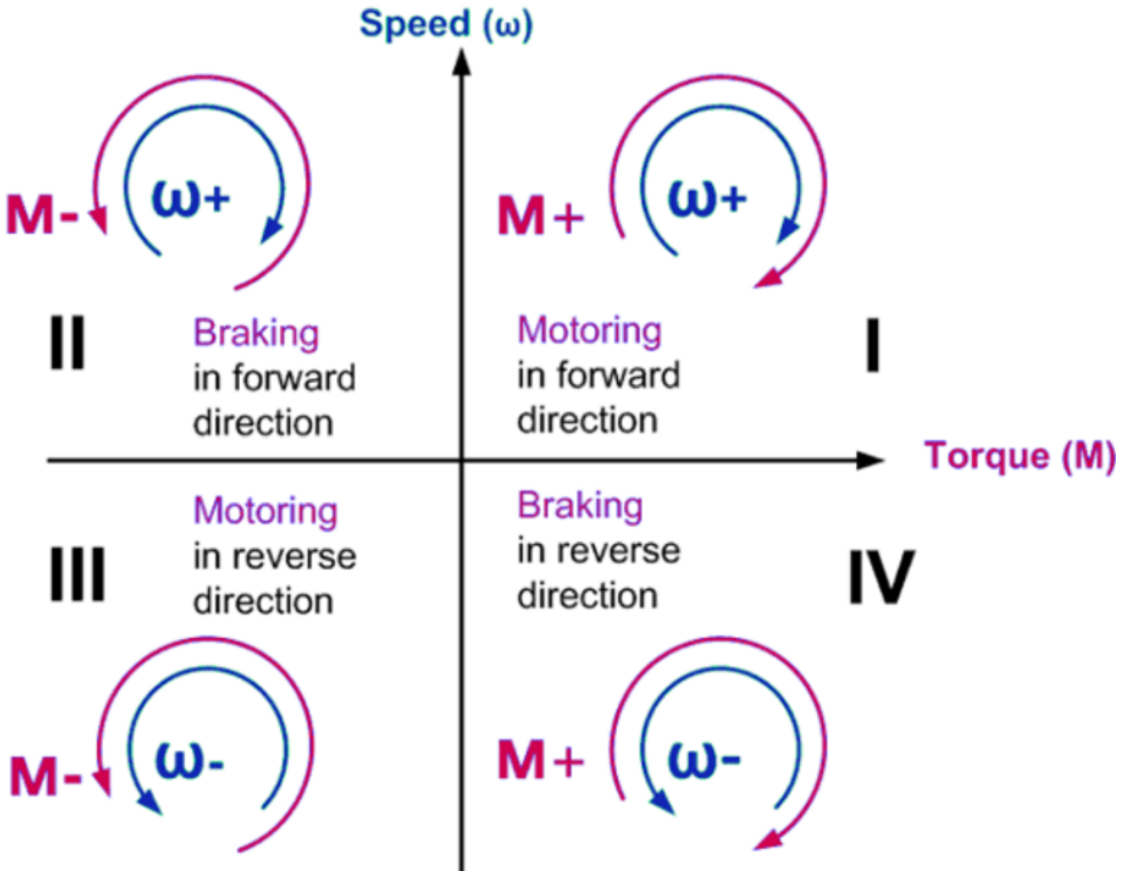


Figura 3.19: Motoring and generating modes

The problem is that our DC power supply does not support four-quadrant operation [Twintex, 2022]. In simple words, it can provide current, but it cannot properly absorb negative current. Therefore, when the machine enters generator mode, the regenerated energy cannot be sent back into the source. If no protection is added, this energy charges

the DC link capacitor and the DC bus voltage can increase. This is dangerous both for the power supply and for the inverter.

For this reason, a braking chopper was implemented. The idea of a braking chopper is simple: when regenerated energy cannot be absorbed by the supply, it is dissipated on a resistor. The resistor converts the electrical energy into heat, avoiding the increase of the DC link voltage and protecting the power supply.

In our test-bench the situation was slightly more delicate because the same DC supply powers both the machine under test and the loading machine. Therefore, it is not sufficient to look only at the current of one machine. One machine could be generating while the other one is consuming power. In that case, part of the regenerated power can be directly used by the other machine, without going back to the power supply.

For this reason, the criterion for activating the braking chopper was based on the sum of the DC currents of the two machines:

$$I_{tot} = I_{MUT} + I_{LM}$$

where I_{MUT} is the DC current related to the machine under test inverter and I_{LM} is the current related to the loading machine. If this total current becomes negative, it means that the global test-bench is injecting current back into the power source. In that condition, the braking chopper has to be activated.

So the logic is:

$$I_{tot} < 0 \quad \Rightarrow \quad \text{braking chopper active}$$

while, if the total current is positive, the power supply is still delivering energy to the system and the braking resistor is not needed.

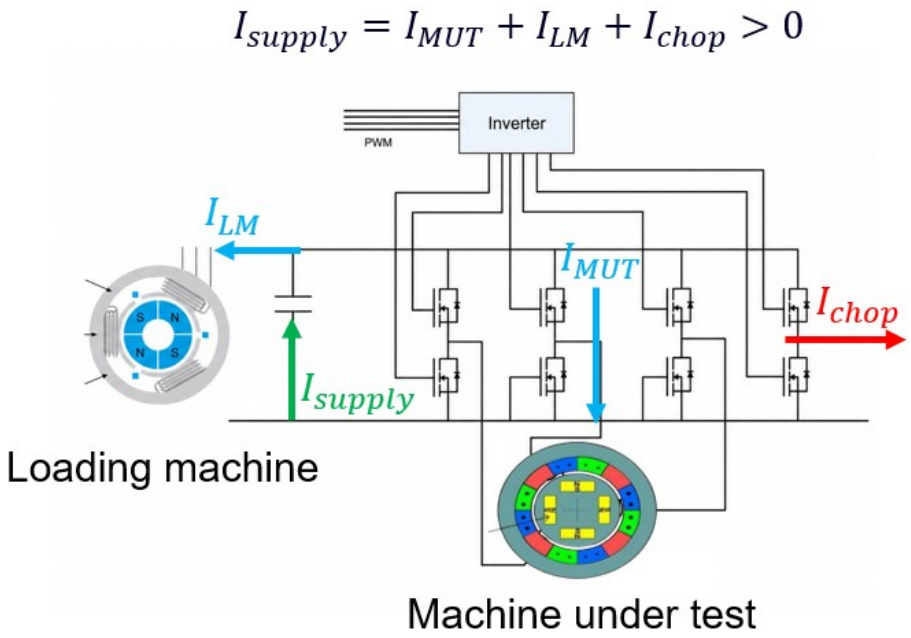


Figura 3.20: Braking chopper integration in the final test-bench configuration

A practical advantage of our hardware was the possibility to reuse the fourth leg of the inverter as the switch of the braking chopper. The inverter was not only used to control the three phases of the machine under test, but one additional leg was available. This leg was connected in such a way to command the current through the braking resistor. Therefore, instead of adding a completely external switching circuit, we exploited a part of the existing power electronics.

When the braking chopper is not active, the resistor is disconnected and no current flows through it. When the condition for activation is reached, the fourth inverter leg is switched and the resistor is connected to the DC link. In this way, the regenerated power is dissipated on the resistor. The resistor must obviously be chosen with a proper power rating, because all the dissipated energy becomes heat.

The implementation of the braking chopper was also linked to the addition of DC current measurements. Without these measurements, the controller would not know if the global system is absorbing or injecting current into the supply. For this reason, current measurements on the DC side were added to the test-bench. These measurements allowed us to check the sign of the current and to decide when the braking chopper had to intervene.

This point is important because, during some tests, the sign of the current can change quite naturally. For example, if the loading machine imposes a speed and the machine under test produces a braking torque, the machine under test can enter generator mode. In the same way, depending on the torque balance, the loading machine can also exchange

power with the DC link. Therefore, the braking chopper was not implemented for one single test only, but as a general protection needed for the whole experimental activity.

The final logic was therefore integrated in the real-time application. The MicroLabBox measures the useful currents, computes the condition for the braking chopper activation, and then commands the additional inverter leg. This made the protection automatic. The operator does not have to manually switch the braking resistor on or off during the test. If the system tries to inject current back into the supply, the chopper is activated and the energy is dissipated.

In the final setup, the braking chopper is one of the elements that allowed the test-bench to move from a simple open-loop validation stage to a more complete experimental platform. Before this implementation, regenerative operation was a limitation, because the DC supply could not accept negative current. After the implementation, tests involving braking torque and generator operation could be performed with a safer behaviour of the DC link.

Together with the braking chopper, a software current fuse was also considered. This was another safety layer, but based on current limitation rather than energy dissipation. Electrical components heat up when current flows through them, and the thermal stress is proportional to:

$$P_J = RI^2$$

where R is the resistance and I is the current. This means that current peaks are dangerous not only because of their instantaneous value, but because they can rapidly increase the thermal stress on the components.

The weakest element in the current path defines the real admissible current of the system. Even if the machine under test has a high rated current, the inverter, sensors and other components may have lower limits. Therefore, the current fuse was introduced to stop the test if the current exceeds a critical value for a certain time. In our case, the limit was set considering the bottleneck of the system and the admissible current duration.

The braking chopper and the current fuse have different roles: the braking chopper protects the DC supply and the DC link when power is regenerated, the current fuse protects the electrical components when the phase currents become too high, but no dissipation is involved in this case. Both were necessary because the test-bench was not an ideal simulation model anymore, but a real hardware setup where wrong commands or unexpected operating conditions could damage components.

To conclude, the braking chopper was implemented because the test-bench had to operate also in generating conditions, while the DC supply could not absorb negative current. Since both the machine under test and the loading machine are connected to the same DC supply, the activation condition was based on the total DC current of the system. The fourth leg of the inverter was reused as the switch for the braking resistor, making the solution practical and well integrated with the existing hardware. This implementation allowed safer four-quadrant-like operation of the test-bench, even if the power supply itself was not a four-quadrant device.

3.4 Experimental torque and flux maps

3.4.1 Aim of the characterization

After the implementation of the current control and after the integration of the loading machine through CANopen, the test-bench could finally be used for a more systematic characterization of the machine under test. Up to this point, the main objective was to make the setup work correctly: phase convention, encoder reliability, current control, loading machine speed control and safety systems. Once these parts were sufficiently stable, the next step was to use the test-bench to obtain experimental maps of the machine.

The quantities of interest were mainly torque and flux linkages as functions of the direct and quadrature currents. In other words, the objective was to obtain:

$$T = T(I_d, I_q)$$

$$\lambda_d = \lambda_d(I_d, I_q)$$

$$\lambda_q = \lambda_q(I_d, I_q)$$

These maps describe the real behaviour of the motor in the operating range and can be compared with the maps obtained from FEA analysis and used to update the simulation model [Armando et al., 2013, Lee et al., 2021].

The machine under test was controlled in current, while the loading machine imposed the desired mechanical speed. By sweeping different values of I_d and I_q, and waiting for the system to reach steady-state conditions, the machine could be characterized over the selected current region.

Torque could be measured directly through the torquemeter placed between the two machines. Flux linkages, instead, could not be measured directly and had to be reconstructed from electrical measurements. An initial estimation was performed directly from the steady-state PMSM voltage equations. The sensitivity of this approach to the stator resistance later motivated the adoption of a Motor–Generator–Motor measurement sequence for the final flux-map construction.

3.4.2 Experimental workflow

The characterization was performed by commanding the machine under test in current control. The current controller receives the desired values of I_d and I_q, and generates the voltage references needed to track them. The loading machine, instead, is controlled in velocity mode and imposes the desired speed. This separation was useful because the machine under test could be forced to operate in a selected electrical point, while the loading machine fixed the mechanical operating condition.

The general workflow used for the acquisition was the following:

1. The system was initialized and the alignment procedure was performed.

2. The current control was activated with zero current references (expected zero torque).
3. The loading machine was enabled and commanded in velocity mode.
4. The desired speed was reached and stabilized.
5. The required current operating point was imposed to the machine under test.
6. After the transient, the steady-state values of currents, voltages, speed and torque were recorded.
7. The procedure was repeated over the selected current grid.

The acquisition was therefore based on a grid of current values. The use of current control was fundamental, because without it we could not directly choose the operating point of the machine. In open-loop voltage control, the currents depend on the machine speed and impedance, so they cannot be imposed precisely. In closed-loop current control, instead, the currents are the controlled variables, therefore the test can be repeated in a much more systematic way.

For each steady-state condition, the first part of the signal was not used, because it includes the transient response of the current controllers and sometimes also a small mechanical transient. After a settling time, a time window was selected and the useful quantities were averaged. The specific sequence later used for the final flux maps is described in the corresponding subsection.

3.4.3 Torque map measurement

The torque was acquired from the torquemeter for every imposed current pair (I_d, I_q), after the operating point had reached steady state. The measured torque was then averaged over the selected acquisition window:

$$T_{dc}(I_d, I_q) = \frac{1}{t_2 - t_1} \int_{t_1}^{t_2} T(t)\, dt$$

where t_1 and t_2 define the steady-state acquisition interval.

A practical issue during the torque map acquisition was the limitation imposed by the loading machine. The machine under test has a larger torque capability than the loading machine. Therefore, not all current points could be tested freely, because some of them would require a torque that the loading machine could not safely withstand. For this reason, the admissible current region was not only limited by the current capability of the inverter or of the machine under test, but also by the mechanical torque limit of the loading machine.

For every tested point, the sign convention had to be checked carefully. A positive I_q should correspond to a positive torque according to the selected phase and angle convention. This was already verified during the previous tests, but it remains important during map generation, because a wrong sign convention would produce a map that is mathematically consistent but physically wrong.

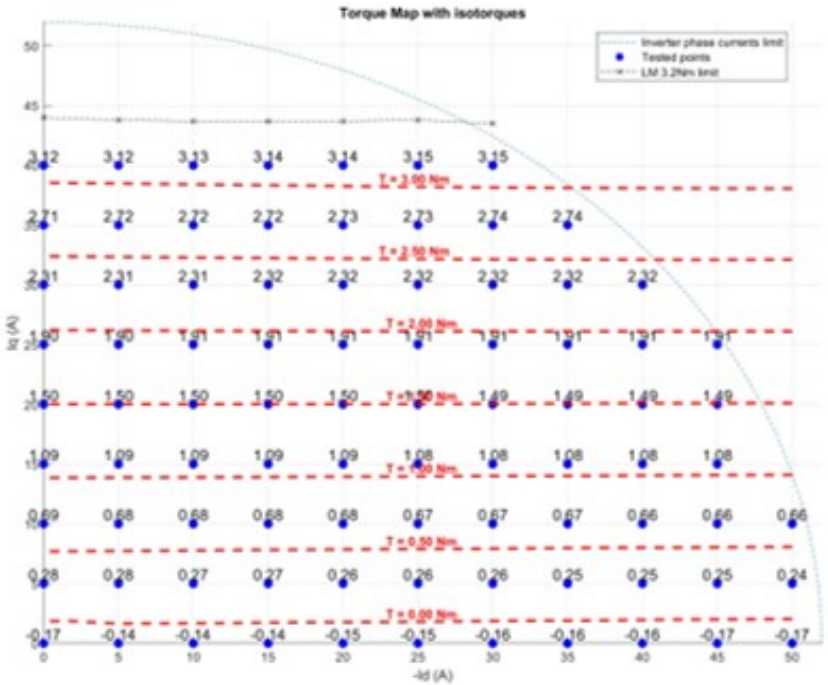


Figura 3.21: Experimental DC torque map obtained from the torquemeter measurement.

The experimental map confirms the behaviour expected for the machine under test: torque varies mainly with I_q, while I_d has a much smaller influence. This is coherent with the low-saliency behaviour previously discussed for a surface-mounted PMSM.

Since the map is built from averaged values, torque ripple components are not included. The limitations related to harmonic characterization are discussed later together with the flux maps.

3.4.4 Initial flux map estimation from voltage equations

As a first approach, we tried to trace back the flux linkages directly from the PMSM voltage equations in the rotating dq reference frame:

$$V_d = R_s I_d + \frac{d\lambda_d}{dt} - \omega_e \lambda_q$$

$$V_q = R_s I_q + \frac{d\lambda_q}{dt} + \omega_e \lambda_d$$

For every operating point, we waited for the transient to end before taking the useful measurements. In this condition, currents, voltages and speed could be considered approximately constant, therefore allowing us to neglect the flux derivatives:

$$\frac{d\lambda_d}{dt} \simeq 0$$

$$\frac{d\lambda_q}{dt} \simeq 0$$

The equations then become:

$$V_d = R_s I_d - \omega_e \lambda_q$$

$$V_q = R_s I_q + \omega_e \lambda_d$$

From these equations, the flux linkages can be obtained as [Liu et al., 2018, Lee et al., 2021]:

$$\lambda_d = \frac{V_q - R_s I_q}{\omega_e}$$

$$\lambda_q = \frac{R_s I_d - V_d}{\omega_e}$$

These equations were initially applied point by point to the acquired data. For every current pair (I_d, I_q), the corresponding averaged values of V_d, V_q, I_d, I_q and ω_e were used to estimate λ_d and λ_q.

The electrical speed was computed from the mechanical speed using:

$$\omega_e = p\omega_m$$

where p is the number of pole pairs. Since the flux estimation requires division by ω_e, the method cannot be applied at zero speed. For this reason, the machine had to be spun by the loading machine at a non-zero speed. At the same time, the speed should not be too high, otherwise the voltage requirement increases and the operating region becomes limited by the available DC link voltage. Therefore, an intermediate speed had to be selected.

The voltage values used in the equations were the reconstructed dq voltages. These voltages are not directly measured as sinusoidal phase voltages, because the machine is supplied through the inverter. The inverter generates PWM voltages, and the control model reconstructs the voltage components starting from the commanded duty cycles and the DC link voltage. This is already an approximation, because the real inverter has dead-time, voltage drops, switching delays and other non-idealities.

The first results were not as reliable as expected. Several approximations could affect the reconstructed flux, but the most relevant problem we experienced was related to the stator resistance.

3.4.5 Problem of the stator resistance

While testing this first estimation procedure, we noticed that the value of R_s was much more critical than initially expected. Since it directly appears in both flux equations, a wrong estimation of the resistance immediately affects the reconstructed flux.

This problem was particularly relevant for the machine under test, which is a small low-voltage, high-current motor. Its stator resistance is very low, while relatively large currents can flow through the windings. Consequently, even short acquisition periods can produce a noticeable temperature increase.

The stator resistance varies with temperature according to:

$$R_s(T) = R_s(T_0)[1 + \alpha(T - T_0)]$$

where T_0 is the reference temperature and α is the temperature coefficient of the winding material. Since the winding temperature was not measured continuously during the characterization, a single value of R_s could not accurately represent all the tested operating points. This effect becomes increasingly relevant at higher currents, where the temperature rise is faster.

Even though the resistance is only in the mΩ range, its effect could not be considered negligible. In fact, several amperes were flowing in the machine while the flux values were only in the order of few mWb. Therefore, the voltage drop R_sI could become significant when compared to the voltage contribution used to reconstruct the flux.

For example, an error ΔR_s produces approximately the following flux estimation errors:

$$\Delta\lambda_d \simeq -\frac{\Delta R_s I_q}{\omega_e}$$

$$\Delta\lambda_q \simeq \frac{\Delta R_s I_d}{\omega_e}$$

Therefore, even an error of only a few mΩ can become significant when multiplied by several amperes. Since the resulting voltage error is then divided by the electrical speed, the effect becomes even more critical at relatively low speeds. If the resistive contribution is estimated inaccurately, the resulting error can become comparable with the flux-related voltage term and can lead to large relative errors or even to a wrong sign in the reconstructed flux at the most critical operating points.

Another effect related to temperature is the variation of the permanent magnet flux. When the rotor magnets heat up, their flux generally decreases. This means that the value of the permanent magnet flux linkage λ_{pm} is also temperature-dependent. In our estimation, this effect was neglected because the rotor temperature was not measured directly and it was not possible to include a reliable thermal correction for the permanent magnets.

These results showed that a direct application of the steady-state voltage equations was not sufficiently robust for the final flux-map identification. In particular, the variation of the stator resistance during the characterization introduced an uncertainty that could not be reliably compensated by using a single constant value of R_s.

For this reason, the final flux maps were obtained using a Motor–Generator–Motor measurement sequence, described in the following subsections.

3.4.6 DC flux maps and limits of harmonic analysis

The characterization considered in this work was focused on the DC, or average, components of the flux linkages in the dq reference frame:

$$\lambda_{d,dc}(I_d, I_q)$$

$$\lambda_{q,dc}(I_d, I_q)$$

An electrical machine can also show harmonic components depending on slotting, magnet shape, winding distribution and inverter behaviour. In a PMSM, the sixth harmonic in the dq frame is particularly interesting because it is associated with flux and torque ripple.

However, this harmonic study was not performed in this work. The main reason is that the available measurement bandwidth was not sufficient for a reliable analysis. The inverter voltages are PWM signals with a switching frequency of 10 kHz. To study harmonics properly, and especially to reconstruct voltage-related quantities, the acquisition system would need a bandwidth and sampling frequency high enough to separate the useful components from the switching behaviour.

In practice, the measured or reconstructed voltages are affected by the PWM nature of the inverter. A proper harmonic study would require a measurement chain designed for this purpose, with high-bandwidth voltage measurements, synchronized acquisition and probably a different post-processing strategy. Otherwise, there is the risk of confusing real electromagnetic harmonics with switching noise, acquisition limitations or filtering artefacts.

For this reason, harmonic characterization, including the sixth harmonic torque component mentioned previously, was considered outside the reliable range of this setup. The maps obtained in this work were therefore limited to their average dq components.

3.4.7 Construction of the maps

Since the direct use of the steady-state voltage equations gave results that were too dependent on the stator resistance, for the final characterization we moved to a Motor–Generator–Motor (MGM) measurement sequence. This is an experimental technique already used for the identification of fundamental dq flux maps [Bojoi, Pescetto, Ferrari, and Pellegrino, 2026], and in our case it was particularly useful to avoid relying on one fixed value of R_s.

For each current operating point, three consecutive steady-state conditions were imposed. The first and third conditions, referred to as M1 and M2, correspond to motor operation at the same current reference:

$$\mathbf{I}_{dq}^{M1} = \mathbf{I}_{dq}^{M2} = \begin{bmatrix} I_d \\ I_q \end{bmatrix}$$

while the intermediate generator condition is obtained by changing the sign of the quadrature current:

$$\mathbf{I}_{dq}^{G} = \begin{bmatrix} I_d \\ -I_q \end{bmatrix}.$$

The three conditions therefore form the M1–G–M2 sequence. During the procedure, the loading machine keeps the mechanical speed approximately constant, while the machine under test remains under closed-loop current control.

The principle of the method is based on the symmetry between the motor and generator operating points. The resistive voltage contribution changes consistently with the current reversal, while the magnetic quantities of the symmetric operating points can be combined. By processing the measurements from the motor and generator conditions together, the contribution associated with the stator resistance can therefore be cancelled from the average flux estimation.

The presence of the second motor operating point, M2, is particularly useful when the stator resistance changes during the test because of winding heating. The two motor measurements are taken before and after the generator condition. If the resistance variation over the short M1–G–M2 sequence is sufficiently smooth, the average of the two motor conditions provides an estimate of the resistance-dependent contribution at the intermediate generator condition. This considerably reduces the influence of the resistance drift on the resulting flux estimate.

This aspect was particularly important for the machine under test because of the combination of very low stator resistance and relatively high operating currents. The MGM sequence therefore allowed the final flux maps to be constructed without relying on a single accurately known value of R_s for the complete characterization campaign.

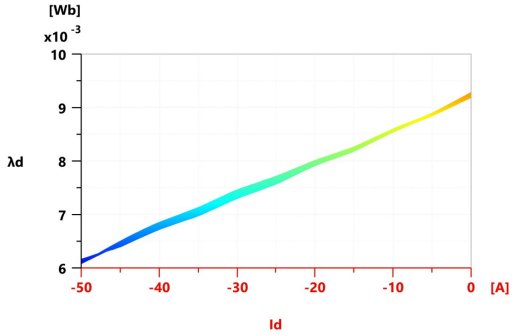
The procedure was repeated over the selected grid of current operating points. For every steady-state condition, the useful quantities were recorded after the transient had disappeared. The main stored quantities were:

- measured I_d and I_q;
- V_d and V_q;
- mechanical speed and electrical speed;
- measured torque;
- estimated λ_d and λ_q.

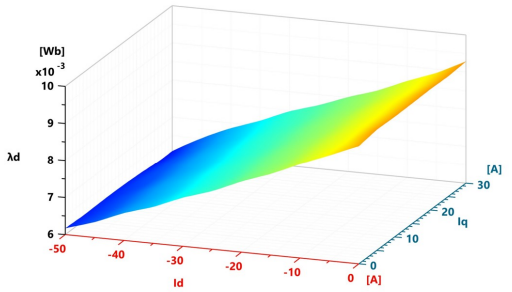
The use of measured currents instead of only reference currents is important. Even if the current control was working correctly, small differences between reference and measured values can still exist. For this reason, the measured values were checked to confirm that the requested operating point had actually been reached.

After repeating the procedure over the complete admissible current grid, the resulting values were organized as surfaces over the $I_d - I_q$ plane. The torque map was obtained directly from the averaged torquemeter measurements, while the d- and q-axis flux-linkage maps were reconstructed from the MGM measurements.

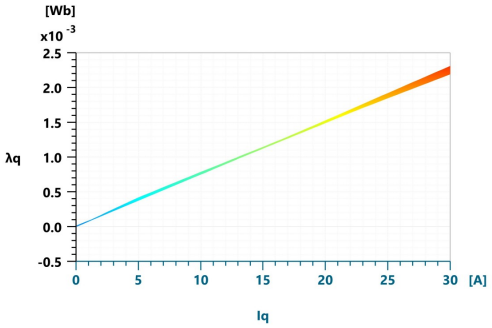
The final surfaces were then compared with the initial FEA maps as a consistency check. They were not expected to be identical, since the purpose of the experiment was precisely to obtain a representation based on the real machine, but their global trend had to remain physically coherent. Large unexpected discontinuities or sign inversions would have suggested problems in the acquisition, current convention or data processing.



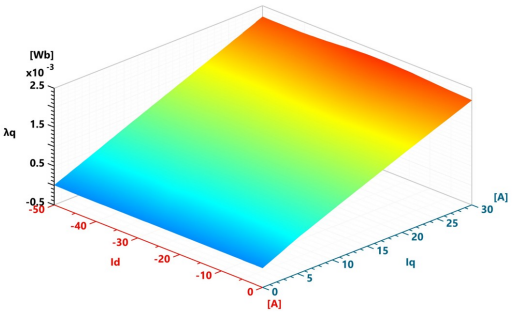
(a) Estimated DC d-axis flux linkage map over I_d



(b) Estimated DC d-axis flux linkage map overview



(c) Estimated DC q-axis flux linkage map over I_q



(d) Estimated DC q-axis flux linkage map overview

Figura 3.22: Estimated DC flux linkage maps: (a) d-axis flux linkage over I_d, (b) d-axis flux linkage overview, (c) q-axis flux linkage over I_q, and (d) q-axis flux linkage overview.

3.4.8 Use of the maps for model updating

The final aim of the characterization was to replace part of the approximate motor data initially used in the simulation with the experimentally obtained maps. In this way, the model could be updated using information coming directly from the real Machine Under Test.

This is especially important for torque control. In current control, the measured current is directly used as feedback, so the controller can correct the error between desired and actual current. In torque control, instead, the requested torque has to be converted into suitable current references using the machine characteristics. An inaccurate map can therefore lead to a torque error even when the current references are tracked correctly.

The flux maps provide additional information about the magnetic behaviour of the machine and can therefore be used together with the torque map for the model update. Such updated characteristics will later be useful when moving from current control to torque control.

Appendice A

Main Inverter Driver Schematic

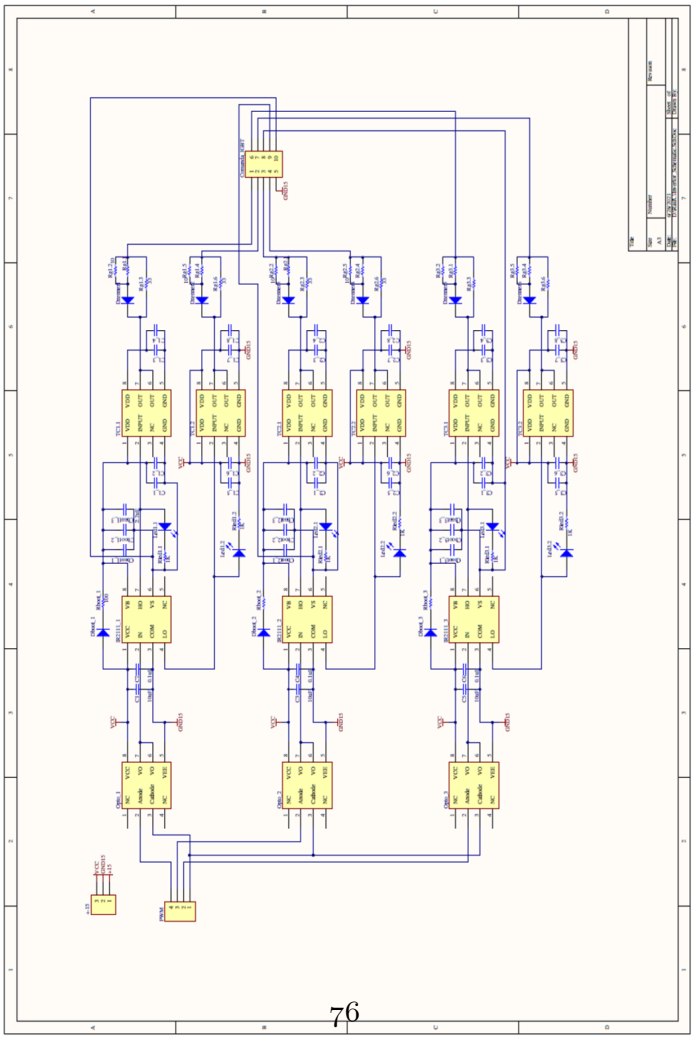


Figura A.1: Complete schematic of the main inverter driver.

Bibliografia

Analog Devices. Mt-095: Emi, rfi, and shielding concepts. Analog Devices Tutorial, 2009. URL https://www.analog.com/media/en/training-seminars/tutorials/MT-095.pdf.

Eric Giacomo Armando, Iustin Radu Bojoi, Paolo Guglielmi, Gianmario Pellegrino, and Michele Angelo Pastorelli. Experimental identification of the magnetic model of synchronous machines. *IEEE Transactions on Industry Applications*, 49(5):2116–2125, 2013. doi: 10.1109/TIA.2013.2258876. URL https://iris.polito.it/handle/11583/2507496.

Belden Inc. Technical information: Shielding. URL https://catalog.belden.com/prodimgs/AllAssets-approved/std.lang.all/58/75/175875.pdf. Shield types, applications and performance.

Andrei Bojoi, Paolo Pescetto, Simone Ferrari, and Gianmario Pellegrino. Torque ripple map of synchronous machines via manipulation of experimental $dq\theta$ flux maps: Implementation and validation, 2026. Politecnico di Torino.

Radu Bojoi and Paolo Pescetto. Lecture 6 – fundamental concepts of the control of power converters and edrives, 2025. Lecture notes, 01TVEQW – Electrical Technologies for eMobility, Politecnico di Torino, Dipartimento Energia "G. Ferraris".

burster. *High-Precision Torque Sensor Model 8661 – Technical Data and Operation Manual*. URL https://www.burster.com/products/sensors/torque-sensors/rotating-applications/8661.

CiA. *CiA 301: CANopen Application Layer and Communication Profile*. CAN in Automation (CiA), version 4.2.0 edition, 2011. URL https://www.can-cia.org/can-knowledge/canopen.

CiA. *CiA 402 Series: CANopen Device Profile for Drives and Motion Control*. CAN in Automation (CiA), version 5.0.0 edition, 2024. URL https://www.can-cia.org/can-knowledge/cia-402-series-canopen-device-profile-for-drives-and-motion-control.

dSPACE. *DS1202 MicroLabBox Platform – User Documentation*. dSPACE GmbH, 2023. URL https://www.dspace.com/en/pub/home/

support/documentation.cfm?Language=en-us&Release=RLS2023-A&externalid=Topic_1eda5995-694e-492e-91cd-cd9c991fe8b7_--_&helpsetid=ModelDeskPlatformManagement. Revision May 2023.

D. Grahame Holmes and Thomas A. Lipo. *Pulse Width Modulation for Power Converters: Principles and Practice*. Wiley-IEEE Press, 2003.

Francesco Iannuzzo. Lecture 3c – power electronic devices, 2025. Lecture notes, 01TVEQW – Electrical Technologies for eMobility, Politecnico di Torino, Dipartimento Energia "G. Ferraris".

Ramu Krishnan. *Permanent Magnet Synchronous and Brushless DC Motor Drives*. CRC Press, 2010.

Joohyun Lee, Yong-Cheol Kwon, and Seung-Ki Sul. Identification of ipmsm flux-linkage map for high-accuracy simulation of ipmsm drives. *IEEE Transactions on Power Electronics*, 36(12):14257–14266, 2021. doi: 10.1109/TPEL.2021.3084558. URL https://ieeexplore.ieee.org/document/9442885.

Current Transducer LA 125-P/SP4 – Datasheet. LEM, 2013. URL https://www.lem.com/sites/default/files/products_datasheets/la%20125-p%20sp4.pdf. Version 15, 25 April 2013.

Kan Liu, Jianghua Feng, Shuying Guo, Lei Xiao, and Zi-Qiang Zhu. Identification of flux linkage map of permanent magnet synchronous machines under uncertain circuit resistance and inverter nonlinearity. *IEEE Transactions on Industrial Informatics*, 14 (2):556–568, 2018. doi: 10.1109/TII.2017.2722470. URL https://ieeexplore.ieee.org/document/7967682.

Murata Manufacturing Co., Ltd. Noise suppression measures at dc power cable connection. URL https://www.murata.com/en-eu/products/emc/emifil/library/knowhow/basic/s2-chapter07-p1. Guidelines for EMI Suppression (EMC).

Nanotec. *PD6-E/-EB CANopen Technical Manual*. Nanotec Electronic GmbH & Co. KG, version 1.1.1 edition, 2025. URL https://www.nanotec.com/fileadmin/files/Handbuecher/Plug_Drive/PD6-E/fir-v2213/PD6E_CANopen_Technical-Manual_V1.1.1.pdf. Valid for firmware version FIR-v2213.

Raul-Octavian Nemeş. *Development and Performance Analysis Platform for Battery Electric Vehicles Power Trains*. PhD thesis, Technical University of Cluj-Napoca, Faculty of Electrical Engineering, 2021.

PeakTech. *PeakTech P 1580 Laboratory DC Power Supply – Operation Manual*. PeakTech Prüf- und Messtechnik GmbH, 2024. URL https://www.peaktech.de/uk/PeakTech-P-1580-Laboratory-power-supply-DC-1-32V-0-30A-USB/P-1580.

Marco Rapelli. Controller area network (can), part i: The basics, 2024a. Lecture notes, Networking Technologies for Connected Vehicles, Politecnico di Torino, Dipartimento di Automatica e Informatica.

Marco Rapelli. Working with controller area network, 2024b. Lecture notes, Networking Technologies for Connected Vehicles, Politecnico di Torino, Dipartimento di Automatica e Informatica.

Seung-Ki Sul. *Control of Electric Machine Drive Systems*. Wiley-IEEE Press, 2011.

Twintex. *PCL Series High Power DC Power Supply – PCL9000-60*. Twintex Electronics Co., Ltd., 2022. URL https://www.twintex-power.com/programmable-high-power-dc-power-supply-9kw-series/.