

Article

A Fault-Tolerant Finite-Control-Set MPC Architecture with Asymmetry-Aware Thermal Balancing for Switched Reluctance Motor Drives

Franklin Sánchez ^{1,2,*} , María Isabel Milanés-Montero ²  and Enrique Romero-Cadaval ² 

¹ Departamento de Eléctrica, Electrónica y Telecomunicaciones, Universidad de las Fuerzas Armadas ESPE, Av. General Rumiñahui S/N, Quito 171103, Ecuador

² Power Electrical, Electronic and Automation Engineering Department, School of Industrial Engineering, University of Extremadura, Avda. de Elvas s/n, 06011 Badajoz, Spain; milanes@unex.es (M.I.M.-M.); eromero@unex.es (E.R.-C.)

* Correspondence: fisanchez@espe.edu.ec

Abstract

Switched reluctance motors (SRMs) are attractive for fault-tolerant drives because their rare-earth-free rotor and intrinsic phase isolation support continued operation after a converter fault. Realising this requires a post-fault control policy that preserves both torque tracking and per-phase thermal balance, with the latter being a safety-relevant design consideration motivated by—though not herein verified against—ISO 26262. This paper proposes and evaluates, by simulation, a three-layer fault-tolerant finite-control-set model predictive control (FCS-MPC) architecture for a four-phase 8/6 SRM under a single open-phase converter fault. The layers are (i) a vector-set reconfiguration from the eight healthy, active vectors to the twenty-six admissible post-fault vectors, which restores controllability of the reduced converter; (ii) soft commutation expressed as a position-dependent penalty inside the MPC cost; and (iii) asymmetry-aware balancing that evens out the accumulated thermal load across the three healthy phases. We additionally analyse an activated-on-demand max-penalty thermal limiter and show, both analytically and in simulation, that it shares its optimiser with the variance-based balancing term and therefore confers no measurable benefit over it; it is consequently retained only as an optional on-demand limiter rather than a separate layer. The architecture is benchmarked against a fault-blind baseline, a rule-based hard fault-tolerant reference (Hard-FT), and intermediate configurations through a deterministic ablation across three critical operating points, complemented by a robustness assessment under measurement noise and parameter mismatch. A six-criteria fault-tolerance scorecard is reported as a methodological observation on the transferability of healthy-mode SRM specifications to post-fault operation.



Academic Editors: Ahmed A. Zaki Diab and Hamdy M. Sultan

Received: 27 May 2026

Revised: 7 July 2026

Accepted: 16 July 2026

Published: 18 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: switched reluctance motor; finite-control-set model predictive control; fault tolerance; open-phase fault; thermally aware load balancing; Monte Carlo robustness; ISO 26262

1. Introduction

Electrification of road transport has renewed industrial and academic interest in electric machines that are robust, inexpensive, and free of rare-earth magnets. Switched reluctance motors (SRMs) meet these three constraints and have re-emerged as a competitive option for traction, auxiliary, and ancillary drives in electric vehicles (EVs) [1–4]. Beyond

their cost and material advantages, two structural properties make SRMs particularly suited to the fault-critical operation demanded by automotive safety standards such as ISO 26262 [5]: each stator phase is magnetically independent of the others, and the typical asymmetric H-bridge converter has two independent switches per phase, so a single switch failure does not disable the neighbouring phases. Fault-tolerant operation after the loss of one phase is therefore achievable in principle by re-routing the torque demand onto the remaining healthy phases [6–8]. For an EV drive, this structural property translates into a safety-relevant limp-home capability—*provided* the control policy running on the healthy phases preserves both torque production and per-phase thermal integrity.

In practice, the quality of the fault ride-through depends strongly on the control policy. A naive policy that continues to demand the pre-fault phase currents—the “fault-blind” baseline—leads to severe torque ripple, speed drop, and current spikes in the neighbouring phases because each healthy phase must cover a larger angular arc of the electromechanical torque profile. Classical fault-tolerant strategies address this through hard-switched compensation: torque reference re-allocation, commutation-angle shifts forced by rule-based logic, and current limits triggered by lookup tables [6–8]. We refer to this rule-based reference design as *Hard-FT* throughout the paper. Mutually coupled and multi-phase variants of the same principle have also been investigated in the literature [9].

Finite-control-set model predictive control (FCS-MPC) has become a popular alternative for SRM current control because it handles the non-linear flux–current–angle relationship natively through a discrete-time plant model and admits arbitrary cost weights [10–13]. Recent FCS-MPC contributions for SRMs have shown improvements in torque ripple, switching effort, and average-current tracking under healthy operation [14–17], including the authors’ optimized dual-phase excitation approach for the 8/6 SRM [18]; a broader review of predictive control for SRM drives is given in [19]. Extensions to fault-tolerant operation are fewer and more recent: the common pattern is to retain the healthy-mode vector set and cost function and add a separate rule-based fault-handling module on top [20,21]. This is functional but leaves three sources of performance loss on the table: (i) the vector set, itself, is wrong after the fault because the admissible post-fault set has 26 non-trivial states rather than 8; (ii) hard turn-on/turn-off angle overrides introduce discontinuous cost-function behaviour that defeats MPC smoothness; and (iii) even with a correct vector set, the three healthy phases do not naturally distribute the missing torque uniformly, so per-phase thermal integrity is not enforced. Beyond SRMs, open-phase fault-tolerant control has advanced rapidly for multi-phase PMSM drives, where recent methods have optimized copper loss and phase-current trajectories under open-phase operation [22,23]; those approaches exploit the sinusoidal multi-phase structure and do not transfer directly to the doubly salient SRM addressed here. Comprehensive treatments of fault-tolerant SRM drives are collected in recent reviews [24]. Other SRM fault modes—notably, stator inter-turn short circuits—are handled by dedicated diagnosis and fault-tolerant methods [25–27], which are complementary to the single open-phase converter fault considered here.

Contributions.

We propose and evaluate, by simulation, a three-layer fault-tolerant FCS-MPC architecture for a four-phase 8/6 SRM under a single open-phase fault, assessed with a 30-seed Monte Carlo simulation under measurement noise and parameter mismatch. The architecture combines (i) *vector-set reconfiguration* with the 26 admissible post-fault vectors, (ii) *soft commutation* via a position-dependent cost-function penalty (replacing rule-based overrides), and (iii) *asymmetry balance* penalizing the variance of RMS currents across the three healthy phases. We additionally analyse an optional max-penalty thermal limiter and show it to be equivalent to layer (iii). We report two intermediate configurations as reference points: *L2-MPC*, which exposes only layers (i) and (ii), and *L3-MPC*, which

adds layer (iii); the limiter is evaluated as *L4-MPC*. The paper's specific contributions are summarized as follows:

- **We introduce a three-layer fault-tolerant FCS-MPC architecture** that restores post-fault speed regulation under a single open-phase fault, which is shown to be robust to measurement noise and $\pm 10\%$ parameter mismatch across three critical operating points, with the largest gains at high speed, where a fault-blind controller loses regulation entirely.
- **An equivalence analysis of the max-penalty limiter is presented.** We show both structurally and empirically that penalizing the *maximum* healthy-phase RMS current against a fixed threshold and penalizing the *dispersion* of currents both require a balanced-state minimizer, so the two coincide wherever the load can be balanced. A weight sweep and a 30-seed Monte Carlo simulation confirm that the limiter adds no measurable benefit at the studied threshold ($p = 1.0$ on all twelve paired comparisons); it is therefore reported as an analysed equivalent rather than a separate layer.
- **We report an ablation that isolates the contribution of each architectural layer** at three operating points under nominal conditions, establishing that vector-set reconfiguration and soft commutation are the dominant contributors to tracking performance, while asymmetry balance adds a smaller but consistent improvement that is largest at low speed.
- **We conduct a paired-seed Monte Carlo evaluation** under measurement noise and parameter mismatch with Wilcoxon signed-rank tests [28] of the proposed controller against the fault-blind baseline and of the limiter variant against the balancing-only controller. The proposed-versus-blind speed-RMSE gain is significant at the mid- and high-speed operating points ($p < 10^{-8}$ at 700/2 and 800/2 Nm) and only marginal at 400/2 Nm; the limiter-versus-balancing comparison is exact ($p = 1.0$).
- **We report a methodological observation on the strict scorecard:** a six-criteria fault-tolerance scorecard, which is typical of healthy-mode SRM control, is not passed at any operating point by any controller under Monte Carlo simulation because post-fault torque ripple intrinsically exceeds the healthy-mode bound when a quarter of the excitation is lost. We interpret this as evidence that healthy-mode SRM specifications must be explicitly relaxed when transferred to post-fault operation—an observation relevant to ISO 26262-style certification of EV drives.

The remainder of the paper is organized as follows. Section 1.1 maps the three-layer architecture onto ISO 26262 safety concepts to make the certification relevance concrete. Section 2 describes the 8/6 SRM plant, the asymmetric bridge converter, the Le-Huy magnetization model, and the open-phase fault. Section 3 derives the three-layer fault-tolerant FCS-MPC architecture and its intermediate references (L2-MPC and L3-MPC), together with the analysed max-penalty limiter (L4-MPC). Section 4 describes the simulation methodology and fault-tolerance scorecard. Section 5 reports the results. Section 6 discusses implications. Section 7 concludes the paper.

1.1. ISO 26262 Context for Post-Fault SRM Drives

ISO 26262 governs the functional safety of electric and electronic systems in passenger vehicles [5]. Traction drives, on which vehicle propulsion and stability depend, are typically classified at Automotive Safety Integrity Level (ASIL) C or D, with the highest class assigned where unintended torque could lead to severe injury at high speed. We make explicit the distinction between the application *context* that motivates this work and the specific *machine* used to demonstrate it. The architecture, its cost-function formulation, and the ISO 26262 mapping developed below are framed for the EV-traction context because it is in traction that post-fault per-phase thermal integrity ceases to be a quality metric and becomes a

quantified technical safety requirement (TSR-2, Section 1.1). However, the laboratory-scale four-phase 8/6 prototype used as the reference machine throughout this paper is a demonstrator whose most immediate realistic deployment is in auxiliary EV functions (Heating, Ventilation and Air Conditioning (HVAC); compressors; oil/coolant pumps; and e-axis accessories), which fall in the ASIL B–C range. The layered reasoning is identical in both cases; what changes when moving to a higher-ASIL traction drive is the tightness of the quantitative budgets (stricter thresholds), not the logical structure of the architecture.

1.1.1. Safety Goals (SGs) Addressable by the Layered Architecture

We identify three high-level safety goals to which post-fault SRM control directly contributes: (SG1) prevent unintended motor shutdown that would compromise vehicle controllability; (SG2) prevent phase-winding over-temperature that could lead to insulation failure or thermal runaway; (SG3) prevent uncontrolled torque ripple that would degrade vehicle drivability and exceed Noise, Vibration and Harshness (NVH) limits.

1.1.2. Technical Safety Requirements (TSR) Implemented per Layer

Each architectural layer implements one or more technical safety requirements derivable from the safety goals above:

- **TSR-1** (from SG1): The drive shall continue producing the commanded average torque within a degraded-mode tolerance after a single open-phase converter fault. *This requirement is implemented by Layers 1 and 2 (vector- reconfiguration and soft commutation).*
- **TSR-2** (from SG2): No healthy-phase RMS current shall exceed the rated thermal current (I_{thermal}) during post-fault operation; RMS current bounds of this order for sustained operation are supported by thermal analyses of comparable SRM constructions [29]. In practice, the permissible continuous post-fault torque is ultimately gated by the winding temperature, typically measured by embedded sensors; the RMS-current bound used here is a control-accessible proxy for that thermal limit, not a substitute for temperature supervision. *This requirement is addressed by Layer 3 (asymmetry balance), which equalizes healthy-phase loading; an optional max-penalty limiter targeting this requirement directly is analysed Section 3.6 and shown to coincide with Layer 3 wherever the load can be balanced.*
- **TSR-3** (from SG3): Post-fault torque ripple shall remain below an explicit fault-mode bound, distinct from the healthy-mode specification. *This requirement is partially addressed by Layers 1–2 and discussed in Section 5.3: the simulation results motivate an explicit relaxation of the healthy-mode ripple specification for post-fault operation, since the standard 80% threshold is empirically infeasible under sensor noise and parameter mismatch.*

1.1.3. Fault-Tolerant Time Interval (FTTI)

ISO 26262-3 defines the FTTI as the maximum interval between fault occurrence and the system reaching a safe state without violating any safety goal. For traction-class drives, the FTTI is typically in the range of 10–50 ms [5], dominated by the safe-state transition rather than fault detection itself. The lightweight residual-based detector Section 3.7 declares a phase faulty within two to three sampling periods (i.e., 20–30 μs at $T_s = 10 \mu\text{s}$), so the dominant component of the FTTI in the proposed architecture is the controller's transient toward the re-distributed phase currents, which the Monte Carlo measurement window starts capturing 10 ms after fault onset (Section 4.3). We emphasize that the present work does not *formally verify* FTTI compliance: the 10 ms measurement-window offset is a conservative lower bound on the post-fault transient duration, not a measured time to the safe state. The post-fault metrics reported in this paper are therefore observed inside a typical traction FTTI budget, which is what makes the post-fault behaviour reported in Section 5.2 relevant to ISO 26262 argumentation within the simulation scope of this study.

1.1.4. Mapping the Layers Onto ISO 26262 Concepts

Table 1 summarizes the correspondence. Layers 1 and 2 implement the safe-state transition and degraded-mode operation associated with TSR-1 (continued torque production); Layer 3 (asymmetry balance) equalizes healthy-phase loading and is the layer most directly associated with the per-phase thermal requirement (TSR-2). The optional max-penalty limiter targets TSR-2 explicitly but, as shown in Section 3.6, coincides with Layer 3 wherever the load can be balanced and is therefore reported as an analysed equivalent rather than a separate functional layer.

Table 1. Roles of the three architectural layers and the analysed max-penalty limiter, mapped onto ISO 26262 concepts. Layers 1–2 implement the safe-state transition and degraded-mode torque production; Layer 3 balances healthy-phase thermal loading. The limiter is an analysed equivalent of Layer 3 (Section 3.6), not a separate functional layer.

Layer	ISO 26262 Concept	Evidence in This Paper
Layer 1 (vector-set reconfig.)	Safe-state transition; expanded admissible set after single-point fault	26-vector set restores controllability of the post-fault plant; dominant contributor to speed-RMSE reduction at the primary OP (Section 5.1)
Layer 2 (soft commutation)	Degraded-mode operation, smoothness preservation	Eliminates the fault-blind-vs-hard-FT NMSE/ripple inversion at 700–800 rpm (Section 5.1)
Layer 3 (asymmetry balance)	Degraded-mode thermal load sharing (TSR-2); tracking-quality refinement	Equalizes healthy-phase RMS current, reducing spread relative to L2-MPC, and restores speed regulation, mostly at low speed (Sections 5.1 and 5.2)
Max-penalty limiter (analysed equivalent)	Explicit per-phase thermal cap (TSR-2)	Shares the balanced-state minimizer with Layer 3; adds no measurable effect at the studied threshold of $p = 1.0$ (Sections 3.6 and 5.2)

2. System Model and Open-Phase Fault Description

2.1. 8/6 SRM Plant

The machine under study is a four-phase 8/6 switched reluctance motor. The per-phase electrical dynamics are

$$v_k = R i_k + \frac{d\lambda_k(\theta, i_k)}{dt}, \quad k \in \{A, B, C, D\}, \quad (1)$$

where v_k , i_k , and λ_k are phase voltage, phase current, and phase flux linkage; θ is the rotor's mechanical angle; and R is the phase resistance. Because the phases are magnetically decoupled, the electromagnetic torque is the sum of the individual phase contributions:

$$T_e(\theta, \mathbf{i}) = \sum_{k=A}^D T_e^{(k)}(\theta, i_k), \quad T_e^{(k)}(\theta, i) = \left. \frac{\partial W_{co}^{(k)}(\theta, i')}{\partial \theta} \right|_{i'=i}, \quad (2)$$

with $W_{co}^{(k)}$ representing the co-energy of phase k . The mechanical equation is $J\dot{\omega} = T_e - B\omega - T_L$, where ω is the mechanical angular velocity, J is the inertia, B is the viscous friction, and T_L is the load torque.

The flux-linkage map ($\lambda_k(\theta, i_k)$) uses the piecewise Le-Huy model [30]:

$$\lambda(\theta, i) = \begin{cases} L(\theta) i, & i \leq i_{\text{sat}}(\theta), \\ L_{\text{sat}} i + A(\theta) \left[1 - e^{-\beta(\theta)(i - i_{\text{sat}}(\theta))} \right], & i > i_{\text{sat}}(\theta), \end{cases} \quad (3)$$

where $L(\theta)$ varies between the unaligned inductance (L_u) and the aligned inductance (L_a) L_{sat} is the deep-saturation slope, and $A(\theta)$ and $\beta(\theta)$ are angle-dependent fitting coefficients derived from the saturation flux (λ_M). The inverse map ($i(\theta, \lambda)$) and the analytical torque ($T_e^{(k)}(\theta, i)$) used in Equation (2) are precomputed on a fine (θ, i) grid and stored as look-up tables.

Reference machine.

The numerical values used throughout this study correspond to a laboratory-scale four-phase 8/6 SRM prototype (Figure 1) that the authors used in earlier work on FCS-MPC for SRM drives [31]. The electrical parameters are a phase resistance of $R = 3.1 \, \Omega$, a DC-link voltage of $V_{\text{dc}} = 120 \, \text{V}$, and a maximum phase current of $I_{\text{max}} = 10 \, \text{A}$. The Le-Huy magnetic parameters are an unsaturated (unaligned) inductance of $L_u = 5.90 \, \text{mH}$, an aligned inductance of $L_a = 23.60 \, \text{mH}$, a deep-saturation slope of $L_{\text{sat}} = 0.15 \, \text{mH}$, a saturation flux linkage of $\lambda_M = 0.486 \, \text{Wb}$, and a saturation current of $i_{\text{sat}} = 10.0 \, \text{A}$ at the aligned position. The mechanical parameters are a moment of inertia of $J = 0.0032 \, \text{kg} \cdot \text{m}^2$ and a viscous friction coefficient of $B = 0.001 \, \text{N} \cdot \text{m} \cdot \text{s} / \text{rad}$. These values are summarized together with the controller weights in Table 2.

Table 2. Plant and controller parameters of the reference laboratory prototype. The MPC sampling period is $T_s = 10 \, \mu\text{s}$ throughout this paper.

Parameter	Symbol	Value
Number of phases/ stator–rotor poles	–	4/8–6
DC-link voltage	V_{dc}	120 V
Phase resistance	R	$3.1 \, \Omega$
Maximum phase current	I_{max}	10 A
Unaligned (unsaturated) inductance	L_u	5.90 mH
Aligned inductance	L_a	23.60 mH
Saturation slope (Le-Huy)	L_{sat}	0.15 mH
Saturation flux (Le-Huy)	λ_M	0.486 Wb
Saturation current (aligned position)	i_{sat}	10.0 A
Inertia	J	$3.2 \times 10^{-3} \, \text{kg} \cdot \text{m}^2$
Viscous friction	B	$1.0 \times 10^{-3} \, \text{N} \cdot \text{m} \cdot \text{s} / \text{rad}$
MPC sampling period	T_s	10 μs
Simulation horizon	T_{sim}	200 ms
Fault onset (nominal)	t_f	60 ms
Current soft limit	I_{lim}	12 A
Rated phase RMS current (thermal)	I_{thermal}	8.0 A
Torque-tracking weight	W_T	1.0
Current-limit weight	W_I	10.0
Switching-effort weight	W_S	0.1
Soft-commutation weight	W_P	5.0
Asymmetry-balance weight (variance)	W_B	2.0
Thermal-barrier weight (max)	W_B^{max}	5.0

The aligned inductance ($L_a = 23.60 \, \text{mH}$) is the small-signal value at the aligned position measured at low current; magnetic saturation sets in above the saturation current ($i_{\text{sat}} = 10.0 \, \text{A}$) and is captured by the deep-saturation slope (L_{sat}) and the saturation flux linkage (λ_M) of the Le-Huy model. The machine is a laboratory-scale unit of the 500 W class

with a rated speed of 1500 rpm (corresponding to a rated torque of approximately 3.2 N·m at rated power). The torque–speed envelope considered in this study spans 400–800 rpm and 1–3 N·m, i.e., below base speed in the constant-torque region, where post-fault torque redistribution across the healthy phases is most demanding. The Monte Carlo load of 2 N·m corresponds to roughly 60% of rated torque, so the healthy phases retain magnetic and thermal headroom for partial compensation; full-load fault compensation is not claimed. The redundancy exploited here is specific to the four-phase topology—in a three-phase 6/4 machine, the loss of one phase removes a third of the excitation, and comparable compensation would be substantially harder.




Figure 1. The physical four-phase 8/6 SRM whose measured parameters define this study.

2.2. Asymmetric Bridge Converter and Fault Model

Each phase k is driven by independent asymmetric H-bridge leg. The three switching states per phase are

$$s_k = \begin{cases} +1 & \text{(both switches ON: } v_k = +V_{dc}), \\ 0 & \text{(freewheel: one switch ON, } v_k \approx 0), \\ -1 & \text{(both switches OFF: } v_k = -V_{dc}, \text{ demagnetization).} \end{cases} \quad (4)$$

For a four-phase machine, this yields $3^4 = 81$ switching combinations in the healthy case, although the healthy FCS-MPC implementation used here evaluates eight single-phase-active vectors (one per phase with $s = +1$) plus the zero vector, which is a standard choice in SRM FCS-MPC [14].

Single open-phase fault.

We consider the loss of one converter leg at time $t = t_f$, modelled as an open circuit in phase $k^* \in \{A, B, C, D\}$. From t_f onwards, the faulty phase current (i_{k^*}) is forced to zero, and s_{k^*} is locked at zero, regardless of the controller command. This reduces the admissible vector set of the converter from 3^4 to $3^3 = 27$ states, among which the 26 non-trivial states (excluding the all-zero vector) comprise the candidate MPC evaluation set after reconfiguration. This models the most restrictive failure (both switches of the leg lost and the phase current forced to zero). If only one switch or its gate drive fails, a zero-voltage freewheel loop through the complementary diode may remain physically available in the faulty phase, enlarging the admissible set beyond the 27 considered states; the 26-vector set used here is therefore the conservative worst case.

3. Proposed Fault-Tolerant FCS-MPC Architecture

3.1. Architectural Overview

The proposed controller combines a healthy-mode FCS-MPC core with three fault-reactive layers that activate upon detection of an open-phase fault (Figure 2). In healthy mode, the controller evaluates the eight active vectors plus zero, selecting the one that minimizes a cost function penalizing torque-tracking error, current excess, and switching effort. Upon fault detection at $t = t_f$ in phase k^* , three changes happen concurrently: the evaluation set expands to the 26 admissible post-fault vectors, a position-dependent penalty activates to implement soft commutation, and an *asymmetry-balance* term activates to penalize the dispersion of RMS currents across the three healthy phases. The three layers reshape the cost landscape continuously. We additionally analyse an optional *max-penalty limiter* that would engage only if the largest healthy-phase RMS current exceeded a rated thermal threshold (I_{thermal}); Section 3.6 shows that this limiter shares its minimizer with the asymmetry-balance layer and is therefore redundant in the operating range studied here. It is reported as an analysed, equivalent alternative rather than as a fourth functional layer and is evaluated in the ablation under the L4-MPC label.




Figure 2. Three-layer fault-tolerant FCS-MPC architecture. Vector-set reconfiguration ($8 \rightarrow 26$ active vectors), soft-commutation via a position-dependent cost penalty, and asymmetry balance across healthy phases are activated simultaneously upon open-phase fault detection (dashed arrows). An optional max-penalty limiter, activated when $\max_k I_k^{\text{rms}} > I_{\text{thermal}}$, is analysed in Section 3.6 and shown to coincide with the asymmetry-balance layer wherever the healthy phases can be balanced. The asterisk (*) denotes a reference value (ω^*) or the specific faulty phase index (k^*)

3.2. FCS-MPC Core and Prediction Model

At each sampling instant, the one-step-ahead prediction is

$$\hat{\lambda}_k[n+1] = \lambda_k[n] + T_s(s_k V_{\text{dc}} - R i_k[n]), \quad \hat{i}_k[n+1] = i(\theta[n+1], \hat{\lambda}_k[n+1]), \quad (5)$$

and the predicted electromagnetic torque ($\hat{T}_e[n+1]$) is evaluated through Equation (2) using the Le-Huy-based analytical-phase torque table. The cost function for a candidate vector ($\mathbf{s} \in \mathcal{V}$) is

$$J(\mathbf{s}) = W_T e_T(\mathbf{s}) + W_I e_I(\mathbf{s}) + W_S e_S(\mathbf{s}) + \underbrace{W_P e_P(\mathbf{s})}_{\text{soft commutation}} + \underbrace{W_B e_B(\mathbf{s})}_{\text{asymmetry balance}} + \underbrace{W_B^{\text{max}} e_B^{\text{max}}(\mathbf{s})}_{\text{thermal barrier}}, \quad (6)$$

where $e_T = (\hat{T}_e - T_e^*)^2$ is the squared torque-tracking error, $e_I = \sum_k \max(0, \hat{i}_k - I_{\text{lim}})^2$ is the current-limit excess, e_S counts per-leg switching transitions relative to the previous command, and the three fault-mode terms (e_P , e_B , and e_B^{max}) are gated by weights (W_P , W_B , and W_B^{max}) that are non-zero only after t_f . Weights are listed in Table 2.

3.3. Layer 1: Vector-Set Reconfiguration

In healthy mode, the controller evaluates the eight active vectors ($\mathcal{V}_0$) plus the zero vector. After the fault in phase k^* , the admissible set becomes

$$\mathcal{V}_{k^*} = \{\mathbf{s} \in \{-1, 0, +1\}^4 : s_{k^*} = 0\} \setminus \{\mathbf{0}\}, \quad |\mathcal{V}_{k^*}| = 26. \quad (7)$$

$\mathcal{V}_{k^*}$ contains vectors with two or three simultaneously active phases, which are needed in the intervals where only one healthy phase would otherwise carry the torque demand and saturate its current. It also exposes the -1 demagnetization option in the healthy phases explicitly, so the optimizer can choose active demagnetization at phase commutation.

3.4. Layer 2: Soft Commutation via Position-Dependent Penalty

Classical fault-tolerant strategies superimpose hard turn-on/turn-off angle limits on the current controller: phase k is allowed to proceed only if $\theta_{\text{on}}^{(k)} \leq \theta_k \leq \theta_{\text{off}}^{(k)}$, and the switch command is clamped to $s_k = 0$ outside that window. This introduces a discontinuity at the window edges and tends to produce audible torque spikes at commutation. The soft-commutation layer instead expresses the window as a quadratic cost penalty applied only to the healthy phases ($k \neq k^*$):

$$e_P(\mathbf{s}) = \sum_{k \neq k^*} \phi_k(\theta) \cdot \tilde{i}_k^2[n+1], \quad \phi_k(\theta) = \begin{cases} 0, & \theta_{\text{on}}^{(k)} \leq \theta_k \leq \theta_{\text{off}}^{(k)}, \\ \alpha_k(\theta), & \text{otherwise,} \end{cases} \quad (8)$$

where $\alpha_k(\theta) \in [0, 1]$ is a linear hinge function rising from zero at the conduction-window edge to unity at a user-specified saturation distance of 5° electrical past the window edge. The 5° value is electrical, not mechanical. For the 8/6 SRM considered here, the rotor pole pitch is $\pi/N_r = 60^\circ$ mechanical, so 5° electrical corresponds to $\approx 0.83^\circ$ mechanical. A wider hinge would smooth the post-window transition further but at the cost of admitting more negative-torque-producing states during the falling-inductance region. The weight of $W_P = 5.0$ is obtained from prior Bayesian optimization work on the same machine [31] and kept fixed so the present paper isolates the effect of placing the commutation knowledge inside versus on top of the MPC cost function.

3.5. Layer 3: Asymmetry Balance Across Healthy Phases

When one phase is lost, the three healthy phases must, together, supply the entire torque. Without awareness of the per-phase thermal limit, the phase angularly adjacent to the faulty one tends to absorb a disproportionate share of the torque burden. Layer 3 penalizes this imbalance by adding the variance of the windowed RMS currents across the three healthy phases to the MPC cost:

$$e_B(\mathbf{s}) = \text{Var}_{k \neq k^*}(I_k^{\text{rms}}) = \frac{1}{3} \sum_{k \neq k^*} (I_k^{\text{rms}} - \bar{I}^{\text{rms}})^2, \quad (9)$$

where $\bar{I}^{\text{rms}}$ is the mean of the three healthy-phase windowed RMS currents computed on a sliding window of one electrical stroke. I_k^{rms} is an internal controller state updated recursively. The weight of $W_B = 2.0$ is obtained from earlier Bayesian optimization work [31].

A natural question is whether penalizing the *dispersion* of the healthy-phase currents, as in Equation (9), differs in effect from penalizing their *maximum* directly. This is the subject of the following subsection, which shows the two coincide wherever the load can be balanced.

3.6. Thermally Aware Load Balancing and the Max-Penalty Limiter

Minimizing the dispersion penalty (e_B) of Equation (9) drives the three healthy phases toward equal accumulated thermal loading, as tracked by an exponentially weighted moving average (EWMA) of the squared phase current whose square root is the per-phase RMS proxy used by the controller.

An optional on-demand limiter and why it is redundant.

A natural alternative is to penalize the *maximum* healthy-phase RMS directly, activating only above a thermal threshold (I_{th} , equal to the rated current of $I_{thermal} = 8.0$ A in Table 2 in the deployed controller and deliberately lowered in the sensitivity study reported in Table 3):

$$e_B^{\max} = [\max_{k \in \mathcal{H}} \sqrt{\hat{i}_k^2} - I_{th}]_+^2, \quad (10)$$

corresponding to a one-sided quadratic hinge ($[\cdot]_+ = \max(\cdot, 0)$) that is exactly zero below I_{th} . We include it for completeness but show that it confers no measurable benefit Equation (9) for a structural reason. For a fixed total healthy-phase loading of $S = \sum_{k \in \mathcal{H}} \hat{i}_k^2$ imposed by the torque demand, the dispersion objective Equation (9) and the peak objective Equation (10) share the *same* minimizer—the balanced state ($\hat{i}_k^2 = S/|\mathcal{H}|$), which simultaneously zeroes the variance and minimizes the maximum. Wherever the controller can balance the load, the two penalties coincide; they can differ only *below* I_{th} , where Equation (10) is dormant Equation (9) continues to act.

Table 3. Effect of the max-penalty weight ($W_B^{\max}$) at 400 rpm/3 Nm (noise-free; $I_{th} = 2.0$ A; $W_B^{\max} = 0$ recovers L3-MPC). The limiter is active here yet nearly inert: a $100\times$ weight increase changes every metric by $<1\%$.

$W_B^{\max}$	$I_{rms,healthy}^{\max}$ [A]	Torque NMSE	Ripple [%]
0 (L3-MPC)	3.134	0.439	65.38
1	3.134	0.439	65.38
2	3.132	0.439	65.34
5	3.132	0.439	65.34
10	3.132	0.439	65.34
20	3.130	0.438	65.28
50	3.126	0.438	65.26
100	3.119	0.436	65.11

This is borne out quantitatively. Table 4 shows that the limiter (L4-MPC variant) reproduces the balancing layer (L3-MPC) to within ≤ 0.004 A on the healthy-phase RMS current at every operating point—identical to three decimal places wherever the limiter is dormant. A weight sweep at the most demanding unsaturated point (400 rpm/3 Nm, where the proxy is ≈ 3.13 A $> I_{th}$ so the limiter is active) confirms it (Table 3): increasing $W_B^{\max}$ over two orders of magnitude reduces the worst-phase RMS by only 0.5% (3.134 $\rightarrow$ 3.119 A) and changes torque-tracking NMSE and ripple by less than 0.5%. Sweeping the threshold (I_{th}) leaves the achieved RMS essentially invariant (3.13 A across $I_{th} \in [1.5, 4.0]$ A): the healthy-phase current is fixed by the torque the speed loop must deliver, and the limiter cannot reduce it below that level without sacrificing torque. We therefore retain the dispersion penalty (9) as the operative thermal-balancing mechanism and report (10) only as an analysed, equivalent alternative; the architecture comprises *three* functional layers.

Table 4. Ablation across three operating points: single seed; no sensor noise; and phase A opened at $t_f = 60$ ms, and $T_s = 10$ μ s. Best value per row in **bold**. The *Healthy* column is the no-fault 8-vector baseline (fault never injected) and serves as the pre-fault reference; at 800 rpm, it already approaches the machine’s high-speed tracking limit. At these three operating points, the thermal barrier (L4-MPC) is *dormant*: $\max_k I_k^{\text{rms}}$ never exceeds $I_{\text{thermal}} = 8.0$ A, so L4-MPC reproduces L3-MPC with numerical precision—the designed behaviour of an activated-on-demand barrier.

OP [rpm/Nm]	Metric	Healthy	Fault-Blind	Hard-FT	L2-MPC	L3-MPC	L4-MPC
400/2	Speed RMSE [rad/s]	0.096	1.563	0.827	0.671	0.667	0.667
	NMSE	0.002	1.106	0.721	0.522	0.518	0.522
	Ripple [%]	3.83	103.02	83.13	70.48	70.21	70.48
	$I_{\text{rms,healthy}}^{\text{max}}$ [A]	1.414	2.652	2.365	2.450	2.400	2.450
700/2	Speed RMSE [rad/s]	0.259	1.578	1.358	0.715	0.712	0.715
	NMSE	0.028	0.773	1.030	0.688	0.681	0.688
	Ripple [%]	14.95	89.08	99.97	78.91	78.55	78.91
	$I_{\text{rms,healthy}}^{\text{max}}$ [A]	1.790	3.254	2.525	2.273	2.261	2.273
800/2	Speed RMSE [rad/s]	6.344	10.368	3.724	3.651	3.635	3.651
	NMSE	0.088	0.765	1.661	1.131	1.136	1.131
	Ripple [%]	20.87	84.03	113.72	91.34	91.38	91.34
	$I_{\text{rms,healthy}}^{\text{max}}$ [A]	2.704	3.123	3.846	2.642	2.606	2.642

3.7. Fault Detection and Mode Transition

A lightweight fault-detection layer is included so the evaluation is end-to-end. Detection is based on the residual between the healthy-model prediction and the measured phase current: if $|\hat{i}_k - i_k| > \delta_{\text{det}}$ persists for more than n_{det} consecutive steps in a single phase, that phase is declared faulty, and the three architectural layers activate. In all reported simulations, detection occurs within two to three sampling periods of the fault onset.

4. Simulation Methodology

All experiments were conducted in a reproducible Python3.11.15 simulation framework with checkpointed batch execution (intermediate per-seed results are persisted to disk so that interrupted batches can resume without recomputation). The plant is simulated with the Le-Huy flux-linkage model using the parameters reported in Table 2. Sensor noise and parameter perturbations are injected between the plant and the controller through hook functions, so the controller source is never modified between runs.

4.1. Controller Variants

Four fault-mode variants and one healthy reference are compared (Table 5). Successive rows differ in exactly one architectural element, enabling the ablation reported in Section 5.1.

Table 5. Controller variants compared in this study. “26v” denotes the reconfigured vector set Equation (7).

Variant	Vector Set	Commutation	Soft Cost (W_P)	Asym. Var. (W_B)	Thermal Barrier (W_B^{max})
Healthy	8v	Intrinsic	Off	Off	Off
Fault-blind	8v	Rule-based	Off	Off	Off
Hard-FT	26v	Hard override	Off	Off	Off
L2-MPC	26v	Soft (Equation (8))	On	Off	Off
L3-MPC	26v	Soft (Equation (8))	On	On	Off
L4-MPC	26v	Soft (Equation (8))	On	On	On (Equation (10))

The *Fault-blind* baseline is the most pessimistic reference: it keeps the healthy-mode eight-vector set and the healthy-mode cost function after the fault. The *Hard-FT* variant corresponds to the Stephens-style reference [6,7]: the vector set is reconfigured to $\mathcal{V}_{k^*}$, but commutation is enforced by zeroing the switch command outside the window. The three-layered variants differ in the asymmetry treatment: *L2-MPC* has none, *L3-MPC* adds the variance-based asymmetry balance (Layer 3), and *L4-MPC* adds the optional max-penalty limiter of Equation (10) on top of Layer 3. *L3-MPC* is the proposed three-layer controller; *L4-MPC* is used to evaluate the limiter, which Section 3.6 shows to be equivalent to Layer 3 in the studied regime.

4.2. Operating Points, Fault Scenarios, and Perturbations

Three critical operating points $\{(400, 2), (700, 2), (800, 2)\}$ (rpm, Nm), spanning low, mid, and high speed at the 2 Nm rated load, are used for both the nominal ablation and the Monte Carlo robustness campaign.

The fault is a single open circuit of phase k^* triggered at $t_f = 60$ ms. Sensor noise measurement-only: zero-mean Gaussian noise is added to the controller's current reading at 3% of the full-scale current ($\sigma_i = 0.3$ A RMS) and to the rotor-position reading at $\sigma_\theta = 1^\circ$ RMS, while the plant integrates the true state. For parameter mismatch, each seed draws multiplicative factors (m_R , m_J , and m_B) from $\text{Unif}(0.9, 1.1)$, applied to the plant resistance, inertia, and friction; the controller predicts using the nominal values, so the perturbation represents a mismatch between the plant and the controller's internal model.

All Monte Carlo runs use $N_{\text{seeds}} = 30$. The robustness campaign of Section 5.2 evaluates the five fault-mode variants at the three critical operating points, totalling $5 \times 3 \times 30 = 450$ runs.

4.3. Metrics

Six metrics are computed on a measurement window starting 10 ms after fault onset (to exclude the detection transient) and ending at T_{sim} : **Speed RMSE** [rad/s] is the RMS error of ω against the constant speed reference; **torque RMS ripple** [%] is the RMS of $T_e - \bar{T}_e$ normalized by $\bar{T}_e$, times 100; **torque-tracking NMSE** is $\langle (T_e - T_e^*)^2 \rangle / \langle T_e^{*2} \rangle$; **peak phase current** is I_{peak} [A], i.e., the max of $|i_k|$ across healthy phases and time; **healthy-phase RMS current** is $I_{\text{rms,healthy}}^{\text{max}}$ [A], i.e., the max over the three healthy phases of their RMS current in the measurement window); and **mean torque** is $\bar{T}_e$ [Nm], i.e., the time-averaged torque, used as a tracking sanity check. Per-step MPC execution-time percentiles ($\{P_{50}, P_{95}, \text{and } P_{99, \text{max}}\}$) are also logged.

4.4. Fault-Tolerance Scorecard

A drive is said to tolerate the fault at a given operating point when it meets the six criteria of Table 6 simultaneously in the measurement window. The scorecard produces a score in the range of $\{0, 1, \dots, 6\}$ per run; a run is fully compliant if its score is 6.

Table 6. The six scorecard criteria. All six must be met simultaneously for a run to be declared fault-tolerant at an operating point.

Criterion	Threshold
Speed tracking	Speed RMSE < 5% of the reference speed
Torque tracking	NMSE < 0.25
Peak-current safety	$I_{\text{peak}} \leq 1.5 \cdot \text{rated phase current}$
Thermal safety	$I_{\text{rms,healthy}}^{\text{max}} \leq 1.2 \cdot \text{rated phase RMS current}$
Ripple	Torque RMS ripple < 80%
No runaway	Mean torque $\geq 0.8 \cdot \text{load torque}$

Threshold derivation.

The six thresholds in Table 6 are taken from healthy-mode SRM control practice, with the following justifications. The 5%-of-reference *speed-tracking* budget follows from EV drivability practice, where speed deviations of 2–5% are at the threshold of driver perception during steady cruising; above this band, drivetrain noise and longitudinal jerk become noticeable. The 0.25 NMSE cap follows from the same drivability considerations expressed as a torque-tracking ratio. The peak-current limit at $1.5\times$ rated provides the customary 50% semiconductor margin used in asymmetric-bridge SRM designs [2,3]. The thermal limit at $1.2\times$ rated phase RMS corresponds to the short-time over-current tolerance of insulation class F and class H windings under IEC 60034-1 [32] and is the certification-relevant constraint mapped to TSR-2 in Section 1.1. The 80% *ripple* threshold is a healthy-mode aspiration in the SRM literature [15,16]; we note in Section 5.3 that this threshold is empirically infeasible during a post-fault transient, and we explicitly recommend relaxation to a fault-mode bound as part of the paper’s contribution. The *no-runaway* criterion ($\bar{T}_e \geq 0.8 T_L$) is included as an implicit composite check that screens out controllers that appear to satisfy the thermal criterion only by under-producing torque; this matters for the under-tracking artefact discussed in Section 6.2.

4.5. Statistical Testing

Paired comparisons use the Wilcoxon signed-rank test [28] on per-seed metric values at each OP (`scipy.stats.wilcoxon` [33]). Seeds are matched, so each pair shares the same noise realization and parameter draw. We report the W statistic, two-sided p -value, and significance at $\alpha = 0.01$. Three OPs $\times$ four primary metrics yield 12 paired tests within each comparison family (e.g., L3-MPC vs. Fault-blind, L4-MPC vs. Fault-blind, and L4-MPC vs. L3-MPC), each treated independently. The Bonferroni-corrected threshold within a family of 12 tests is $\alpha_B = 0.01/12 \approx 8.3 \times 10^{-4}$, used consistently in the significance testing of Section 5. A less conservative alternative, the Holm step-down procedure [34], would relax this floor at the cost of greater interpretive complexity; we report the Bonferroni threshold for transparency.

A note on terminology. We use the absence of a significant Wilcoxon result ($p > \alpha$) as evidence *against the presence of a detectable difference at the achieved sample size*, not as evidence *for* statistical equivalence. Formal equivalence testing requires the Two One-Sided Tests (TOST) procedure [35] with pre-specified equivalence margins on each metric (e.g., ± 0.5 rad/s on speed RMSE, ± 0.1 on NMSE, ± 0.2 A on $I_{\text{rms,healthy}}^{\text{max}}$, and ± 5 percentage points on torque ripple). Where the present paper reports “no detectable difference” between L4-MPC and L3-MPC, we mean this in the absence-of-evidence sense and do not claim formal equivalence.

5. Results

This section presents the simulation-based evaluation of the proposed three-layer architecture. The evaluation is organized as follows. Section 5.1 reports the nominal ablation at three operating points (single seed, no sensor noise) to isolate the contribution of each architectural layer under ideal conditions. Section 5.2 reports the 30-seed Monte Carlo robustness campaign at the three critical operating points under measurement noise and $\pm 10\%$ parameter mismatch, with paired-seed Wilcoxon significance tests against the fault-blind baseline and the empirical confirmation that the max-penalty limiter (variant L4-MPC) is exactly equivalent to the asymmetry-balancing controller (L3-MPC) in this regime. Section 5.3 reports the scorecard pass rate. Section 5.4 reports algorithmic complexity and the reference Python implementation timing as a complexity indicator.

5.1. Nominal Ablation Across Three Operating Points

Table 4 reports the four fault-mode variants at three operating points with a single seed and no sensor noise. Successive columns add one architectural layer at a time, so the column-wise differences quantify the individual contributions, as shown in the (Figure 3).




Figure 3. Nominal ablation across the three critical operating points (single seed, no sensor noise, and $T_s = 10 \mu s$). Each panel shows horizontal bars for the five variants on speed RMSE; the gold star at the start of the L3-MPC bar marks the value that L4-MPC reproduces exactly because the thermal barrier is dormant ($\max_k I_k^{rms} < I_{thermal}$ at all three OPs; *not* the L2-MPC column—the asymmetry layer of L3-MPC has already acted, and L4-MPC inherits its values, not those of L2-MPC). The figure visualises the monotonic tracking improvement at 400/2 Nm and the Hard-FT-worse-than-Fault-blind NMSE inversion at 700/2 and 800/2 Nm that the soft-commutation layer recovers.

At the primary operating point (400 rpm/2 Nm), the ablation is clean. The transition from Fault-blind to Hard-FT bundles vector-set reconfiguration with a rule-based commutation override (cf., Table 5), so the 47% speed-RMSE reduction ($1.56 \rightarrow 0.83$ rad/s) attributable to that transition reflects the *joint* effect of these two changes rather than vector-set reconfiguration in isolation; the further reduction of Hard-FT $\rightarrow$ L2-MPC (19%, $0.83 \rightarrow 0.67$ rad/s) isolates the contribution of soft commutation alone. The asymmetry balance (L2-MPC $\rightarrow$ L3-MPC) adds a further 0.5% ($0.671 \rightarrow 0.667$). The same pattern holds for NMSE (53% cumulative reduction, with asymmetry balance contributing 0.7%) and for ripple (31.8% cumulative, 0.4% from asymmetry balance).

At 700 rpm/2 Nm, a counter-intuitive feature appears: Hard-FT is *worse* than the Fault-blind baseline on NMSE (1.03 vs. 0.77) and on ripple (100% vs. 89%) because the hard turn-off override forces premature demagnetization in the expanded vector set. The soft-commutation layer then recovers the loss (Hard-FT $\rightarrow$ L2-MPC reduces NMSE by 33% and ripple by 21%). This is the clearest evidence in the ablation that placing the commutation logic inside the MPC cost function (rather than on top of it) is the right architectural choice at mid-to-high speeds. Asymmetry balance, again, contributes less than 1% to every metric, including a 0.5% reduction in $I_{rms,healthy}^{max}$ ($2.273 \rightarrow 2.261$ A) at this OP.

At 800 rpm/2 Nm, the Fault-blind baseline reports a speed RMSE of 10.37 rad/s—a tracking failure—but the lowest NMSE of all four variants. This is the first appearance of an artefact that recurs in the Monte Carlo results: a controller that fails to track the load reference produces an artificially low NMSE simply because $\hat{T}_e$ collapses toward zero, along with T_e^* . The mean torque ($\bar{T}_e$) at 800/2 (omitted from Table 4 for space) confirms partial torque collapse under Fault-blind, while the three expanded-set variants produce stronger torque and incur higher NMSE values in absolute terms. We discuss this under-tracking artefact and its consequences for thermal-metric interpretation once in Section 6.2.

In summary, the ablation establishes three findings: (i) the dominant contributions come from the joint vector-set reconfiguration plus soft-commutation pair, not asymmetry balance; (ii) tracking metrics alone can be misleading under torque under-production;

and (iii) the thermal barrier is dormant under nominal conditions (L4-MPC $\equiv$ L3-MPC to numerical precision), as designed.

5.2. Robustness Under Measurement Noise and Parameter Mismatch

The nominal ablation isolates each layer's contribution but uses a single noise-free seed. To assess robustness, we run a Monte Carlo campaign at the three critical operating points (400/2, 700/2, and 800/2 Nm) with independent seeds each. Two perturbations are applied per seed and held paired across variants: *measurement-only* sensor noise, in which the controller reads a corrupted current ($\sigma_i = 0.3$ A, 3% of $I_{\max}$) and position ($\sigma_\theta = 1^\circ$) while the plant integrates the true state, and a $\pm 10\%$ perturbation of the plant resistance, inertia, and friction (R , J , and B), with the controller predicting using the nominal R . The fault is applied at $t_f = 60$ ms, and metrics are computed over $[t_f + 10 \text{ ms}, T_{\text{sim}}]$. Five variants are reported: Fault-blind; Hard-FT; and the cumulative MPC variants, i.e., L2-MPC, L3-MPC, and L4-MPC.

Fault tolerance restores speed regulation.

Across all three operating points, the proposed controller restores speed regulation relative to the Fault-blind baseline (Table 7). The improvement is modest at low speed (400 rpm/2 Nm: 1.25 vs. 1.57 rad/s) and pronounced at high speed, where the Fault-blind controller loses regulation entirely (800 rpm/2 Nm: 2.53 vs. 12.61 rad/s). Paired Wilcoxon signed-rank tests of L3-MPC against Fault-blind on matched seeds ($N = 30$) confirm that the speed-RMSE gain is significant at the two higher-speed operating points ($p < 10^{-8}$ at both 700/2 and 800/2 Nm), where the Fault-blind baseline loses regulation; at 400/2 Nm, where the improvement is modest, the gain is only marginal ($p = 0.014$, significant at $\alpha = 0.05$ but not at the $\alpha = 0.01$ threshold used here).

Table 7. Monte Carlo robustness under measurement noise and $\pm 10\%$ parameter mismatch (30 seeds per cell, reported as the mean). Single open-phase fault, phase A opened at $t_f = 60$ ms, and $T_s = 10 \mu\text{s}$. L4-MPC and L3-MPC are statistically identical at every operating point (paired Wilcoxon, $p = 1.0$; Table 8); the L4-MPC row equals L3-MPC and is omitted. Hard-FT is retained as a cautionary reference.

OP [rpm/Nm]	Variant	Speed RMSE [rad /s]	Ripple [%]	I_{pk} [A]	$\bar{T}_e$ [Nm]
400/2	Fault-blind	1.57	102.9	12.19	2.041
	Hard-FT	0.84	83.8	10.55	2.042
	L2-MPC	1.59	83.1	11.25	2.055
	L3-MPC	1.25	78.0	10.97	2.062
700/2	Fault-blind	2.49	81.5	8.81	2.109
	Hard-FT	1.65	110.6	9.85	2.058
	L2-MPC	0.69	79.2	7.93	2.050
	L3-MPC	0.69	78.7	7.89	2.049
800/2	Fault-blind	12.61	81.1	8.79	2.120
	Hard-FT	4.85	110.9	14.20	2.413
	L2-MPC	2.53	95.0	9.56	2.355
	L3-MPC	2.53	95.4	9.58	2.357

Asymmetry balance helps most at low speed.

The asymmetry-aware layer (L3-MPC) matches or improves on the balancing-free variant (L2-MPC) and never degrades it. The clearest margin is at 400 rpm/2 Nm, where L2-MPC alone barely separates from the Fault-blind baseline on speed RMSE (1.59 vs. 1.57 rad/s) while L3-MPC reaches 1.25 rad/s; at 700 and 800 rpm, the two converge (≈ 0.69 and ≈ 2.53 rad/s, respectively). L3-MPC ripple is at or below L2-MPC at every point.

The max-penalty limiter is exactly equivalent here.

With $I_{th} = 8$ A, the limiter never binds: the largest healthy-phase RMS current the proposed controller draws anywhere in the campaign is 3.2 A. As a result, L4-MPC reproduces L3-MPC *exactly*. All twelve paired Wilcoxon comparisons (four metrics $\times$ three operating points) return $p = 1.0$, with per-seed means identical to twelve significant figures (Table 8). This is the empirical counterpart of the structural equivalence of Section 3.6: wherever the healthy phases can be balanced, penalizing dispersion and penalizing the maximum coincide.

Table 8. Paired Wilcoxon signed-rank tests (L4-MPC vs. L3-MPC) on matched seeds ($N = 30$). With the limiter dormant at $I_{th} = 8$ A, the two controllers are bit-identical at every operating point and for every metric.

OP [rpm/Nm]	Metrics (Speed RMSE, NMSE, Ripple, and $I_{rms,healthy}^{max}$)	p -Value
400/2	all identical	1.000
700/2	all identical	1.000
800/2	all identical	1.000

Limitations of the robustness campaign.

Three caveats accompany these results see (Figure 4). First, the absolute torque ripple is high (78–110%) because a single open-phase fault in a four-phase machine removes a quarter of the available excitation; this is intrinsic to the fault, not a controller artefact. At high speed, the fault-tolerant variants exhibit *higher* ripple than Fault-blind ($\approx 95\%$ vs. 81% at 800/2 Nm) because the Fault-blind controller under-produces torque while the proposed controller pushes to maintain it—a trade of increased ripple for restored regulation, not a ripple reduction. Second, the Hard-FT baseline is the least favourable on both counts: it raises ripple to $\approx 110\%$ and drives peak current to 14.2 A (against the 15 A peak-current limit, $1.5\times$ rated) at 800 rpm, which motivates the soft-commutation and balancing layers rather than a rule-based override. Third, the $\pm 10\%$ parameter mismatch drives peak current above the 10 A nominal threshold at low speed (up to 12.2 A for Fault-blind and 11.0 A for L3-MPC at 400/2 Nm), entering mild magnetic saturation while operation remains stable; this is the expected consequence of predicting with a nominal model under a de-rated plant and bounds the current margin a deployment would need to allow.




Figure 4. Monte Carlo robustness at the three critical operating points (30 seeds per cell, measurement-only sensor noise, and $\pm 10\%$ parameter mismatch) (a) Speed RMSE (error bars: ± 1 std): the Fault-blind baseline loses regulation at 800 rpm (12.6 rad/s), while the proposed controller holds 2.5 rad/s (b) Torque ripple: the fault-tolerant variants trade higher ripple at high speed for restored regulation; Hard-FT is the least favourable. L4-MPC is omitted, as it is identical to L3-MPC (Table 8).

5.3. Scorecard Pass Rate Under Monte Carlo

Applying the healthy-mode fault-tolerance scorecard of Table 6 to the robustness campaign of Section 5.2 shows that its six criteria cannot be passed simultaneously during

a post-fault transient under the tested uncertainty, for any variant. The binding constraints are the speed-regulation criterion—the 5%-of-reference band is routinely exceeded once 3% current noise is injected—and, decisively, the 80% ripple criterion: post-fault ripple sits between 78% and 110% at every tested operating point (Table 7), so the ripple criterion alone precludes an all-pass outcome. The healthy-phase RMS-current criterion, by contrast, is comfortably met—the largest healthy-phase RMS current the proposed controller draws across the campaign is 3.2 A, against the 8 A rated threshold.

We interpret this not as a controller failure but as evidence that the thresholds of Table 6—typical of healthy-mode SRM control [6,8]—are too strict to pass simultaneously under sensor noise and parameter mismatch during a post-fault transient, when a quarter of the machine’s excitation is unavailable. post-fault-specific scorecard with relaxed ripple and speed bands is the appropriate evaluation instrument; its calibration is left to future work, alongside hardware validation. The implications for ISO 26262 post-fault certification are discussed in Section 6.

5.4. Algorithmic Complexity and Computational Cost

Algorithmic complexity (FLOP count per MPC step).

The post-fault MPC step evaluates $|\mathcal{V}_{k^*}| = 26$ candidate vectors. Each candidate evaluation comprises four phase predictions through Equation (5) (each requiring two multiplies and one add for the flux update, then a bilinear interpolation in the $i(\theta, \lambda)$ table costing approximately eight FLOPs), one electromagnetic-torque evaluation through Equation (2) (a four-phase summation of bilinear interpolations in the $T_e^{(k)}(\theta, i)$ tables ≈ 36 FLOPs in total), and a cost evaluation summing the six terms in Equation (6) (≈ 12 FLOPs). Therefore, the per-candidate budget is $4 \cdot (3 + 8) + 36 + 12 \approx 92$ FLOPs, yielding $26 \cdot 92 \approx 2400$ FLOPs per MPC step for the full post-fault evaluation. The healthy-mode eight-vector evaluation is proportionally smaller, at ≈ 740 FLOPs per step.

Reference Python timing.

Table 9 reports per-step MPC execution-time percentiles of the reference Python implementation on a Colab CPU runtime, aggregated across 30 seeds at each operating point. These values are not real-time measurements: P_{99} exceeds $T_s = 10 \mu\text{s}$ by roughly a factor of 450 (i.e., $P_{99} \approx 4.5 \text{ ms}$, while $T_s = 10 \mu\text{s}$), reflecting Python interpreter overhead in Colab rather than the algorithmic cost of evaluating 26 candidate vectors. The distribution is informative as a complexity indicator: P_{50} is roughly constant at 1.8–1.9 ms across the operating points, showing that operating-point-dependent branching does not meaningfully affect the average case; $P_{99}/P_{50} \approx 2.4$ is consistent with a fixed work budget plus interpreter-level variability. A compiled implementation (e.g., the Numba JIT acceleration we used in an earlier study, which gave a $\sim 50\times$ speed-up on a closely related SRM simulator [31]) would reduce these timings by an order of magnitude or more, but compiled-runtime real-time validation is outside the scope of this simulation-based study.

Table 9. Per-step MPC execution-time percentiles in microseconds for L3-MPC at the three critical operating points for a reference Python implementation on a Colab CPU runtime, aggregated across 30 seeds per OP, reported as a complexity indicator for the simulation framework.

rpm	Load [Nm]	P_{50}	P_{95}	P_{99}	max
400	2	1864.1	3682.9	4547.1	10,341.4
700	2	1830.7	3616.9	4427.7	10,679.1
800	2	1808.6	3557.1	4347.8	13,547.6

6. Discussion

6.1. Where the Architectural Contributions Come From

Viewed across the ablation (Table 4) and the Monte Carlo robustness campaign (Table 7), the three layers contribute on distinct axes. Layers 1 and 2 drive tracking performance; Layer 3 (asymmetry balance) adds a smaller but consistent improvement, which is largest at low speed, and equalizes healthy-phase thermal loading. The optional max-penalty limiter targets the per-phase thermal cap directly but coincides with Layer 3 wherever the load can be (Section 3.6).

Vector-set reconfiguration plus soft commutation do the heavy lifting.

The transition from the 8-vector healthy set to the 26-vector post-fault admissible set, taken together with the move from rule-based to in-cost commutation, accounts for the bulk of the speed-RMSE reduction at 400 rpm/2 Nm in the nominal ablation (Fault-blind $\rightarrow$ L2-MPC: 1.56 $\rightarrow$ 0.67 rad/s, a 57% reduction) and continues to dominate the Fault-blind-to-L2-MPC improvement under Monte Carlo simulation at every OP. This is consistent with the observation that the post-fault plant admits a larger admissible set than the healthy plant; the optimizer exploits the extra degrees of freedom only when the commutation window is not used to discard them.

Soft commutation matters most at mid-to-high speed.

At 700 rpm/2 Nm and 800 rpm/2 Nm, the Hard-FT variant is *worse* than the Fault-blind baseline on NMSE and ripple: the rule-based turn-off override forces premature demagnetization that the Fault-blind dispatch avoids. Soft commutation recovers most of this difference and, at 700/2 Nm, reduces NMSE by 33% and ripple by 21% over Hard-FT. This is the clearest evidence in the paper that placing the commutation logic inside the MPC cost function is the right architectural choice—a cautionary observation that may help explain the operating-point-specific performance variability reported in prior FCS-MPC fault- work [20,21], which combines reconfiguration with rule-based commutation.

Dispersion versus maximum: why the limiter is redundant.

The asymmetry-balance penalty of Equation (9) equalizes the accumulated loading of the three healthy phases. A natural concern is that equalizing dispersion need not bound the *maximum*: a controller could equalize the phases while pushing all of them toward the thermal limit. This motivates the explicit max-penalty limiter of Equation (10). As shown in Section 3.6, however, the two objectives share the balanced-state minimizer whenever the load can be balanced, so the limiter is redundant wherever it would act. The robustness campaign confirms this directly: with $I_{th} = 8$ A, the limiter never binds (the largest healthy-phase RMS current the proposed controller draws is 3.2 A), and L4-MPC reproduces L3-MPC exactly—all twelve paired Wilcoxon comparisons return $p = 1.0$. We therefore report the limiter as an analysed, equivalent alternative and frame the per-phase thermal bound as safety-oriented motivation only; no thermal-compliance guarantee or fault-tolerant time-interval budget is claimed.

6.2. Torque Under-Delivery and Ripple Interpretation

One interpretive caution applies to the high-speed results. At 800 rpm/2 Nm, the Fault-blind baseline reports the *lowest* torque ripple of any variant (81%, Table 7) despite the *largest* speed RMSE by a wide margin (12.6 rad/s). The low ripple is not evidence of good fault-tolerant behaviour: it reflects a baseline that has partially surrendered speed regulation, so its torque trajectory is smoother simply because it tracks the demand less aggressively. Therefore, torque ripple should be read alongside speed RMSE and mean torque rather than in isolation; the fault-tolerant variants accept higher ripple as the cost of

holding speed regulation under the fault. The same caution applies to the healthy-phase RMS current, which is monotone in delivered torque to the first order: a controller that under-delivers torque also draws less current, so absolute current levels are comparable only between controllers that deliver matched mean torque.

6.3. Implications of the Scorecard Outcome

The absence of an all-pass scorecard outcome under Monte Carlo simulation (Section 5.3) reflects the binding of the speed-RMSE and ripple budgets under realistic uncertainty, not a thermal failure: the healthy-phase RMS-current criterion is the most-often-passed of the six. Read literally, healthy-mode SRM control specifications cannot all be preserved during a post-fault transient under sensor noise and parameter mismatch. The fault-tolerant SRM literature reports post-fault specifications evaluated mostly under nominal conditions [6–8]. Two practical responses follow: explicit relaxation of fault-mode thresholds (e.g., a 15–20% speed-RMSE budget and a 150% ripple budget during the post-fault transient) or reporting of per-criterion pass-rates rather than the six-way intersection. We recommend the first as language for future post-fault SRM evaluation.

6.4. Scope and Threats to Validity

The scope of this study should be made explicit. First, all results are simulation-based: the plant is modelled with the Le-Huy flux-linkage map identified from a laboratory-prototype reference, sensor noise is injected synthetically, and parameter mismatch is sampled from a uniform $\pm 10\%$ box around the nominal value. Therefore, the findings are claims about the controller architecture under these modelling assumptions and not measurements on a physical drive. Second, the controller weights (W_p , W_B , and $W_B^{\max}$) are fixed at values obtained from earlier Bayesian optimization work [31] at a single OP or from the load-balancing analysis Section 3.6; per-OP gain scheduling is outside the scope of this paper. Third, the limiter threshold of $I_{th} = 8.0$ A coincides with the simulated machine's rated phase RMS current and lies below $I_{max} = 10$ A; because the limiter remained dormant at this threshold throughout the campaign (the largest healthy-phase RMS current drawn by the proposed controller was 3.2 A), the present data do not quantify how its activation frequency and the associated tracking penalty would scale as I_{th} is reduced toward the sustained thermal rating. Fourth, the robustness perturbations are synthetic and a conservative lower bound on real measurement variability; truly independent runs with varied winding temperatures, mechanical wear, and component ageing would likely exhibit greater distributional shift. Fifth, the compute-time measurements in Table 9 reflect Python interpreter overhead and Colab host variability and should be read only as a complexity indicator, not as a real-time budget. Sixth, the scorecard criteria of Table 6 are a synthesis of healthy-mode SRM control practice and were not adjusted for post-fault operation, as discussed in Section 5.3. Seventh, the scope is restricted to single open-phase converter faults: multi-phase faults and short-circuit faults are not addressed by the cost-function formulation considered here. Eighth, the plant model treats the four phases as magnetically decoupled (Equation (1)); real SRMs exhibit weak inter-phase mutual coupling that is most pronounced near the aligned position, which is not captured here and which would modestly redistribute the post-fault healthy-phase currents.

7. Conclusions

This paper proposed and systematically evaluated, by simulation, a three-layer fault-tolerant FCS-MPC architecture for a four-phase 8/6 SRM under single open-phase faults, motivated by the use of SRMs in EV traction, where per-phase thermal integrity after a fault is a hard certification requirement. The architecture combines vector-set reconfiguration

(8 → 26 active vectors), soft commutation via a position-dependent cost-function penalty, and variance-based asymmetry balance across the three healthy phases; an explicit max-penalty thermal limiter is analysed as an equivalent alternative rather than a separate layer.

Within the simulation study, the explicit max-penalty thermal limiter was found to be equivalent to the variance-based balancing term: for a fixed healthy-phase loading, the two penalties share the balanced-state optimiser, and a noise-free sweep across the three critical operating points confirmed that the limiter reproduces the variance-based controller to within ≤ 0.004 A on the worst healthy-phase RMS current, with a two-order-of-magnitude change in its weight altering torque ripple by less than 1%. Under the 30-seed Monte Carlo configuration the two controllers were statistically identical at every operating point ($p = 1.0$ on all paired comparisons). Therefore, the limiter is reported as an analysed, equivalent alternative rather than a distinct performance layer, and the architecture comprises three functional layers. We frame the per-phase thermal bound as safety-oriented motivation only; no fault-tolerant time-interval budget or ISO 26262 compliance is claimed.

The deterministic nominal ablation confirms a monotonic improvement at the primary operating point (400 rpm/2 Nm): speed RMSE drops from 1.56 rad/s (Fault-blind) to 0.67 rad/s for the full controller. The combination of vector-set reconfiguration and the move from rule-based to in-cost commutation accounts for the majority of this gain; asymmetry-aware balancing adds a smaller but consistent improvement and evens the loading of the healthy phases. At 700 rpm/2 Nm, rule-based hard commutation is worse than the Fault-blind baseline in NMSE and ripple, while soft commutation recovers this loss, supporting the choice of placing commutation knowledge inside the MPC cost function.

Under a 30-seed Monte Carlo simulation with measurement-only sensor noise and $\pm 10\%$ parameter mismatch, the proposed controller restores post-fault speed regulation at every operating point—markedly so at high speed, where the fault-blind baseline loses regulation entirely (speed RMSE 2.5 vs. 12.6 rad/s at 800 rpm/2 Nm; paired Wilcoxon $p < 10^{-8}$). A strict six-criterion scorecard is not met at any operating point, which we interpret as evidence that healthy-mode SRM specifications must be explicitly relaxed for post-fault operation. All findings are simulation-based and anchored to a laboratory-measured machine model; hardware validation of the post-fault converter behaviour is the natural next step.

Author Contributions: Conceptualization, F.S., M.I.M.-M. and E.R.-C.; methodology, F.S.; software, F.S.; validation, F.S., M.I.M.-M. and E.R.-C.; formal analysis, F.S.; investigation, F.S.; resources, M.I.M.-M. and E.R.-C.; data curation, F.S.; writing—original draft preparation, F.S.; writing—review and editing, F.S., M.I.M.-M. and E.R.-C.; visualization, F.S.; supervision, M.I.M.-M. and E.R.-C.; project administration, M.I.M.-M. and E.R.-C.; funding acquisition, M.I.M.-M. and E.R.-C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The controller source code, simulation notebook, and raw experimental data (parquet files) are available upon request from the corresponding author. The numerical analysis pipeline is built on SciPy [33] and NumPy [36].

Acknowledgments: The authors would like to thank Universidad de las Fuerzas Armadas ESPE for supporting this work, including covering the registration fee related to the publication of this paper. This work was also supported in part by the European Regional Development Fund—“A way to make Europe” under GR24040 project.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ASIL	Automotive Safety Integrity Level
EV	Electric Vehicle
FCS-MPC	Finite Control Set Model Predictive Control
FTTI	Fault-Tolerant Time Interval
HVAC	Heating, Ventilation and Air Conditioning
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization
MC	Monte Carlo
MPC	Model Predictive Control
NMSE	Normalized Mean-Square Error
NVH	Noise, Vibration and Harshness
OP	Operating Point
RMSE	Root-Mean-Square Error
SG	Safety Goal
SRM	Switched Reluctance Motor
TOST	Two One-Sided Tests
TSR	Technical Safety Requirement

References

1. Miller, T.J.E. *Electronic Control of Switched Reluctance Machines*; Newnes: Oxford, UK, 2001.
2. Krishnan, R. *Switched Reluctance Motor Drives: Modeling, Simulation, Analysis, Design, and Applications*; CRC Press: Boca Raton, FL, USA, 2001.
3. Bilgin, B.; Jiang, J.W.; Emadi, A. (Eds.) *Switched Reluctance Motor Drives: Fundamentals to Applications*; CRC Press: Boca Raton, FL, USA, 2019. [[CrossRef](#)]
4. Bilgin, B.; Emadi, A. Electric motors in electrified transportation: A step toward achieving a sustainable and highly efficient transportation system. *IEEE Power Electron. Mag.* **2014**, *1*, 10–17. [[CrossRef](#)]
5. ISO 26262:2018; Road Vehicles—Functional Safety, Parts 1–12. International Organization for Standardization (ISO): Geneva, Switzerland, 2018.
6. Stephens, C.M. Fault detection and management system for fault-tolerant switched reluctance motor drives. *IEEE Trans. Ind. Appl.* **1991**, *27*, 1098–1102. [[CrossRef](#)]
7. Mir, S.; Husain, I.; Elbuluk, M.E. Switched reluctance motor modeling with online parameter identification. *IEEE Trans. Ind. Appl.* **1998**, *34*, 776–783. [[CrossRef](#)]
8. Gameiro, N.S.; Marques Cardoso, A.J. A new method for power converter fault diagnosis in SRM drives. *IEEE Trans. Ind. Appl.* **2012**, *48*, 653–662. [[CrossRef](#)]
9. Ding, W.; Hu, Y.; Wu, L. Investigation and experimental test of fault-tolerant operation of a mutually coupled dual three-phase SRM drive under faulty conditions. *IEEE Trans. Power Electron.* **2015**, *30*, 6857–6872. [[CrossRef](#)]
10. Rodríguez, J.; Pontt, J.; Silva, C.A.; Correa, P.; Lezana, P.; Cortés, P.; Ammann, U. Predictive current control of a voltage source inverter. *IEEE Trans. Ind. Electron.* **2007**, *54*, 495–503. [[CrossRef](#)]
11. Kouro, S.; Cortés, P.; Vargas, R.; Ammann, U.; Rodríguez, J. Model predictive control—A simple and powerful method to control power converters. *IEEE Trans. Ind. Electron.* **2009**, *56*, 1826–1838. [[CrossRef](#)]
12. Vázquez, S.; Rodríguez, J.; Rivera, M.; Franquelo, L.G.; Norambuena, M. Model predictive control for power converters and drives: Advances and trends. *IEEE Trans. Ind. Electron.* **2017**, *64*, 935–947. [[CrossRef](#)]
13. Brosch, A.; Hanke, S.; Wallscheid, O.; Böcker, J. Data-driven recursive least-squares estimation for model predictive current control of permanent magnet synchronous motors. *IEEE Trans. Power Electron.* **2021**, *36*, 2179–2190. [[CrossRef](#)]
14. Li, X.; Shamsi, P. Model predictive current control of switched reluctance motors with inductance auto-calibration. *IEEE Trans. Ind. Electron.* **2016**, *63*, 3934–3941. [[CrossRef](#)]
15. Peng, F.; Emadi, A. A digital PWM current controller for switched reluctance motor drives. In *Proceedings of the 2014 IEEE Transportation Electrification Conference and Expo (ITEC), Dearborn, MI, USA, 15–18 June 2014*; IEEE: Piscataway, NJ, USA, 2014; pp. 1–6. [[CrossRef](#)]

16. Xue, X.D.; Cheng, K.W.E.; Ho, S.L. Optimization and evaluation of torque-sharing functions for torque ripple minimization in switched reluctance motor drives. *IEEE Trans. Power Electron.* **2009**, *24*, 2076–2090. [\[CrossRef\]](#)
17. Fang, G.; Ye, J.; Xiao, D.; Xia, Z.; Emadi, A. Computational-efficient model predictive torque control for switched reluctance machines with linear-model-based equivalent transformations. *IEEE Trans. Ind. Electron.* **2022**, *69*, 5465–5477. [\[CrossRef\]](#)
18. Sánchez, F.; Milanés-Montero, M.I.; Romero-Cadaval, E.; Llanos, J.; Moreano, G. Electric vehicle-oriented predictive control for SRMs 8/6 with optimized dual-phase excitation vectors. *Energies* **2025**, *18*, 6246. [\[CrossRef\]](#)
19. Valencia, D.F.; Tarvirdilu-Asl, R.; García, C.; Rodríguez, J.; Emadi, A. A review of predictive control techniques for switched reluctance machine drives. Part I: Fundamentals and current control. *IEEE Trans. Energy Convers.* **2021**, *36*, 1313–1322. [\[CrossRef\]](#)
20. Hennen, M.D.; Niessen, M.; Heyers, C.; Brauer, H.J.; De Doncker, R.W. Development and control of an integrated and distributed inverter for a fault-tolerant five-phase switched reluctance traction drive. *IEEE Trans. Power Electron.* **2012**, *27*, 547–554. [\[CrossRef\]](#)
21. Chen, H.; Xu, D.; Deng, X. Control for power converter of small-scale switched reluctance wind power generator. *IEEE Trans. Ind. Electron.* **2021**, *68*, 3148–3158. [\[CrossRef\]](#)
22. Tao, R.; Wang, W.; Luo, C.; Sun, L.; Cheng, M. Full-torque copper-loss minimum adaptive fault-tolerant control of dual three-phase PMSM under open-phase faults. *IEEE Trans. Transp. Electr.* **2026**, *12*, 1761–1774. [\[CrossRef\]](#)
23. Tao, R.; Wang, W.; Jiang, Y.; Tang, C.; Fan, Y.; Hua, W. Free-current-curve fault-tolerant control of dual three-phase permanent magnet synchronous motor under open-phase fault. *IEEE Trans. Transp. Electr.* **2025**, *11*, 11619–11631. [\[CrossRef\]](#)
24. Ali, N.; Narimani, M. Fault-tolerant SRM drives—A review. *IEEE Trans. Power Electron.* **2024**, *39*, 10261–10275. [\[CrossRef\]](#)
25. Alam, M.; Payami, S. A novel control-independent online fault diagnosis of interturn short circuits in switched reluctance motors using signal injection technique. *IEEE Trans. Ind. Electron.* **2023**, *70*, 2157–2167. [\[CrossRef\]](#)
26. Zhang, P.; Li, K.; Yu, S.; Yu, D. A novel fault diagnosis technique of interturn short-circuit fault for switched reluctance motors in current chopper mode. *IEEE Trans. Ind. Electron.* **2022**, *69*, 3037–3046. [\[CrossRef\]](#)
27. Hamouda, M.; Al-Amyal, F.; Elsherbiny, H.; Odinaev, I.; Menaem, A.A.; Alluhaybi, K.; Zaky, A.A. A novel interturn fault tolerant-based average torque control of switched reluctance motors for electric vehicles. *IEEE Access* **2024**, *12*, 111769–111781. [\[CrossRef\]](#)
28. Wilcoxon, F. Individual comparisons by ranking methods. *Biometrics Bull.* **1945**, *1*, 80–83. [\[CrossRef\]](#)
29. Arbab, N.; Wang, W.; Lin, C.; Hearron, J.; Fahimi, B. Thermal modeling and analysis of a double-stator switched reluctance motor. *IEEE Trans. Energy Convers.* **2015**, *30*, 1209–1217. [\[CrossRef\]](#)
30. Le-Huy, H.; Brunelle, P. A versatile nonlinear switched reluctance motor model in Simulink using realistic and analytical magnetization characteristics. In *Proceedings of the 31st Annual Conference of the IEEE Industrial Electronics Society (IECON 2005), Raleigh, NC, USA, 6–10 November 2005*; IEEE: Piscataway, NJ, USA, 2005. [\[CrossRef\]](#)
31. Sánchez, F.; Milanés-Montero, M.I.; Romero-Cadaval, E.; García, C. Power-invariant $\alpha\beta$ transformation for FCS-MPC of 8/6 SRMs: A torque ripple reduction approach. In *Proceedings of the 2025 IEEE 8th Student Conference on Electric Machines and Systems (SCEMS), Busan, Republic of Korea, 20–22 November 2025*; IEEE: Piscataway, NJ, USA, 2025. [\[CrossRef\]](#)
32. IEC 60034-1; Rotating Electrical Machines—Part 1: Rating and Performance. International Electrotechnical Commission (IEC): Geneva, Switzerland, 2017.
33. Virtanen, P.; Gommers, R.; Oliphant, T.E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Peterson, P.; Weckesser, W.; Bright, J.; van der Walt, S.J.; et al. SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nat. Methods* **2020**, *17*, 261–272. [\[CrossRef\]](#) [\[PubMed\]](#)
34. Holm, S. A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* **1979**, *6*, 65–70.
35. Schuirmann, D.J. A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability. *J. Pharmacokinet. Biopharm.* **1987**, *15*, 657–680. [\[CrossRef\]](#) [\[PubMed\]](#)
36. Harris, C.R.; Millman, K.J.; van der Walt, S.J.; Gommers, R.; Virtanen, P.; Cournapeau, D.; Wieser, E.; Taylor, J.; Berg, S.; Smith, N.J.; et al. Array programming with NumPy. *Nature* **2020**, *585*, 357–362. [\[CrossRef\]](#) [\[PubMed\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.