

DISSERTATION

Numerical Analysis and Innovative Simulation Techniques for Designing Advanced MRAM

ausgeführt zum Zwecke der Erlangung des akademischen Grades
eines Doktors der technischen Wissenschaften

unter der Betreuung von

Associate Prof. Victor Sverdlov, Ph.D. MSc.
O.Univ.Prof. Dipl.-Ing. Dr.techn. Dr.h.c. Siegfried Selberherr

eingereicht an der Technischen Universität Wien
Fakultät für Elektrotechnik und Informationstechnik
von

Dipl.-Ing. Mario Bendra
Matrikelnummer: 51820045

Wien, im April 2026

Abstract

Scaling spin-transfer torque magnetoresistive random access memory (STT-MRAM) to sub-20 nm dimensions poses major challenges for thermal stability, switching reliability, and energy efficiency. As devices shrink, the macro-spin approximation breaks down, requiring a rigorous treatment of non-uniform magnetization dynamics and transport phenomena. This thesis presents a computational framework for modeling and optimizing ultra-scaled, multi-layered magnetoresistive random access memory (MRAM) devices.

First, a three-dimensional micromagnetic solver based on finite element method (FEM) is developed. To address the computationally expensive demagnetization field in complex geometries, a hybrid finite element-boundary element method (FE-BEM) is used. Numerical stiffness from the exchange interaction is handled by a *tangent-plane time integration scheme*, which reformulates the nonlinear Landau–Lifshitz–Gilbert equation as a linear saddle-point problem. This approach is benchmarked against *adaptive higher-order backward differentiation formula (BDF) and implicit-explicit (IMEX) methods*, demonstrating superior stability and efficiency for the stiff dynamics of ultra-scaled magnetic elements.

Central to this work is the extension of the micromagnetic framework to include a coupled spin and charge drift-diffusion formalism. Unlike standard models that assume continuity of spin currents, this work derives and implements specialized boundary conditions. These are used to rigorously describe spin filtering across magnesium oxide (MgO) tunnel barriers and spin dephasing at metallic interfaces. Furthermore, a novel numerical approach for interlayer exchange coupling (IEC) is introduced. This leverages an interface-mapping algorithm to efficiently simulate synthetic antiferromagnet (SAF) without the prohibitive computational cost of volumetric meshing of angstrom-scale spacers.

Applying this unified solver reveals that reliability issues in ultra-scaled devices, such as back-hopping, are deterministic consequences of the inter-segment torque hierarchy within composite free layers. They are not stochastic thermal effects. Crucially, this thesis demonstrates that by engineering the relative tunnel barrier polarizations and free layer geometry, this instability can be either suppressed for reliable binary switching or deliberately exploited. This allows realization of multi-level cell (MLC) capable of storing multiple bits per physical cell through distinct intermediate magnetization configurations. Additionally, double-spin torque magnetic tunnel junction (ds-MTJ) architectures are investigated. In these, the cooperation

between giant magnetoresistance (GMR) and tunneling magnetoresistance (TMR) torques enables sub-nanosecond switching. When IEC is introduced in advanced multi-layer stacks, it emerges as a unifying design parameter. Specific coupling windows govern the stability of the SAF-enhanced reference layer. Hybrid free layers with metallic spacers exploit the spacer material's spin-flip length to enhance switching speed. The joint optimization of spacer material and thickness determines write efficiency.

Kurzfassung

Die Skalierung von Spin-Transfer Torque Magnetoresistive Random Access Memory (STT-MRAM) in den Sub-20 nm-Bereich bringt tiefgreifende Herausforderungen hinsichtlich thermischer Stabilität, Schaltzuverlässigkeit und Energieeffizienz mit sich. Mit schrumpfenden Bauteilabmessungen wird die Makro-Spin-Näherung unzureichend, was eine rigorose Behandlung der räumlich nicht-uniformen Magnetisierungsdynamik und Transportphänomene erfordert. Diese Arbeit präsentiert ein umfassendes computergestütztes Framework für die Modellierung und Optimierung von ultra-skalierten, mehrschichtigen MRAM-Bauteilen.

Zunächst wird ein dreidimensionaler mikromagnetischer Löser basierend auf der Finiten-Elemente-Methode (FEM) entwickelt. Um das rechenintensive Entmagnetisierungsfeld in komplexen Geometrien aufzulösen, wird eine hybride Finite-Elemente-Randelemente-Methode (FE-BEM) eingesetzt. Der numerischen Steifheit, die der Austauschwechselwirkung inhärent ist, wird durch die Implementierung eines *Tangentialebenen-Zeitintegrationsschemas* begegnet, welches die nichtlineare Landau-Lifshitz-Gilbert Gleichung als lineares Sattelpunktproblem neu formuliert. Dieser Ansatz wird gegen *adaptive Backward Differentiation Formulas (BDF) höherer Ordnung und implizit-explicit (IMEX) Methoden* verglichen und zeigt überlegene Stabilität und Effizienz für die steife Dynamik ultra-skalierten magnetischer Elemente.

Im Zentrum dieser Arbeit steht die Erweiterung des mikromagnetischen Frameworks um einen gekoppelten Spin- und Ladungs-Drift-Diffusions-Formalismus. Anders als Standardmodelle, die eine Kontinuität der Spinströme annehmen, leitet diese Arbeit spezialisierte Randbedingungen her und implementiert diese, um die Spinfiltration durch Magnesiumoxid (MgO)-Tunnelbarrieren und die Spindephasierung an metallischen Grenzflächen rigoros zu beschreiben. Darüber hinaus wird eine neuartige numerische Behandlung der Interlayer Exchange Coupling (IEC) eingeführt, die einen Algorithmus zur Grenzflächenabbildung nutzt. Dies ermöglicht die effiziente Simulation von Synthetic Antiferromagnets (SAF) ohne den prohibitiven Rechenaufwand einer volumetrischen Vernetzung von Abstandshaltern im Angström-Bereich.

Die Anwendung dieses vereinheitlichten Lösert zeigt, dass Zuverlässigkeitsprobleme in ultra-skalierten Bauteilen, wie etwa das „Back-Hopping“, deterministische Konsequenzen der Inter-Segment-Torque-Hierarchie innerhalb zusammengesetzter freier Schichten sind und nicht auf stochastische thermische Effekte zurückzuführen sind. Entscheidend ist der Nachweis, dass diese Instabilität durch gezielte Abstimmung der relativen Tunnelbarrieren-Polarisationen und der Geometrie der freien Schicht

entweder für zuverlässiges binäres Schalten unterdrückt oder gezielt ausgenutzt werden kann, um Multi-Level Cells (MLC) zu realisieren, die über unterschiedliche intermediäre Magnetisierungskonfigurationen mehrere Bits pro physischer Zelle speichern können. Zusätzlich werden Double Spin-Torque (dsMTJ)-Architekturen untersucht, bei denen die Synergie zwischen Giant Magnetoresistance (GMR) und Tunneling Magnetoresistance (TMR) Torques ein Schalten im Sub-Nanosekundenbereich ermöglicht. Wird Interlayer Exchange Coupling (IEC) in fortgeschrittenen Mehrschichtstrukturen eingeführt, erweist es sich als vereinheitlichender Designparameter: Spezifische Kopplungsfenster bestimmen die Stabilität SAF-verstärkter Referenzschichten, hybride freie Schichten mit metallischen Spacern nutzen die Spin-Flip-Länge des Spacer-Materials zur Verbesserung der Schaltgeschwindigkeit, und die gemeinsame Optimierung von Spacer-Material und -Dicke bestimmt die Schreibeffizienz.

Sažetak

Skaliranje Spin-Transfer Torque Magnetoresistive Random Access Memory (STT-MRAM) u područje ispod 20 nm donosi duboke izazove u pogledu termalne stabilnosti, pouzdanosti prebacivanja i energetske učinkovitosti. Kako se dimenzije uređaja smanjuju, aproksimacija makro-spina postaje nedostatna, što zahtijeva rigorozan tretman prostorno neuniformne dinamike magnetizacije i transportnih fenomena. Ova disertacija predstavlja sveobuhvatan računalni okvir za modeliranje i optimizaciju ultra-skaliranih, višeslojnih MRAM uređaja.

Prvo, razvijen je trodimenzionalni mikromagnetski rješavač temeljen na metodi konačnih elemenata (FEM). Za rješavanje računalno zahtjevnog polja demagnetizacije u složenim geometrijama koristi se hibridna metoda konačnih elemenata i rubnih elemenata (FE-BEM). Numerička krutost svojstvena interakciji izmjene rješava se implementacijom *scheme vremenske integracije u tangentnoj ravnini*, koja reformulira nelinearnu Landau–Lifshitz–Gilbertovu jednadžbu kao linearni problem sedlaste točke. Ovaj pristup uspoređen je s *adaptivnim formulama diferencijacije unatrag (BDF) višeg reda i implicitno-eksplicitnim (IMEX) metodama*, pokazujući superiornu stabilnost i učinkovitost za krutu dinamiku ultra-skaliranih magnetskih elemenata.

Središnji dio ovog rada je proširenje mikromagnetskog okvira uključivanjem spregnutog formalizma drifta i difuzije spina i naboja. Za razliku od standardnih modela koji pretpostavljaju kontinuitet spinskih struja, ovaj rad izvodi i implementira specijalizirane rubne uvjete za rigorozan opis filtriranja spina kroz tunelske barijere od magnezijevog oksida (MgO) i defaziranja spina na metalnim sučeljima. Nadalje, uveden je novi numerički tretman za Interlayer Exchange Coupling (IEC), koristeći algoritam mapiranja sučelja koji omogućuje učinkovitu simulaciju sintetičkih antiferomagneta (SAF) bez prohibitivnog računalnog troška volumetrijskog umrežavanja razmačnika na razini angstrema.

Primjena ovog objedinjenog rješavača otkriva da su problemi pouzdanosti u ultra-skaliranim uređajima, poput „back-hoppinga“, determinističke posljedice hijerarhije inter-segmentnih momenata unutar kompozitnih slobodnih slojeva, a ne stohastičkih termalnih efekata. Ključno je da ova teza pokazuje kako se preciznim podešavanjem relativnih polarizacija tunelskih barijera i geometrije slobodnog sloja ova nestabilnost može ili potisnuti za pouzdano binarno prebacivanje ili ciljano iskoristiti za realizaciju višerazinskih ćelija (MLC) sposobnih za pohranu više bitova po fizičkoj ćeliji putem različitih intermedijarnih konfiguracija magnetizacije. Dodatno, istražene su arhitekture Double Spin-Torque (dsMTJ), gdje sinergija između momenata divovskog

magnetootpora (GMR) i tunelskog magnetootpora (TMR) omogućuje prebacivanje ispod jedne nanosekunde. Kada se Interlayer Exchange Coupling (IEC) uvede u napredne višeslojne strukture, on postaje objedinjujući projektni parametar: specifični prozori vezanja određuju stabilnost SAF-pojačanih referentnih slojeva, hibridni slobodni slojevi s metalnim razmačnicima iskorištavaju duljinu spin-flipa materijala razmačnika za poboljšanje brzine prebacivanja, a zajednička optimizacija materijala i debljine razmačnika određuje učinkovitost zapisivanja.

Acknowledgement

First and foremost, I would like to express my deepest gratitude to my supervisor, Prof. Viktor Sverdlov, for the opportunity to pursue my doctorate under his guidance. I greatly value the freedom he gave me to explore different directions in my research. His genuine trust in letting me find my own path stood out. His unwavering support extended beyond academia. When I became a father, his understanding, patience, and encouragement meant more to me than I can put into words. I am deeply grateful for having a supervisor who cared not just about my scientific growth but also about my well-being. It has been a privilege to work with him within the Christian Doppler Laboratory for Nonvolatile Magnetoresistive Memory and Logic.

I would also like to sincerely thank Prof. Siegfried Selberherr, who founded the Institute for Microelectronics more than three decades ago, for his steady support of my work, for fostering a culture of open scientific exchange, especially through the weekly seminars, and for cultivating an excellent and inspiring working environment.

I am grateful for the opportunity to work at the Christian Doppler Laboratory for Nonvolatile Magnetoresistive Memory and Logic, which provided an outstanding research environment. My sincere thanks go to my colleagues Roberto, Johannes, Simone, Tomáš, Nils, Bernhard, and all other members of the laboratory for their collaboration, friendship, and the many stimulating discussions we shared. I would also like to thank Wolfgang Goes for the fruitful collaboration with our laboratory.

Furthermore, I thank all my colleagues at the Institute for Microelectronics for the warm atmosphere and many enriching conversations. In particular, I would like to thank Diana for her tireless help with organizational and administrative matters. I am also grateful to Manfred, Cerv, and the entire technical and administrative staff for ensuring the institute's infrastructure has always run smoothly.

My progress relied on a wide and invaluable network of family and friends. Without them, science would be far less meaningful and certainly far less enjoyable. I am deeply grateful to my parents and my entire family for their unconditional support at every turn throughout my life and studies.

My heartfelt thanks also go to our closest friends, Stjepan and Ana. Your steadfast support, trust, and friendship have accompanied us through every stage of this journey. Your presence has meant more than words can express.

Above all, I would like to express my deepest and most personal gratitude to my beloved wife, Gabrijela, for her boundless love, patience, and encouragement throughout this demanding journey. She supported me, especially during the long and

difficult months of the pandemic. Her strength, understanding, and unwavering belief in me made this work possible in ways that cannot be measured.

Finally, I dedicate this achievement to my daughter, Aurelia-Izabel. You are my greatest motivation and my greatest joy. I hope that one day you will look at this work and see not only a scientific accomplishment, but also a reminder that curiosity, perseverance, and kindness can go hand in hand. I strive to be a role model for you and to encourage you to follow your own path with confidence and courage.

Without the love and support of my family and friends, none of this would have been the same. I am deeply thankful to all of you.

Contents

Abstract	i
Kurzfassung	iii
Sažetak	v
Acknowledgement	vii
List of Figures	xiv
List of Tables	xxv
List of Abbreviations	xxvii
List of Symbols	xxx
1 Introduction	1
1.1 Outline of the Thesis	4
1.2 Research Setting	5
2 Spin-Transfer Torque Magnetic Random Access Memory	7
2.1 Introduction to Magnetoresistive Random Access Memory	7
2.1.1 Historical Context	8
2.2 Magnetoresistance Phenomena	8
2.2.1 Giant Magnetoresistance	8
2.2.2 Tunnel Magnetoresistance	10
2.3 Spin-Transfer Torque Switching	12
3 Micromagnetic Modelling	15
3.1 Landau-Lifshitz-Gilbert Equation	16
3.2 Effective Magnetic Field	18
3.2.1 Exchange Field	19
3.2.2 Zeeman Field	21
3.2.3 Anisotropy Field	21
3.2.3.1 Uniaxial Anisotropy	21

3.2.3.2	Planar Anisotropy	22
3.2.3.3	Cubic Anisotropy	22
3.2.4	Demagnetizing Field	24
3.2.5	Ampère Field	25
3.2.6	Thermal Field	25
3.2.7	Interlayer Exchange Coupling	26
3.3	Modelling Spin-Transfer Torque	28
4	Computational Methods	31
4.1	Finite Difference Method	32
4.1.1	Discretization of the Landau-Lifshitz-Gilbert Equation	36
4.1.1.1	External Field	36
4.1.1.2	Exchange Field	36
4.1.1.3	Anisotropy Field	36
4.1.1.4	Demagnetizing Field	37
4.1.1.5	Ampère Field	38
4.1.1.6	Thermal Field	38
4.1.1.7	Time Integration Scheme	39
4.2	Finite Element Method	39
4.2.1	Weak Formulation and Variational Form	40
4.2.2	Finite Element Discretization	42
4.3	Finite Element Micromagnetic Solver	45
4.3.1	Efficient Algorithms for Demagnetizing Field Computation	46
4.3.1.1	Truncation Approach	46
4.3.1.2	Hybrid Finite Element–Boundary Element Method	47
4.3.1.3	Verification and Benchmarking	49
5	Numerical Time Integration of the LLG Equation	55
5.1	Challenges in LLG Time Integration	55
5.1.1	Stiffness and Disparate Time Scales	55
5.1.2	The Unit-Sphere Constraint	57
5.1.3	Structural Properties	58
5.2	Classification of Time Integrators	58
5.2.1	Explicit Runge–Kutta Methods	58
5.2.2	Implicit Methods and A-Stability	59
5.2.3	Multistep Methods: BDF Schemes	60
5.2.4	Implicit-Explicit Method	61
5.3	The Tangent-Plane Scheme	61
5.3.1	Reformulation of the LLG Equation	62
5.3.2	Weak Formulation	62
5.3.3	Complete Weak Form with Spin-Transfer Torque	62
5.3.4	Time Discretization: The θ -Method	63
5.3.5	The Saddle-Point System	63
5.3.6	Constraint Enforcement and Magnetization Update	64
5.4	Adaptive Multistep Methods via <i>SUNDIALS</i>	64

5.4.1	CVODE: Variable-Order BDF	65
5.4.2	ARKStep: IMEX Runge–Kutta	70
5.5	Verification and Validation	72
5.5.1	The μ MAG Standard Problem 4	72
5.5.2	Solver Comparison on Standard Problem 4	73
5.5.3	Constraint Preservation and Solver Accuracy	76
5.5.4	Adaptive Time Stepping	77
5.6	Summary	78
6	Spin-Transfer Torques in Ultra-Scaled MRAM Cells	81
6.1	Spin and Charge Drift-Diffusion Formalism	82
6.1.1	Derivation of the Transport Equations	82
6.1.2	Three-Dimensional Formulation	83
6.1.3	Spin-Accumulation Equation	83
6.1.4	Spin Dephasing and Complete Transport Equations	84
6.2	Finite Element Implementation	85
6.2.1	Coupled LLG and Drift-Diffusion Framework	85
6.2.2	Treatment of the Tunnel Barrier	86
6.3	Tunneling Spin-Current Boundary Conditions	86
6.3.1	Tunneling Current Expressions	87
6.3.2	Finite Element Boundary Terms	87
6.3.3	Implementation of Interface Mapping	88
6.4	Numerical Implementation of Interlayer Exchange Coupling	88
6.4.1	IEC Energy and Effective Field	88
6.4.2	IEC Boundary Term in Weak Formulation	89
6.4.3	Mapping Strategy for IEC	89
6.5	Validation Results	90
6.5.1	Spin-Accumulation and Spin Torque Solutions	90
6.5.2	Effect of Dephasing Length on Torque Localization	90
6.6	GMR Effect in Spin-Valves	91
6.7	Torques in Elongated Structures	93
6.7.1	Comparison with Drift-Diffusion Solution	94
6.7.2	Interaction Between Tunneling and Bulk Torques	95
6.8	Preconditioner Evaluation	96
6.9	Conclusion	101
7	Advanced Switching Dynamics and Novel Device Structures	103
7.1	Introduction	103
7.2	Computational Setup	103
7.3	Ultra-Scaled MRAM Cell Switching	105
7.3.1	Device Architecture and Geometry	105
7.3.2	Elongated FL: Domain Wall Failure	107
7.3.3	Composite FL Solution	107
7.3.4	Switching Mechanisms in Composite FLs	108
7.3.5	FL Length Optimization	109

7.3.6	Material Parameter Robustness	112
7.3.7	Switching Characteristics in Asymmetric FLs	113
7.4	Back-Hopping and Multi-Level Cell Operation	115
7.4.1	Torque Analysis	115
7.4.1.1	AP→P: Cooperative Torques	116
7.4.1.2	P→AP: Competing Torques and Back-Hopping Origin	117
7.4.2	Cyclic Switching Visualization	120
7.4.3	Multi-Level Cell Realization	121
7.4.4	FL Thickness Variation	122
7.4.5	Voltage Polarity Robustness	123
7.5	Interlayer Exchange Coupling Effects	125
7.5.1	Interlayer Exchange Coupling as Control Parameter	125
7.5.2	Systematic IEC Sweep Across Three Stack Configurations	127
7.6	Advanced Structures	132
7.6.1	SAF Architecture and Operating Principle	132
7.6.2	Reliability Requirements: The Critical Role of Coupling Strength	133
7.6.3	Spin-Torque Analysis and IEC Design Criterion	134
7.6.4	Hybrid FL with Metallic Spacers	134
7.6.5	Four-Stage Torque Decomposition	135
7.6.6	Write Efficiency Analysis	137
7.6.7	Device Concept and Dual Torque Mechanism	138
7.6.8	Influence of Spacer Material	138
7.6.9	Role of Interlayer Exchange Coupling	139
7.6.10	SAF-Enhanced Double Spin-Torque MTJ	140
7.7	Synthesis	144
7.7.1	Design Guidelines for Ultra-Scaled Composite FL Structures	144
7.7.2	Design Guidelines for Multi-Level Cell Implementation	144
7.7.3	Design Guidelines for ds-MTJ and SAF-Enhanced Structures	145
7.7.4	Cross-Architecture Performance Comparison	145
8	Summary and Outlook	147
8.1	Summary	147
8.1.1	Micromagnetic Solver and Time Integration	147
8.1.2	Coupled Spin and Charge Drift-Diffusion Transport	148
8.1.3	Interlayer Exchange Coupling and Interface-Mapping	148
8.1.4	Application to Advanced MRAM Architectures	148
8.2	Outlook: Advanced Numerical Methods for Micromagnetics	149
8.2.1	Implicit Treatment of Spin-Transfer Torque in IMEX Schemes	150
8.2.2	Higher-Order Finite Elements for Exchange Resolution	150
8.2.3	Parallel-in-Time Integration	150
8.2.4	Unified Space-Time Finite Element Methods	151
8.3	Conclusion	151
	Bibliography	153

List of Publications**171**

List of Figures

- 1.1 Schematics of different existing memory types (gray pyramid) and prospective spintronic technologies based on MTJs that could replace them (colored pyramid). L1-L4 denote various cache levels; SRAM and DRAM represent static and dynamic random-access memory, respectively; and HDD represents magnetic hard disk drives. Mechanisms for writing bit information in MTJ devices include the Oersted field, SOT-MRAM, STT-MRAM, and DW motion. 2
- 2.1 Illustration of spin-dependent electron scattering in a metallic multilayer structure demonstrating the GMR effect. Red arrows represent *spin-up* electrons, blue arrows denote *spin-down* electrons, and black arrows indicate the magnetization direction. In the AP configuration (left), both spin channels experience greater scattering, leading to higher resistance. In the P configuration (right), majority-spin electrons encounter less scattering, resulting in lower resistance. 9
- 2.2 Illustration of spin-dependent tunneling responsible for the TMR effect. In the P configuration (left), spin-polarized electrons tunnel into available states of the same spin orientation in the second ferromagnetic layer, resulting in high conductance. In the AP configuration (right), spin mismatch reduces the tunneling probability, leading to low conductance. 11
- 2.3 Illustration of magnetization orientations in MTJs, showing the difference between in-plane (iMTJ) and perpendicular (pMTJ) configurations. 13
- 3.1 Magnetization dynamics on the unit sphere. (a) Pure precession: $\mathbf{m}$ rotates around $\mathbf{H}_{\text{eff}}$ at constant polar angle (no damping, $\alpha = 0$). (b) Pure damping: $\mathbf{m}$ relaxes directly toward $\mathbf{H}_{\text{eff}}$ along a meridian (no precession). (c) Combined motion: damped precession spirals toward equilibrium. 17
- 3.2 Anisotropy functions visualized as energy surfaces. (a) Uniaxial anisotropy with $K_u = 1$ and $\mathbf{a} = \mathbf{e}_3$. (b) Planar anisotropy with $K_p = 1$ and $\mathbf{a} = \mathbf{e}_3$. (c) Cubic anisotropy with $K_{c1} = K_{c2} = 1$, and $\mathbf{a}_i = \mathbf{e}_i$ for all $1 \leq i \leq 3$. 23

3.3	Schematic depiction of a trilayer structure composed of left and right semi-infinite FM regions, separated by a NMS layer. The interface magnetization on the left FM interface points into an arbitrary direction, whereas the interface magnetization in the right FM forms an angle $\Delta\phi$ relative to the magnetization of the left FM. The magnetization vectors originate at the respective interface regions $\partial\Omega_1$ and $\partial\Omega_2$, which are highlighted by hatching. Figure based on the figure originally published in [116].	28
4.1	Discretization of a unit sphere using two common methods in micromagnetics: (a) structured Cartesian mesh used in FDM, and (b) unstructured tetrahedral mesh used in FEM.	32
4.2	Mesh points for the FDM grid in one (a) and two (b) spatial dimensions. The grid spacing is uniform with h between adjacent points. In (b), the 2D indexing convention for a 3×3 grid is shown.	33
4.3	Representation of the original solution u and the finite element approximation u_h in one (a) and two (b) spatial dimensions. The basis functions for all nodes are shown at the bottom of the 1D plot. The highlighted basis function $\varphi_2(x)$ and corresponding solution value u_2 at node x_2 are shown in orange. The 2D plot shows the smooth function $u(x, y)$, its piecewise approximation $u_h(x, y)$, and the basis function $\varphi_{5,6}(x, y)$ centered at node (x_5, y_6)	44
4.4	Flowcharts of the micromagnetic simulation implementations.	45
4.5	Finite-element mesh used for computing the demagnetization field with the truncation approach. (a) The magnetic region is shown in blue, while the surrounding external region is shown in light blue. The external domain is chosen to be approximately five times larger than the magnetic region in each spatial direction. (b) The magnetic region represents the area of interest and is embedded within the larger external domain used to model open boundary conditions.	48
4.6	Comparison of demagnetizing energy for three distinct magnetization states: (a) uniform, (b) flower, and (c) vortex. For each case, the top subfigures show the corresponding magnetic potential, and the bottom subfigures show the magnetization distribution.	50
4.7	Comparison of demagnetizing energy for unit cubes with different magnetization configurations (uniform, flower, vortex), calculated using the truncation approach and the hybrid FE-BEM approach against analytical reference values [153]. The hybrid FE-BEM shows excellent agreement with analytical results, while the truncation approach significantly underestimates the energy. Based on the figure originally published in [154].	51

- 4.8 Two-disc configuration with parallel magnetizations. Left: demagnetizing field $H_{d,x}$ along the symmetry axis. Right: relative error with respect to the analytical reference (4.84). The FE-BEM maintains accuracy across the axis, while the truncation approach shows significant deviations. 52
- 4.9 Two-disc configuration with anti-parallel magnetizations. Left: demagnetizing field $H_{d,x}$ along the symmetry axis. Right: relative error with respect to the analytical reference. The FE-BEM reproduces the analytical field within a few percent, while the truncation approach yields larger errors. 52
- 4.10 Demagnetizing field in an STT-MRAM cell with P magnetization configuration. Arrows indicate the demagnetization field direction, while color-coding represents the magnetic scalar potential u . Left: hybrid FE-BEM. Right: truncation approach. Both methods yield similar results, validating the implementation for disconnected magnetic domains. Figure originally published in [154]. 54
- 4.11 Demagnetizing field in an STT-MRAM cell with AP magnetization configuration. Arrows indicate the demagnetization field direction, while color-coding represents the magnetic scalar potential u . Left: hybrid FE-BEM. Right: truncation approach. Both methods yield similar results, validating the implementation for disconnected magnetic domains. Figure originally published in [154]. 54
- 5.1 Hierarchy of time scales in micromagnetic simulations. The exchange interaction imposes the CFL stability limit ($\tau_{\text{exch}} \sim 10^{-12}$ s for $h = 2$ nm mesh). Demagnetization dynamics occur at $\tau_{\text{demag}} \sim 10^{-11}$ s. Precession due to external fields operates at $\tau_{\text{prec}} \sim 10^{-9}$ s. Damping and switching occur at $10^{-9} - 10^{-8}$ s, far above the explicit stability limit. Material: Permalloy ($A = 1.3 \times 10^{-11}$ J/m, $M_s = 8 \times 10^5$ A/m, $\alpha = 0.02$). . . . 57
- 5.2 Stability regions in the complex plane $z = \lambda \Delta t$ with physical eigenvalues. The *Physics of Interest* (Precession, Damping) lies near the origin and could be stable for explicit methods. The exchange eigenvalue position depends on the time step: for $\Delta t = 10^{-13}$ s, $\lambda_{\text{exch}} \Delta t \approx -1$ lies within the RK4 stability region (black dot). Increasing Δt moves it left (gray dot), eventually exiting explicit stability regions. Explicit solvers (Red/Blue) are forced to take tiny time steps ($\Delta t \rightarrow 0$) just to keep the black dot inside their stability boundary. The Implicit solver (Green) remains stable for any Δt , allowing us to step over the stiff exchange modes. IMEX methods combine both approaches: exchange modes are treated implicitly (Green), while precession and damping remain explicit (Red/Blue). 60

5.3	Block structure of the tangent-plane scheme saddle-point system corresponding to (5.31). The (1,1)-block A combines the damping, precession, and implicit exchange terms. The off-diagonal blocks B and B^T enforce the tangent-plane constraint $\mathbf{v} \cdot \mathbf{m} = 0$, while the solution vector contains both the velocity $\mathbf{v}$ and Lagrange multipliers $\boldsymbol{\lambda}$	65
5.4	Setup for the μ MAG Standard Problem 4. (a) Permalloy film geometry ($500 \times 125 \times 3 \text{ nm}^3$) discretized by an unstructured tetrahedral mesh with 45391 elements. (b) Initial S -state magnetization configuration showing the characteristic S pattern in the xy -plane, colored by the x -component of magnetization.	73
5.5	Magnetization dynamics for Standard Problem 4 (Field 1) computed with different time integrators. All solvers agree with the μ MAG reference solutions [174, 175].	74
5.6	Deviation of computed magnetization trajectories from the μ MAG reference solutions: (left) d'Aquino <i>et al.</i> [174], (right) Albuquerque <i>et al.</i> [175]. All solvers exhibit comparable deviation from both references, with differences attributable to the finite element discretization rather than time integration errors.	75
5.7	(a) Deviation from the unit-sphere constraint $ \mathbf{m} = 1$ before renormalization. (b) Difference between solver trajectories and the TPS reference solution. (c) Time step evolution for adaptive solvers (BDF, IMEX) compared to the fixed time step used by TPS and RK4.	77
6.1	(Left) Current density distribution in an MTJ biased under a constant voltage for a non-uniform magnetization configuration of the FL. The magnetization varies from parallel to the reference layer in the center to anti-parallel on the sides. (Right) The current density is larger in the center, where the FL magnetization is parallel to that of the RL and the magnetization-dependent conductivity is highest. Based on the figure originally published in [182].	86
6.2	Validation of the spin drift-diffusion model with tunneling spin-current boundary conditions for semi-infinite ferromagnetic layers. (a) spin-accumulation and (b) spin-transfer torque components computed along the stack direction, comparing the analytical solution with the FE implementation. The magnetization is oriented along x in the RL and along z in the FL. (c) Angular dependence of the damping-like torque for several values of the RL spin polarization ($P_{\text{RL}} = 0.4, 0.7, 0.9$), reproducing the expected sinusoidal behavior characteristic of MTJs. Figure (c) adapted from [182].	90
6.3	Spin-accumulation and torque in a symmetric MTJ structure including NM contacts. (a) and (b) present results for $\lambda_{\varphi} = 0.4 \text{ nm}$. The shorter dephasing length guarantees faster decay of the transverse spin-accumulation components, so that the torque acts only in the proximity of the TB interface. Adapted from the figure [182].	92

6.4	Angular dependence of the charge current for (a) an MTJ with MgO tunnel barrier at two TMR ratios ($V = -1.3$ V) and (b) a spin valve with metallic spacer at two GMR ratios ($V = -0.2$ V). Markers represent simulation data and solid/dashed lines show analytical fits using Equation (2.7) and Equation (6.34), respectively.	93
6.5	(a) Non-uniform magnetization texture with the magnetization orientation changing from x to $-z$. (b) Comparison of the spin torque $\mathbf{T}_S$ to the ZL torque $\mathbf{T}_{ZL}$ for an approximately 100 nm long magnetization texture, with the parameters of Table 6.1. (c) Comparison for an approximately 3 nm long magnetization texture with $\lambda_\varphi = 0.4$ nm. The shorter dephasing length quickly dissipates the transverse spin-accumulation, ensuring agreement between the ZL approximation and the drift-diffusion solution. The two approaches show good agreement in both cases.	94
6.6	(a) Torque computed for an MRAM cell with elongated RL and FL and a magnetization profile in the FL similar to that of Figure 6.5 (a), with a width of approximately 3 nm. The brown vectors indicate the magnetization direction in the RL and in two parts of the FL. (b) Close-up of the spin torque $\mathbf{T}_S$ compared to the ZL torque $\mathbf{T}_{ZL}$. The presence of the MTJ also influences the bulk torque component, making the unified approach the most suitable for ultra-scaled MTJs with elongated ferromagnetic layers. Figure (a) originally published in [182].	95
6.7	Spatially-averaged magnetization dynamics during STT-driven switching computed with different time integrators. The TPS solution with $\Delta t = 0.01$ ps serves as the reference.	97
6.8	Quantitative solver analysis: (a) deviation from the unit-sphere constraint before renormalization; (b) trajectory deviation from the TPS reference solution; (c) time step evolution for the adaptive solvers compared to the fixed $\Delta t = 0.01$ ps baseline.	98
7.1	Three-dimensional renderings of (a)–(b) conventional large-diameter MTJ cells and (c)–(e) ultra-scaled 2.3 nm diameter pillar geometries with different FL configurations: elongated single FL, composite two-segment FL with middle TB, and composite FL with non-magnetic metallic spacer [182].	106
7.2	Magnetization dynamics in an ultra-scaled cell with a 15 nm elongated single FL (cf. Figure 7.1 (c)) at ± 2 V bias. (a), (b) Time evolution of magnetization components for AP to P and P to AP initial configurations showing failed switching. (c) Three-dimensional magnetization snapshot at 2 ns revealing the domain wall that prevents reversal. Figures based on [182].	108

- 7.3 Successful switching of composite FL structures at -2 V bias: (a) two 5 nm FL segments separated by a 0.9 nm MgO TB (cf. Figure 7.1 (d)), reaching switching after approximately 2 ns, (b) two 5 nm FL segments separated by a 2 nm non-magnetic metallic spacer (cf. Figure 7.1 (e)), achieving switching after approximately 1 ns. Figures based on [182]. 109
- 7.4 Three-dimensional magnetization snapshots of the sequential switching process in ultra-scaled composite FL cells (cf. Figure 7.1 (d)), showing initial, intermediate (plateau), and final states for (a) AP $\rightarrow$ P and (b) P $\rightarrow$ AP transitions. The switching order is direction-dependent: FL₁ switches first in AP $\rightarrow$ P, while FL₂ switches first in P $\rightarrow$ AP, reflecting the asymmetric torque configurations at each interface [123]. 109
- 7.5 Magnetization trajectories during (a) AP $\rightarrow$ P and (b) P $\rightarrow$ AP switching for FL segment lengths of 2, 3, 4, 5, and 7.5 nm with standard polarization ($P_1 = 0.62$, $P_2 = 0.57$, $P_3 = 0.2$) and Set A material parameters (Table 7.2). Critical design constraints emerge: segments shorter than approximately 5 nm exhibit back-hopping instability during P $\rightarrow$ AP switching, because the anisotropy barrier scales with segment length ($\mathcal{E}_B \propto K_{\text{eff}} \cdot l$) while the interfacial spin-transfer torque is approximately length-independent, so shorter segments satisfy the instability condition (7.1) more readily. Segments longer than approximately 6 nm fail to complete switching at moderate bias. The optimal range of 3 nm to 5 nm provides reliable switching in both directions. Based on Figure [193]. 110
- 7.6 Same configuration as Figure 7.5 but with Set B material parameters (Table 7.2), confirming robustness of the FL length constraints across material parameter sets. Based on Figure [194]. 111
- 7.7 Magnetization trajectories for FL segment lengths of 2, 3, 4, and 5 nm with enhanced TB₂ polarization ($P_1 = 0.5$, $P_2 = 0.9$, $P_3 = 0.2$) and Set A material parameters. The enhanced TB₂ polarization modifies the back-hopping threshold compared to the standard case (Figure 7.5). Based on Figure [193]. 111
- 7.8 Same enhanced TB₂ polarization as Figure 7.7 but with Set B material parameters, demonstrating that the polarization-dependent switching behavior persists across material choices. Based on Figure [194]. . . . 112
- 7.9 Magnetization trajectories with standard polarization ($P_1 = 0.62$, $P_2 = 0.57$, $P_3 = 0.2$) and asymmetric FL segment lengths for (a) AP $\rightarrow$ P and (b) P $\rightarrow$ AP switching across FL lengths of 2 nm to 5 nm. Based on Figure [188]. 114
- 7.10 Magnetization trajectories with enhanced TB₂ polarization ($P_1 = 0.5$, $P_2 = 0.9$, $P_3 = 0.2$) and asymmetric FL segment lengths for (a) AP $\rightarrow$ P and (b) P $\rightarrow$ AP switching across FL lengths of 2 nm to 5 nm. The enhanced TB₂ polarization intensifies the cyclic transitions compared to the standard polarization case (Figure 7.9), confirming deterministic back-hopping. Based on Figure [188]. 114

- 7.11 Spin-transfer torque heatmap showing $T_{S,x}$ (top) and $T_{S,z}$ (bottom) as a function of stack position and magnetization angle for standard polarization ($P_1 = 0.62$, $P_2 = 0.57$, $P_3 = 0.2$). Left column: FL₂ rotation, right column: FL₁ rotation. Contour lines indicate torque magnitude within ferromagnetic layers. 116
- 7.12 Spin-transfer torque heatmap for enhanced TB₂ polarization ($P_1 = 0.5$, $P_2 = 0.9$, $P_3 = 0.2$). Same layout as Figure 7.11. The enhanced polarization introduces a pronounced torque imbalance: $T_{S,x}$ on FL₁ is significantly enhanced by the high $P_2 = 0.9$, while the torque on FL₂ is reduced by the lower $P_3 = 0.2$ 118
- 7.13 Three-dimensional surface representation of $T_{S,x}$ (a, b) and $T_{S,z}$ (c, d) for standard polarization ($P_1 = 0.62$, $P_2 = 0.57$, $P_3 = 0.2$). Panels (a, c): FL₂ rotation, (b, d): FL₁ rotation. The in-plane torque $T_{S,x}$ is negative throughout the angular range with a minimum near 90°. A Gaussian filter suppresses numerical noise from the FE mesh. 119
- 7.14 Three-dimensional torque surfaces for enhanced TB₂ polarization ($P_1 = 0.5$, $P_2 = 0.9$, $P_3 = 0.2$). Compared to the standard case (Figure 7.13), $T_{S,x}$ on FL₁ (b) reaches approximately twice the magnitude of FL₂ (a), reflecting the $P_2/P_3 = 0.9/0.2$ ratio. This imbalance drives sequential switching and back-hopping. 120
- 7.15 Decomposition of damping-like and field-like spin-transfer torques on FL₁ and FL₂ during (a) AP→P and (b) P→AP switching. The torque hierarchy between TB₁ and TB₂ governs the sequential switching order and determines back-hopping onset [199]. 121
- 7.16 Cyclic switching visualization during P→AP transition: three-dimensional magnetization snapshots at five characteristic stages (i)–(vi) of the FL reversal cycle, colored by the m_x magnetization component. Based on [199]. 122
- 7.17 Resistance (purple, left axis) and average magnetization m_x (orange dashed, right axis) as a function of time for four pulse durations: (a) $t_{\text{pulse}} = 1.0$ ns, (b) 1.3 ns, (c) 2.0 ns, and (d) 3.0 ns. The vertical dotted line marks pulse termination t_{off} . The four resistance levels are clearly distinguishable with >20% separation between adjacent states. Based on Figure [194]. 123
- 7.18 Average magnetization $\langle m_x \rangle$ versus time for P→AP switching at 1.5 V bias with write pulse durations ranging from 0 to 3.5 ns (color-coded from red for short pulses to blue for long pulses). By terminating the write pulse at different times, distinct final magnetization configurations are accessed: short pulses ($t_{\text{pulse}} \lesssim 0.8$ ns) return the composite FL to its initial state, intermediate pulses stabilize partial-switching configurations, and the longest pulses ($t_{\text{pulse}} > 2.5$ ns) drive a complete back-hopping cycle. Based on [199]. 124

- 7.19 Magnetization response under reversed voltage polarity (opposite to the polarity required for switching): (a) AP state with positive voltage, stable at $m_x \approx -1$ for $V \leq 2.5$ V, but cyclic oscillations between $m_x = -1$ and $+1$ emerge at $V \geq 3.0$ V with increasing frequency, (b) P state with negative voltage, stable at $m_x = +1$ across the entire range (2.0 V to 4.5 V). The asymmetry arises from the cooperative torque configuration of the P state: both TB_1 and TB_2 torques act to stabilize the parallel alignment (see Section 7.4.1), whereas the AP state suffers from competing torques that enable destabilization at sufficient voltage. 124
- 7.20 Influence of IEC strengths $J_{iec,1}$ and $J_{iec,2}$ on the AP→P switching dynamics in the ultra-scaled composite FL with MgO TB separator (Figure 7.1 (d)). The IEC is swept at the FL_1 – FL_2 interface through the middle MgO barrier: (a) negative $J_{iec,1}$ (AFM), (b) $J_{iec,1} = 0$, and (c) positive $J_{iec,1}$ (FM), each producing distinct switching dynamics. Based on Figures [203]. 126
- 7.21 Systematic IEC sweep for the two ultra-scaled composite FL structures from [208]: (a–c) AP→P and (d–f) P→AP transitions for the composite FL with MgO TB separator (Figure 7.1 (d)), and (g–i) the same analysis for the composite FL with NMS separator (Figure 7.1 (e)). In each row, the left panel applies negative $J_{iec,1}$ (AFM), the middle panel sets $J_{iec,1} = 0$, and the right panel applies positive $J_{iec,1}$ (FM). The coupling strength range reflects experimentally reported values for CoFeB|MgO interfaces [204,205] and metallic spacers [206,207]. Based on Figures [208]. 127
- 7.22 Layer stacks of the three configurations studied in [116]: (a) Stack A (ultra-scaled composite FL, same architecture as Figure 7.1 (d)), (b) Stack B (SAF without PL), and (c) Stack C (SAF with PL). IEC acts at different interfaces in each stack: FL_1 – FL_2 coupling through TB_2 in Stack A, HL–RL coupling through Ru in Stack B, and both HL–RL (AFM) and RL–PL (FM) coupling in Stack C. 128
- 7.23 Switching trajectories during P→AP transition for the three stack configurations: (a) Stack A (ultra-scaled composite FL), (b) Stack B (SAF without PL), and (c) Stack C (SAF with PL). Weak FM coupling suppresses back-hopping in Stack A, strong AFM coupling maintains RL stability in Stack B, and PL reversal in Stack C depends on the RL–PL coupling strength [116]. 129
- 7.24 Spin-torque distribution at the moment of back-hopping onset for (a) Stack A, (b) Stack B, and (c) Stack C. The field-like component $T_{S,x}$ and damping-like component $T_{S,z}$ reveal distinct torque patterns at each ferromagnetic interface, with IEC modifying both the magnitude and spatial distribution of the torques [116]. 129

- 7.25 Three-dimensional magnetization states during P→AP switching in Stack B (SAF without PL), contrasting weak ($J_{\text{iec}} = -0.5 \text{ mJ m}^{-2}$) and strong ($J_{\text{iec}} = -1.32 \text{ mJ m}^{-2}$) AFM HL–RL coupling. (a)–(c) weak coupling, full view and (d)–(f) weak coupling, FL clipped to reveal RL, (g)–(i) strong coupling, full view and (j)–(l) strong coupling, FL clipped to reveal RL. Columns show the initial state, mid-switching ($\sim 0.4 \text{ ns}$), and final state ($\sim 1 \text{ ns}$). With strong coupling, the RL remains uniformly oriented throughout switching, with weak coupling, the RL develops domain walls and its magnetization reverses almost completely [215]. 130
- 7.26 Three-dimensional magnetization states during P→AP switching in Stack C (SAF with PL), contrasting weak ($J_{\text{iec}} = 0.2 \text{ mJ m}^{-2}$) and strong ($J_{\text{iec}} = 0.8 \text{ mJ m}^{-2}$) FM RL–PL coupling. (a)–(c) weak coupling, full view and (d)–(f) weak coupling, FL clipped to reveal PL, (g)–(i) strong coupling, full view and (j)–(l) strong coupling, FL clipped to reveal PL. Columns show the initial state, mid-switching ($\sim 0.4 \text{ ns}$), and final state ($\sim 1 \text{ ns}$). With strong coupling, the PL remains stable, with weak coupling, the PL reverses almost completely due to parasitic spin-torques overcoming the IEC [215]. 131
- 7.27 Layer stack of the SAF-enhanced composite FL configuration from [216]. Stack D and Stack E share the same architecture, differing only in the NMS material between the FL segments: tungsten (W, $\lambda_{sf} = 2.4 \text{ nm}$) for Stack D, tantalum (Ta, $\lambda_{sf} = 1.9 \text{ nm}$) for Stack E. IEC acts at the HL–RL interface (through Ru, AFM coupling) and at the FL₁–FL₂ interface (through the metallic spacer, FM coupling). Stack C employs a single FL (same architecture as Figure 7.22 (c)). 132
- 7.28 Switching duration with different FM coupling for three stack configurations: (a) Stack C with single FL at strong (0.8 mJ m^{-2}) and weak (0.2 mJ m^{-2}) FM coupling, (b) Stack D with W spacer composite FL at $J_{\text{iec}} = 0.37 \text{ mJ m}^{-2}$ and 0.8 mJ m^{-2} , and (c) Stack E with Ta spacer composite FL at $J_{\text{iec}} = 0.2 \text{ mJ m}^{-2}$ and 0.8 mJ m^{-2} . Weaker coupling delays switching in all three configurations, with the Ta spacer (c) exhibiting the longest switching times [216]. 133
- 7.29 Three-dimensional magnetization states during P→AP switching in Stack D (W spacer composite FL), contrasting weak ($J_{\text{iec}} = 0.37 \text{ mJ m}^{-2}$) and strong ($J_{\text{iec}} = 0.8 \text{ mJ m}^{-2}$) FM FL₁–FL₂ coupling. (a)–(c) weak coupling, full view and (d)–(f) weak coupling, FL₂ clipped to reveal FL₁, (g)–(i) strong coupling, full view and (j)–(l) strong coupling, FL₂ clipped to reveal FL₁. Rows show the initial state, mid-switching ($\sim 0.5 \text{ ns}$), and final state ($\sim 1.7 \text{ ns}$). With weak coupling, FL₁ reverses first producing a transient plateau at $\langle m_x \rangle \approx 0.5$, with strong coupling, synchronized reversal eliminates the plateau [216]. 135

- 7.30 Three-dimensional magnetization states during $P \rightarrow AP$ switching in Stack E (Ta spacer composite FL), contrasting weak ($J_{\text{iec}} = 0.2 \text{ mJ m}^{-2}$) and strong ($J_{\text{iec}} = 0.8 \text{ mJ m}^{-2}$) FM FL_1 – FL_2 coupling. (a)–(c) weak coupling, full view and (d)–(f) weak coupling, FL_2 clipped to reveal FL_1 , (g)–(i) strong coupling, full view and (j)–(l) strong coupling, FL_2 clipped to reveal FL_1 . Rows show the initial state, mid-switching ($\sim 0.75 \text{ ns}$), and final state ($\sim 1.7 \text{ ns}$). The shorter spin-flip length of Ta (1.9 nm) compared to W (2.4 nm) produces weaker inter-segment torques, eliminating the transient plateau and slowing overall switching [216]. 136
- 7.31 Four-stage torque decomposition during $P \rightarrow AP$ switching for (a)–(d) Stack C with single FL, (e)–(h) Stack D with W spacer composite FL, and (i)–(l) Stack E with Ta spacer composite FL. Rows show the spin-transfer torque distribution at the initial P state ($\theta \approx 0^\circ$), $\theta = 45^\circ$, $\theta = 135^\circ$, and the final AP state ($\theta \approx 180^\circ$). The composite FL stacks show additional torque peaks at the NMS interfaces absent in the single-FL configuration [216]. 137
- 7.32 Layer stack of the double spin-torque MTJ (ds-MTJ) architecture, in which the FL is sandwiched between two reference layers with anti-parallel magnetizations separated by a TB and a non-magnetic metallic spacer. The anti-parallel RL configuration produces cooperative torques from both interfaces that constructively add, and the NMS material choice (Ru vs. Ta) controls the secondary torque magnitude via its spin-flip length [218]. 139
- 7.33 Switching performance of ds-MTJ vs. conventional pMTJ structures, showing (a) the switching time ratio and (b) absolute switching times for both transition directions. The ds-MTJ with Ru-NMS achieves a $2\text{--}3\times$ speedup and sub-nanosecond switching at moderate voltages, while Ta-NMS provides a $1.2\text{--}1.5\times$ enhancement, consistent with the ratio of their spin-flip lengths (4 nm vs. 1.9 nm) [218]. 140
- 7.34 Influence of IEC on ds-MTJ switching dynamics with Ru-NMS, exploring a 3×5 parameter matrix of AFM, zero, and FM coupling through TB_1 (columns) with varying NMS coupling (curves within each panel) for (a) $AP \rightarrow P$ and (b) $P \rightarrow AP$ transitions. Zero coupling ($J_{\text{iec},1} = J_{\text{iec},2} = 0$) generally yields the fastest switching, while strong coupling can introduce energy landscape distortions that trap the magnetization in unintended states [218]. 141
- 7.35 Influence of IEC on ds-MTJ switching dynamics with Ta-NMS, exploring a 3×5 parameter matrix of AFM, zero, and FM coupling through TB_1 (columns) with varying NMS coupling (curves within each panel) for (a) $AP \rightarrow P$ and (b) $P \rightarrow AP$ transitions. Zero coupling ($J_{\text{iec},1} = J_{\text{iec},2} = 0$) generally yields the fastest switching, while strong coupling can introduce energy landscape distortions that trap the magnetization in unintended states [218]. 142

- 7.36 Switching trajectories for SAF-enhanced ds-MTJ structures comparing Ru and Ta NMS at varying thicknesses: (a) P→AP and (b) AP→P transitions. Optimal performance at Ru 1.0 nm ($J_{\text{iec}} = -0.65 \text{ mJ m}^{-2}$, AFM) and Ta 0.25 nm ($J_{\text{iec}} = 0.4 \text{ mJ m}^{-2}$, FM). Ru at 0.4 nm (-2.1 mJ m^{-2}) exhibits oscillatory behavior from excessively strong AFM coupling [217]. 143
- 7.37 Spin-transfer torque distribution across the full multilayer stack for (a) single SAF ds-MTJ and (b) double SAF ds-MTJ, comparing Ru (AFM) and Ta (FM) NMS. The Ru spacer generates stronger torques on RL_2 due to its longer spin-flip length, while the double SAF configuration provides enhanced reference layer stability with minimal switching speed penalty [217]. 144

List of Tables

5.1	Comparison of preconditioner strategies for the C-language Variable-coefficient Ordinary Differential Equation solver (CVODE) Newton system.	69
5.2	Summary of implemented time integrators.	72
5.3	Quantitative comparison of solvers on μ MAG Standard Problem 4. Errors are computed against both reference solutions.	75
6.1	Parameters used in the drift-diffusion simulations with the tunnel barrier (TB) boundary terms.	91
6.2	Quantitative performance comparison for spin-transfer torque (STT)-driven switching. The trajectory error is measured by e_∞ relative to the tangent-plane scheme (TPS) reference (baseline row). Speedup is computed as $\text{Speedup} = \text{CPU}_{\text{TPS,ref}}/\text{CPU}_{\text{method}}$	99
7.1	Standard simulation parameters for the switching studies presented in this chapter. Values are representative of CoFeB MgO-based magnetic tunnel junctions and are sourced from the Boris Computational Spintronics material database and published experimental literature. .	104
7.2	Comparison of material parameter sets used to verify robustness of the back-hopping mechanism.	113
7.3	Multi-level cell state definitions and characteristics.	121
7.4	Material and geometry parameters for the three MDPI 2024 stack configurations [116] and the MME 2025 synthetic antiferromagnet (SAF)-enhanced stacks [216]. Common parameters are listed in Table 7.1.	131
7.5	Layer parameters of the double-spin torque magnetic tunnel junction (ds-MTJ) architecture [218]. The free layer (FL) is sandwiched between RL ₁ (MgO tunnel barrier) and RL ₂ (metallic spacer). Two nonmagnetic spacer (NMS) materials are compared.	140
7.6	SAF-enhanced ds-MTJ spacer parameters and coupling values [217]. The optimal thickness for each spacer material is highlighted.	143

7.7	Cross-architecture comparison of the seven device configurations studied in this chapter. Each column represents a distinct architecture family, and rows compare key design metrics and the role of interlayer exchange coupling (IEC). Switching speedup is relative to a conventional single-RL perpendicular magnetic tunnel junction (pMTJ). Dashes indicate metrics not applicable or not directly quantified for that architecture. The ultra-scaled and multi-level cell (MLC) studies (Sections 7.3 and 7.4) use $J_{\text{iec}} = 0$, the IEC role listed refers to the effect observed when coupling is subsequently introduced (Section 7.5).	146
-----	--	-----

List of Abbreviations

AFM	antiferromagnet
AMG	algebraic multigrid
AP	anti-parallel
ARK	additive Runge–Kutta
BDF	backward differentiation formula
BEM	boundary element method
BiCGSTAB	Biconjugate Gradient Stabilized
CFL	Courant–Friedrichs–Lewy
CMOS	complementary metal oxide semiconductor
CoFeB	cobalt iron boron
CRMT	Continuous Random Matrix Theory
CVODE	C-language Variable-coefficient Ordinary Differential Equation solver
DOF	degrees of freedom
DOS	density of states
DRAM	dynamic random access memory
ds-MTJ	double-spin torque magnetic tunnel junction
DW	domain-wall
FDM	finite difference method
FE	finite element
FE-BEM	finite element-boundary element method
FEM	finite element method
FeRAM	ferroelectric RAM
FL	free layer
FM	ferromagnetic
GMR	giant magnetoresistance
GMRES	Generalized Minimal Residual
HDD	hard-disk drive
HL	hard layer
IEC	interlayer exchange coupling
ILU	incomplete LU
IMA	in-plane magnetic anisotropy

IMEX	implicit-explicit
iMTJ	in-plane magnetic tunnel junction
LL	Landau–Lifshitz
LLG	Landau–Lifshitz–Gilbert
MgO	magnesium oxide
MLC	multi-level cell
MRAM	magnetoresistive random access memory
MTJ	magnetic tunnel junction
NEGF	Nonequilibrium Green Functions
NM	non-magnetic
NMS	nonmagnetic spacer
NVM	non-volatile memory
ODE	ordinary differential equation
P	parallel
PCG	preconditioned conjugate gradient
PCRAM	phase-change RAM
PDE	partial differential equation
PL	polarization layer
PMA	perpendicular magnetic anisotropy
pMTJ	perpendicular magnetic tunnel junction
RAM	random access memory
RHS	right-hand side
RILU	relaxed incomplete LU
RK4	fourth-order Runge–Kutta
RKKY	Rudderman–Kittel–Kasuya–Yosida
RL	reference layer
RRAM	resistive RAM
SAF	synthetic antiferromagnet
SOR	successive over-relaxation
SOT-MRAM	spin-orbit torque magnetoresistive random access memory
SPD	symmetric positive-definite
SRAM	static random access memory
STT	spin-transfer torque
STT-MRAM	spin-transfer torque magnetoresistive random access memory
SW-MRAM	Stoner–Wohlfarth magnetoresistive random access memory
TB	tunnel barrier
TCAD	technology computer aided design
TMR	tunneling magnetoresistance
TPS	tangent-plane scheme
VCMA	voltage-control of magnetic anisotropy

ZL Zhang and Li

List of Symbols

A	Exchange coefficient (J m^{-1})
D_S	Diffusion coefficient in the tunnel barrier ($\text{m}^2 \text{s}^{-1}$)
D_e	Electron diffusion coefficient ($\text{m}^2 \text{s}^{-1}$)
D_p	Diffusion polarization
E_b	Energy barrier (J)
G	Electrical conductance (S)
H_K	Anisotropy field (A m^{-1})
I	Total charge current (A)
J_1	Bilinear interlayer exchange coupling constant (J m^{-2})
J_2	Biquadratic interlayer exchange coupling constant (J m^{-2})
J_{iec}	Interlayer exchange coupling energy density (J m^{-2})
J_{c0}	Critical switching current density (A m^{-2})
K_{int}	Interface anisotropy coefficient (J m^{-2})
K_i	Interfacial anisotropy energy density (J m^{-2})
K_u	Uniaxial anisotropy energy density (J m^{-3})
K	Anisotropy coefficient (J m^{-3})
M_S	Saturation magnetization (A m^{-1})
P	Spin polarization at tunnel barrier
R_{AP}	Resistance in anti-parallel state (Ω)
R_{P}	Resistance in parallel state (Ω)
R	Electrical resistance (Ω)
V	Electrical potential (V)
Δ	Thermal stability factor
α	Gilbert damping constant
β_D	Diffusion spin asymmetry parameter
β_σ	Conductivity spin asymmetry parameter
η_{IP}	In-plane torque reduction factor
η_{OP}	Out-of-plane polarization
η	Spin-transfer torque efficiency
γ_0	Rescaled gyromagnetic ratio ($\text{m A}^{-1} \text{s}$)
γ	Gyromagnetic ratio ($\text{rad T}^{-1} \text{s}$)
$\hbar$	Reduced Planck constant (J s)
λ_J	Exchange length (m)

λ_{sf}	Spin-flip length (m)
λ_{ϕ}	Dephasing length (m)
λ_m	Momentum relaxation length (m)
λ	Slonczewski polarization factor
$\mathbf{E}$	Electric field vector (V m^{-1})
$\mathbf{H}_{\text{eff}}$	Effective magnetic field vector (A m^{-1})
$\mathbf{J}_C$	Charge current vector (A m^{-2})
$\mathbf{S}_s$	Spin current polarization density tensor (A s^{-1})
$\mathbf{S}$	Spin accumulation vector (A m^{-1})
$\mathbf{T}_S$	Spin torque vector ($\text{A m}^{-1} \text{s}^{-1}$)
$\mathbf{m}$	Unit magnetization vector
μ_0	Vacuum permeability (N A^{-2})
μ_B	Bohr magneton (J T^{-1})
ρ_m	Spin mixing resistivity (Ωm)
σ_m	Spin mixing conductivity (S m^{-1})
σ_p	Conductivity polarization
σ	Electrical conductivity (S m^{-1})
e	Elementary charge (C)
g	Electron g-factor
h	Planck constant (Js)

Introduction

The relentless demand for faster, smaller, and more energy-efficient digital storage and processing technologies continues to drive advances in spintronics. This field uses electron spin rather than its charge for information storage and manipulation. This need has been further intensified by the rise of big data and artificial intelligence, which require rapid data writing, high-density storage, and reduced energy consumption [1]. In particular, the ability to store information reliably and access it quickly while minimizing power usage has become critical for both embedded and large-scale computing systems.

Since the invention of the transistor [2, 3], several breakthroughs have shaped the evolution of computing and memory systems. Among these, the introduction of complementary metal oxide semiconductor (CMOS) technology more than 50 years ago [4] enabled the transition from magnetic-core memories [5] to charge-based transistor memories like static and dynamic random access memory (RAM). Flash memory further extended the capabilities of non-volatile storage. These developments led to the establishment of a hierarchical memory system, where each level is optimized for speed, power, cost, and location (see Figure 1.1).

However, the scaling of silicon CMOS is approaching its physical and economic limits [6, 7]. As energy efficiency becomes a key constraint, particularly for mobile, wearable, and edge computing devices, alternative technologies are under active exploration. A crucial characteristic for energy saving is non-volatility: the ability to retain data even when power is off. This is especially important in devices that remain idle for extended periods, as well as in high-performance systems such as supercomputers and data centers, where improving energy efficiency without compromising speed or endurance is essential.

Several types of non-volatile memory (NVM) technologies have been proposed, including variants of RAM such as resistive RAM (RRAM) [8], ferroelectric RAM (FeRAM) [9], phase-change RAM (PCRAM) [10], and magnetoresistive random access memory (MRAM). Each offers a unique balance of speed, endurance, scalability, and energy efficiency, but also faces distinct challenges. Among these, MRAM has emerged as a particularly promising candidate [11, 12, 13] due to its non-volatility, fast operation, and excellent endurance.

Unlike conventional charge-based memories, MRAM exploits the quantum-mechanical degree of freedom associated with electron spin. The core storage element is the MTJ, which consists of a thin insulating barrier, typically magnesium oxide (MgO),




Figure 1.1: Schematics of different existing memory types (gray pyramid) and prospective spintronic technologies based on magnetic tunnel junctions (MTJs) that could replace them (colored pyramid). L1-L4 denote various cache levels; static random access memory (SRAM) and dynamic random access memory (DRAM) represent static and dynamic random-access memory, respectively; and hard-disk drive (HDD) represents magnetic hard disk drives. Mechanisms for writing bit information in MTJ devices include the Oersted field, spin-orbit torque magnetoresistive random access memory (SOT-MRAM), spin-transfer torque magnetoresistive random access memory (STT-MRAM), and domain-wall (DW) motion.

sandwiched between two ferromagnetic layers, such as cobalt iron boron (CoFeB). One layer has a fixed magnetization and serves as a reference, while the other (the free layer) can switch its magnetization direction. Binary information is encoded in the relative alignment of these magnetizations: the parallel (P) configuration results in low resistance, and the anti-parallel (AP) configuration yields high resistance. This difference, known as the tunneling magnetoresistance (TMR) effect [14, 15, 16, 17], enables simple, reliable electrical readout of the stored bit.

Early generations of MRAM relied on magnetic fields for data writing, in Stoner-Wohlfarth magnetoresistive random access memory (SW-MRAM), for instance, currents driven through orthogonal lines produced Oersted fields that switched the magnetization of the free layer. While functional, this method faced limitations in scalability and write selectivity. Toggle MRAM enhanced stability and improved write

margins, yet it remained dependent on external fields, making large-scale integration into standard semiconductor processes challenging.

A major leap was the development of STT-MRAM [18, 19, 20]. In STT-MRAM, a spin-polarized current is passed directly through the MTJ, allowing the magnetization of the free layer to be switched without external magnetic fields. This simplified the cell architecture and improved scalability. STT-MRAM also features fast switching and endurance on par with SRAM or DRAM [21]. However, its shared read/write current paths lead to tunnel barrier wear and write disturbances, particularly at high operating frequencies. Despite this, refined device designs have demonstrated pulse switching as short as 0.3 ns [22], highlighting its potential for ultra-fast operation.

MRAM is therefore positioned to function both as an embedded memory technology and as a candidate for universal memory. It offers a unique balance of speed, non-volatility, low power, and endurance, and can complement or replace other memory types across various levels of the memory hierarchy. As illustrated in Figure 1.1, MTJ-based MRAM can be deployed from last-level cache to system memory, bridging the gap between volatile and non-volatile storage in conventional architectures, and even enabling in-memory computing [23, 24].

The success of MRAM has been driven in part by material innovations. The CoFeB/MgO/CoFeB system, owing to its high TMR [25, 26] and interfacial perpendicular magnetic anisotropy (PMA) [27, 28], has become a standard for MTJ design. It enables efficient switching, strong thermal stability, and excellent scalability, which are critical for both embedded and high-density applications.

Complementing experimental developments, the rapid progress in MRAM has been supported by advanced computational modeling. Simulation tools are essential for designing and optimizing devices with high efficiency and reliability. The Landau–Lifshitz–Gilbert (LLG) equation, extended with spin-transfer torque terms, provides the theoretical foundation for modeling magnetization dynamics in MTJs. Standard numerical methods include the finite difference method (FDM), known for parallelization efficiency, and the finite element method (FEM), which excels at handling complex geometries and material boundaries [29, 30].

As device dimensions shrink to the sub-20 nm regime, the assumptions underlying simple macro-spin models break down [31]. Nonuniform current distributions arise from the angle-dependent tunneling conductance of the MgO barrier, leading to spatially varying spin torques that spatially averaged treatments cannot capture. Composite free layers introduce additional interfaces whose spin-filtering and spin-dephasing properties must be modeled explicitly [32]. Moreover, the increasing adoption of synthetic antiferromagnet (SAF) structures and interlayer exchange coupling (IEC)-mediated multilayers [33, 34] demands numerical frameworks capable of resolving the coupling across ultrathin nonmagnetic spacers without resorting to prohibitively fine volumetric meshes. Addressing these challenges requires a unified computational approach that couples three-dimensional micromagnetic solvers with self-consistent spin and charge transport models [35].

This thesis develops such a framework and applies it to investigate the switching dynamics, reliability phenomena, and design optimization of next-generation STT-MRAM devices. By combining a FEM-based micromagnetic solver with a coupled

drift-diffusion transport formalism, the work provides physical insights into ultra-scaled device behavior that are inaccessible to conventional macro-spin or one-dimensional analytical models.

1.1 Outline of the Thesis

This thesis focuses on the development and application of advanced computational techniques to analyze and optimize the behavior of modern STT-MRAM devices. Specifically, it addresses the modeling of various torque contributions and their impact on magnetization switching dynamics, leveraging simulation frameworks built on C++ libraries. The thesis is organized as follows.

Chapter 2 provides a concise overview of the evolution, current architectures, and key operating principles of MRAM technologies, with particular emphasis on STT-MRAM cells. The magnetoresistance phenomena underlying device readout, including giant magnetoresistance (GMR) and TMR, are reviewed, and the physical mechanisms governing thermal stability and spin-transfer torque switching are discussed.

Chapter 3 introduces the micromagnetic modeling framework employed to simulate magnetization dynamics in spintronic devices. Central to this approach is the LLG equation, whose derivation and physical interpretation are outlined in detail. The chapter discusses the primary contributions to the effective magnetic field, including exchange, anisotropy, demagnetization, IEC, and external fields. It presents simplified expressions for spin-transfer torques (STTs) as they arise in both spin valve structures and MTJs.

Chapter 4 presents two numerical methods for spatially discretizing the computational domain in magnetization dynamics simulations: the FDM and the FEM. The chapter begins by introducing the principles underlying both approaches and illustrates how the effective magnetic field components are discretized. Special emphasis is placed on the computation of current density redistribution in MTJs with nonuniform magnetization, which significantly influences the spin-transfer torque. Furthermore, the in-house FDM-based solver is discussed in detail, including its implementation of the LLG equation.

Chapter 5 addresses the numerical time integration of the spatially discretized LLG equation. The key challenges of stiffness, the unit-sphere constraint, and geometric structure preservation are identified. Three classes of time-stepping schemes are developed and compared: adaptive higher-order backward differentiation formula (BDF) methods, implicit-explicit (IMEX) schemes that treat the exchange interaction implicitly, and a tangent-plane formulation that recasts the nonlinear constrained problem as a sequence of linear saddle-point systems. Verification is carried out using the μ MAG Standard Problem 4.

Chapter 6 extends the micromagnetic framework to include a coupled spin and charge drift-diffusion formalism for computing spin-transfer torques in ultra-scaled MRAM cells. Specialized boundary conditions are derived to describe spin filtering at MgO tunnel barriers and spin dephasing at metallic interfaces. A novel numerical treatment of IEC is introduced via an interface-mapping algorithm, enabling efficient simulation of SAF structures and composite free layers without volumetric meshing of angstrom-scale spacers.

Chapter 7 applies the complete simulation framework to investigate three interconnected phenomena in next-generation STT-MRAM devices. First, the switching dynamics in ultra-scaled cells with composite free layers are analyzed, revealing the role of sequential switching and interface effects. Second, the back-hopping phenomenon is shown to be a deterministic consequence of composite torque dynamics, and its deliberate exploitation for multi-level cell (MLC) operation is demonstrated. Third, advanced device architectures incorporating double spin-torque MTJs and SAF-enhanced reference layers are systematically investigated, and quantitative design thresholds for reliable operation are identified.

Chapter 8 summarizes the principal contributions of this work, presents the main conclusions, and discusses directions for future research, including the advancement of the numerical methods themselves toward higher-order finite elements, parallel-in-time integration, and fully coupled space-time formulations.

1.2 Research Setting

The research described in this work was carried out as part of the Christian Doppler Laboratory for Nonvolatile Magnetoresistive Memory and Logic (NOVOMEMLOG). The Christian Doppler Association fosters collaboration between research institutions and companies engaged in application-oriented basic research. In this laboratory, a partnership was established between the Institute for Microelectronics at TU Wien and Silvaco Inc.. This company develops and provides electronic device automation and software tools for technology computer aided design (TCAD).

Spin-Transfer Torque Magnetic Random Access Memory 2

2.1 Introduction to Magnetoresistive Random Access Memory

Magnetoresistive random access memory (MRAM) is a non-volatile memory technology that stores binary information through the orientation of magnetic moments rather than electrical charge [12]. Each memory cell typically consists of a magnetoresistive element, either an magnetic tunnel junction (MTJ) or a spin valve, composed of two ferromagnetic (FM) layers: a fixed reference layer (RL) and a switchable free layer (FL), separated by a thin spacer layer. The electrical resistance of the cell depends on the relative alignment of these magnetic layers: parallel (P) alignment yields low resistance, while anti-parallel (AP) alignment results in high resistance [12, 36]. These states correspond to logic '0' and '1'. Data retention relies on an energy barrier provided by magnetic anisotropy, whereas readout operations exploit magnetoresistive effects. Write operations are predominantly performed using spin-transfer torque (STT), while device designs incorporating synthetic antiferromagnet (SAF) structures are also gaining interest. SAFs suppresses stray magnetic fields and dipolar coupling between layers, thereby improving scalability and enabling denser cell integration. Moreover, their intrinsic stability and reduced sensitivity to external perturbations enhance device reliability, making them particularly suitable for large-scale complementary metal oxide semiconductor (CMOS) integration [33, 37].”

With advantages such as non-volatility, high endurance, rapid access times, and compatibility with existing CMOS technologies, spin-transfer torque magnetoresistive random access memory (STT-MRAM) is increasingly considered a promising candidate to replace traditional dynamic random access memory (DRAM) and static random access memory (SRAM) in future memory systems [36, 37, 38]. This chapter outlines the fundamental principles of MRAM, describes its historical progression, and examines physical mechanisms, including magnetoresistance, thermal stability, and STT switching [39, 40].

2.1.1 Historical Context

The development of random access memory began in the mid-20th century, initially storing bits as electrically charged spots on cathode-ray tubes [41]. Magnetic memory technologies significantly advanced in the 1950s with magnetic core memory, where information was stored in the magnetization orientation of ferrite rings [42]. This robust form of non-volatile memory was extensively employed until it was superseded by semiconductor-based technologies such as DRAM and SRAM, offering superior scalability, integration density, and ease of manufacturing [43].

Interest in magnetic memory resurged following the discovery of giant magnetoresistance (GMR) in 1988, independently reported by Fert *et al.* [44] and Grünberg *et al.* [45]. GMR introduced substantial resistance variations based on the relative magnetization alignment of multilayer structures, fundamentally altering magnetic storage readout techniques and catalyzing the advent of spintronics [45, 46].

Further advancements were achieved through the discovery of tunneling magnetoresistance (TMR) in MTJs, which featured electron tunneling through insulating barriers. Significant TMR effects were first observed in junctions involving aluminum oxide barriers [47, 48]. Later, crystalline magnesium oxide (MgO)-based MTJs significantly enhanced TMR ratios, enabling practical and scalable MRAM technology [25, 26, 49, 50].

Early MRAM devices required external magnetic fields for switching, which limited scalability and device density. This limitation was overcome by introducing STT, enabling purely electrical magnetization reversal within the MTJ. Since switching relies only on charge currents delivered through standard access transistors and interconnects, STT-MRAM is inherently compatible with conventional CMOS fabrication processes, while also offering improved scalability and potential for faster switching [39, 51, 52].

Contemporary MRAM research continues to evolve, exploring innovations such as perpendicular magnetic anisotropy (PMA), SAF, and interlayer exchange coupling (IEC), promising further improvements in storage density, power efficiency, and integration capability [53, 54, 55, 56, 57].

2.2 Magnetoresistance Phenomena

2.2.1 Giant Magnetoresistance

GMR is a quantum mechanical magnetoresistive effect observed in metallic multilayer structures composed of alternating FM layers separated by nonmagnetic conducting spacers. The electrical resistance of such stacks depends on the relative alignment of the magnetizations in the ferromagnetic layers. When the magnetizations are aligned P, majority-spin electrons traverse the structure with minimal scattering, resulting in lower resistance. In contrast, spin-dependent scattering increases in the AP configuration, leading to higher resistance. This type of structure is commonly referred to as a *spin valve* [58].

The magnitude of the effect is characterized by the GMR ratio, which can be expressed in terms of resistance, resistivity, or conductivity:

$$\text{GMR} = \frac{R_{\text{AP}} - R_{\text{P}}}{R_{\text{P}}} = \frac{\rho_{\text{AP}} - \rho_{\text{P}}}{\rho_{\text{P}}} = \frac{\sigma_{\text{P}}}{\sigma_{\text{AP}}} - 1, \quad (2.1)$$

where R_{P} and R_{AP} are the resistances in the P and AP states, ρ the respective resistivities, and σ the corresponding conductivities.

In the simplest model, electron transport occurs via two parallel spin channels, *spin-up* and *spin-down*, each experiencing different scattering probabilities [59]. When the magnetizations of adjacent FM layers are aligned, the majority-spin channel experiences reduced scattering and can travel relatively unhindered, yielding low resistance. In the AP configuration, both spin channels are subject to frequent scattering events in at least one of the layers, resulting in higher resistance [15].

This spin-dependent scattering mechanism underpins the GMR effect, visually summarized in Figure 2.1. The phenomenon was independently discovered in 1988 by Fert *et al.* [44] and Grünberg [45], who demonstrated significant changes in resistance in Fe/Cr and Fe/Cu multilayers, respectively. Their pioneering work earned Albert Fert and Peter Grünberg the Nobel Prize in Physics in 2007.

The discovery of GMR sparked a revolution in spintronics and magnetic data storage. It was rapidly adopted in hard disk drive read heads in the 1990s, dramatically increasing areal storage density. The underlying principles later inspired magnetoresistive memory concepts such as GMR-based MRAM [60]. Nevertheless, due to the limited resistance contrast achievable with metallic multilayers, GMR spin valves were eventually superseded by MTJ-based architectures that offer higher signal-to-noise ratios and better scalability.




Figure 2.1: Illustration of spin-dependent electron scattering in a metallic multilayer structure demonstrating the GMR effect. Red arrows represent *spin-up* electrons, blue arrows denote *spin-down* electrons, and black arrows indicate the magnetization direction. In the AP configuration (left), both spin channels experience greater scattering, leading to higher resistance. In the P configuration (right), majority-spin electrons encounter less scattering, resulting in lower resistance.

2.2.2 Tunnel Magnetoresistance

TMR is a magnetoresistive phenomenon observed in MTJs, where a thin insulating barrier, most commonly crystalline MgO, separates two FM electrodes. The junction resistance depends on the relative alignment of the magnetization in the ferromagnetic layers, analogous to the GMR effect. In the P configuration, spin-polarized electrons can tunnel more efficiently due to the matching density of states (DOS) at the Fermi level in both layers, resulting in higher conductance and lower resistance. In contrast, when the magnetization is AP, the spin states become misaligned, reducing the tunneling probability and thus increasing the resistance.

TMR can be understood as the counterpart of the GMR effect, observed when the metallic spacer layer in a spin valve is replaced with a thin insulating tunnel barrier, forming an MTJ. The phenomenon was initially discovered by Jullière in 1975 using Fe/GeO_x/Co junctions, where a conductance change of approximately 14 % was observed at 4.2 K [14]. Due to its limited magnitude, the effect initially attracted little attention. However, after the discovery of GMR, TMR was revisited, and significantly stronger effects (10–20 % at room-temperature) were soon reported using amorphous Al₂O₃ barriers between Co or CoFe ferromagnetic layers [48, 61]. For comparison, modern MTJs based on ferromagnets typically exhibit TMR ratios between 50–200 % at room-temperature, with record values exceeding 600 % [53].

The key functional element of a MRAM cell based on the TMR effect is the MTJ. In such junctions, spin-dependent tunneling occurs because the probability of an electron tunneling through the barrier depends on the availability of the same spin states in the destination layer. Since electron spin is conserved mainly during tunneling, electrons can only tunnel into electronic states of the same spin orientation (see Figure 2.2). A transition from the P configuration (Figure 2.2 (left panel)) to the AP configuration (Figure 2.2 (right panel)) effectively inverts the spin bands in one electrode, leading to a mismatch in the spin-resolved DOS between the two layers. As a result, the overall tunneling current is reduced, leading to a measurable change in resistance.

The theoretical foundation of TMR was first described by Jullière [14], who proposed a simple model based on two key assumptions:

- (i) The electron spin is conserved in the tunneling process. This implies that the tunneling of *spin-up* and *spin-down* electrons is an independent process, and that the carriers tunnel from the first film are accepted by empty states of the same spin orientation in the second film.
- (ii) The tunneling conductance G_P and G_{AP} are proportional to the product of the spin-resolved density of states of the two ferromagnetic layers.

Based on these assumptions, the conductance in the P and AP configurations are expressed as:

$$G_P \propto N_1^\uparrow N_2^\uparrow + N_1^\downarrow N_2^\downarrow = 1 - N_1^\uparrow - N_2^\uparrow + 2N_1^\uparrow N_1^\uparrow, \quad (2.2)$$

$$G_{AP} \propto N_1^\uparrow N_2^\downarrow + N_1^\downarrow N_2^\uparrow = N_1^\uparrow + N_2^\uparrow - 2N_1^\uparrow N_1^\uparrow, \quad (2.3)$$




Figure 2.2: Illustration of spin-dependent tunneling responsible for the TMR effect. In the P configuration (left), spin-polarized electrons tunnel into available states of the same spin orientation in the second ferromagnetic layer, resulting in high conductance. In the AP configuration (right), spin mismatch reduces the tunneling probability, leading to low conductance.

where $N^\uparrow$ and $N^\downarrow = 1 - N^\uparrow$ are the *spin-up* and *spin-down* fractions of the density of states at the Fermi level, and subscripts 1 and 2 refer to the first and second ferromagnet, respectively.

Introducing the spin polarization of the conduction electrons as:

$$P = \frac{N^\uparrow - N^\downarrow}{N^\uparrow + N^\downarrow} = 2N^\uparrow - 1, \quad (2.4)$$

we can write the TMR ratio as:

$$\text{TMR} = \frac{G_P - G_{AP}}{G_{AP}} = \frac{R_{AP} - R_P}{R_P} = \frac{2P_1P_2}{1 - P_1P_2}. \quad (2.5)$$

For symmetric junctions where $P_1 = P_2 = P$, this simplifies to:

$$\text{TMR} = \frac{2P^2}{1 - P^2}. \quad (2.6)$$

While Jullière's model provides intuitive insight into the spin-filtered nature of tunneling, it assumes incoherent tunneling and does not account for band symmetries. In practice, the tunneling process significantly depends on the symmetry of the electron wavefunctions. In MTJs with amorphous Al_2O_3 barriers, tunneling is incoherent, and Jullière's model is applicable. However, for crystalline barriers such as MgO , only electrons with specific Bloch state symmetries (e.g., Δ_1) contribute to tunneling. This symmetry filtering effect strongly suppresses conductance in the AP state and leads to much higher TMR ratios [62].

This behavior is reflected in the angular dependence of the conductance, modeled by Slonczewski [63]:

$$G(\theta) = \frac{G_P + G_{AP}}{2} + \frac{G_P - G_{AP}}{2} \cos \theta = \frac{G_P + G_{AP}}{2} \left(1 + \frac{\text{TMR}}{2 + \text{TMR}} \cos \theta \right), \quad (2.7)$$

where θ is the angle between the magnetizations in the two ferromagnetic layers. Maximum conductance is achieved at $\theta = 0$ (P), and minimum at $\theta = \pi$ (AP) [36,64,65].

Although discovered before GMR, practical applications of TMR became viable only in the mid-1990s. In 1995, Moodera *et al.* [48] and Miyazaki *et al.* [66] independently demonstrated room-temperature TMR in junctions with amorphous Al_2O_3 barriers, achieving TMR ratios of 20-30 % [61]. This milestone catalyzed extensive research into MTJ-based memory devices. Further improvements in materials and deposition processes enabled TMR values of up to 70.4 % with Al_2O_3 barriers [67].

A major theoretical breakthrough came in 2001, when Butler *et al.* [49] and Mathon *et al.* [50] predicted that coherent tunneling through crystalline MgO barriers could yield TMR ratios exceeding 1000 %. These predictions were experimentally validated in 2004 by Parkin *et al.* [25] and Yuasa *et al.* [26], who reported room-temperature TMR values of 180–220 % in epitaxial Fe/MgO/Fe junctions. Subsequent work highlighted the critical role of Ar pressure during MgO sputtering, which affects barrier crystallinity and overall performance [68]. Continued optimization of interface engineering and thermal annealing led to record-breaking TMR values of 410 % [69] and 604 % at room-temperature (and up to 1144 % at 4.2 K) in CoFeB/MgO stacks [53].

Today, MgO-based MTJ structures are central to commercial STT-MRAM technologies due to their high TMR, scalability, and CMOS compatibility.

2.3 Spin-Transfer Torque Switching

MRAM cells are built around an MTJ, which consists of a pinned RL and an FL, separated by a thin insulating tunnel barrier (TB), typically MgO. The magnetization orientation of both the FL and RL can either lie in-plane or perpendicular to the film plane due to PMA, as illustrated in Figure 2.3. When both magnetic layers exhibit perpendicular anisotropy, the device is referred to as a perpendicular magnetic tunnel junction (pMTJ), while in-plane magnetized structures are called in-plane magnetic tunnel junction (iMTJ).

The first MTJs used amorphous Al_2O_3 barriers and achieved modest TMR values in the double digits. However, theoretical predictions [18, 19] and experimental demonstrations [39, 70] of crystalline MgO-based MTJs showed a dramatic increase in TMR, reaching up to 600 % at room-temperature under optimized laboratory conditions. Today, typical TMR values in perpendicularly magnetized MTJs range between 100 % and 200 % [71].

Data retrieval from MTJs is performed by applying a small bias voltage (typically 0.1-0.2 V) and measuring the resulting resistance across the junction. Operating at such low bias avoids disturbing the stored magnetization state and ensures that the




Figure 2.3: Illustration of magnetization orientations in MTJs, showing the difference between in-plane (iMTJ) and perpendicular (pMTJ) configurations.

TMR remains high. The TMR is highest at low bias and decreases with increasing voltage [49].

The thermal stability of MRAM cells, critical for ensuring data retention over years, is characterized by the dimensionless factor:

$$\Delta = \frac{\mathcal{E}_B}{k_B T} = \frac{\mu_0 M_s H_k V}{2 k_B T}, \quad (2.8)$$

where $\mathcal{E}_B$ is the energy barrier separating the P and AP states for out-of-plane magnetized films, V is the volume of the free layer, k_B is the Boltzmann constant, T the absolute temperature, and $H_k = 2K_{\text{eff}}/(\mu_0 M_s)$ is the effective anisotropy field. The effective anisotropy $K_{\text{eff}} = K_u - \mu_0 M_s^2/2$ accounts for the reduction of the intrinsic uniaxial anisotropy K_u by the demagnetizing field $H_d = M_s$ of the thin film. To ensure reliable retention over ten years, Δ values of 40–60 are typically required [72, 73].

The energy barrier is estimated by:

$$\mathcal{E}_B \approx \frac{\mu_0 M_s H_K V}{2}, \quad (2.9)$$

where μ_0 is the vacuum permeability, M_s the saturation magnetization, H_K the anisotropy field, and V the volume of the FL. Optimization of H_K and the FL geometry are thus key to MRAM scaling [28].

STT was independently proposed by Slonczewski and Berger in 1996 [18, 19], enabling purely electrical switching in MTJs. The first experimental demonstration followed in metallic spin valves [39] and later in MTJs [70].

Electrons flowing from the RL to the FL generate a spin-polarized current. Due to exchange interaction, spins align with the magnetization of the FL, and conservation of angular momentum causes the transverse spin component to exert a torque on the FL, known as STT [18, 19]. Reversing current direction induces AP switching via reflected minority spins [35]. Since STT depends on transverse spin components, no torque is exerted when the magnetizations are collinear, leading to an incubation delay. Thermal fluctuations are then needed to initiate switching, and if the current pulse duration is shorter than the delay, write errors may occur [74, 75].

The critical current at $T = 0$ is:

$$I_{c0} = \frac{2e\alpha}{\hbar\eta} \cdot \frac{\mu_0 M_s H_K V}{2}, \quad (2.10)$$

and the corresponding current density:

$$J_{c0} = \frac{I_{c0}}{A} = \frac{2e\alpha}{\hbar\eta} \cdot \frac{\mu_0 M_S H_K t_{FL}}{2}, \quad (2.11)$$

where A is the cross-sectional area, t_{FL} is the FL thickness, α is the Gilbert damping factor, and η is the STT efficiency parameter [72, 76].

In addition to the switching current, data retention in STT-MRAM is governed by the thermal stability factor:

$$\Delta = \frac{\mathcal{E}_B}{k_B T} = \frac{\mu_0 M_S H_K V}{2k_B T}, \quad (2.12)$$

where $\mathcal{E}_B$ is the energy barrier separating stable magnetization states P or AP, k_B is the Boltzmann constant, and T is the absolute temperature. A larger Δ ensures long-term retention by suppressing the thermally activated switching. However, increasing Δ also raises the required switching current, creating a trade-off between writability and data retention.

Early MRAM relied on in-plane magnetic anisotropy (IMA), with shape anisotropy in elliptical FLs [11, 77]. In these devices, switching involves out-of-plane precession, as the reversal trajectory does not follow thermal fluctuations directly [78]. In-plane devices were typically larger than 60 nm to maintain thermal stability, which limited scaling [79]. Furthermore, both the anisotropy field H_K and the demagnetizing field H_d must be overcome for switching, increasing energy requirements [80]:

$$J_{c0}^{IMA} = \frac{2e\alpha}{\hbar\eta} \cdot \mu_0 M_S \left(|H_{ext}| + |H_K| + \frac{|H_d|}{2} \right) t_{FL}. \quad (2.13)$$

Additionally, stronger in-plane demagnetizing fields oppose magnetization changes, necessitating higher switching currents [72]. These drawbacks hindered scalability and energy efficiency.

PMA structures, dominated by interfacial anisotropy (e.g., CoFeB/MgO), mitigate these challenges by eliminating H_d and orienting magnetization perpendicular to the film plane, allowing for more energy-efficient STT switching [81]. PMA devices also support cylindrical layouts with better scaling and enhanced thermal stability. The corresponding critical current density simplifies to:

$$J_{c0}^{PMA} = \frac{e\alpha}{\hbar\eta} \cdot \mu_0 M_S H_K t_{FL}. \quad (2.14)$$

Ikeda *et al.* [28, 53] demonstrated robust PMA in MgO/CoFeB stacks, a pivotal advancement in the development of scalable and high-performance MRAM [82]. Annealing also significantly impacts both TMR and anisotropy optimization [83].

Further improvements arise from engineering multilayered FLs, SAF RLs [84, 85, 86], and tailored interfaces. SAF designs use two antiferromagnetically coupled ferromagnetic layers (often separated by Ru) to cancel dipolar stray fields on the FL, enhancing write symmetry and error tolerance [34]. Advanced strategies include ionic liquid gating to modulate Ruderman–Kittel–Kasuya–Yosida (RKKY) coupling [87] and re-doped Ru spacers for thermal robustness during CMOS backend processing [88].

Micromagnetic Modelling

The theory of micromagnetism, first formalized by William Fuller Brown Jr. in the mid-20th century [89], provides a semi-classical framework for modeling magnetization dynamics in ferromagnetic materials. It bridges classical field theory and quantum mechanics by incorporating quantum-derived effects, such as the exchange interaction, into a continuum description [90]. This approach is particularly suitable for modern spintronic devices, such as STT-MRAM, where atomistic *ab initio* simulations are computationally infeasible due to the device scale, and macroscopic models fail to resolve magnetization behavior on the required length and time scales.

The central assumption of micromagnetism is that, on length scales larger than the exchange length $l_{\text{ex}} = \sqrt{2A/(\mu_0 M_s^2)}$, where A is the exchange stiffness constant, the discrete distribution of atomic magnetic moments $\mathbf{S}_i$ can be approximated by a continuous magnetization vector field $\mathbf{M}(\mathbf{x})$. This approximation is justified when averaging over spatial regions Ω whose volume is much greater than the characteristic exchange volume λ^3 , since in such areas the exchange interaction enforces near-uniform alignment of the magnetic moments:

$$\sum_i \mathbf{S}_i \approx \int \Omega \mathbf{M}(\mathbf{x}) d\mathbf{x}. \quad (3.1)$$

As a result, micromagnetics enables the simulation of complex magnetization patterns in real devices by using continuous field variables rather than tracking individual spins. Another fundamental assumption of the micromagnetic model is that magnetization is a constant norm throughout space and time. The magnetization vector is thus written as:

$$\mathbf{M}(\mathbf{x}) = M_s \mathbf{m}(\mathbf{x}) \quad \text{with} \quad |\mathbf{m}(\mathbf{x})| = 1, \quad (3.2)$$

where $\mathbf{m}(\mathbf{x})$ is a unit vector field representing the direction of the local magnetization, and M_s is the saturation magnetization.

Micromagnetic models are effective for simulating time-dependent magnetization dynamics, which are crucial in applications such as magnetic sensors, memory elements, and spin-torque oscillators. In these systems, the magnetization is not static but evolves in response to external fields, spin-polarized currents, or thermal fluctuations.

3.1 Landau-Lifshitz-Gilbert Equation

The dynamics of the magnetization vector $\mathbf{M}(\mathbf{x}, t)$ are governed by the Landau–Lifshitz–Gilbert (LLG) equation, a semi-classical differential equation incorporating both precessional motion and damping. The origins of this equation can be traced to fundamental work on magnetic moments and torque. A magnetic moment $\boldsymbol{\mu}$ experiences a torque $\boldsymbol{\tau}$ in an external magnetic induction field $\mathbf{B}$, described by:

$$\boldsymbol{\tau} = \boldsymbol{\mu} \times \mathbf{B}. \quad (3.3)$$

This torque induces rotational motion according to Newton’s second law for rotation:

$$\frac{d\mathbf{L}}{dt} = \boldsymbol{\tau}, \quad (3.4)$$

where $\mathbf{L}$ denotes angular momentum. For an electron, the magnetic moment is related to its angular momentum by:

$$\boldsymbol{\mu} = \gamma \mathbf{L}, \quad (3.5)$$

γ being the electron’s gyromagnetic ratio. Combining these expressions yields the classical equation of motion for an individual magnetic moment:

$$\frac{d\boldsymbol{\mu}}{dt} = \gamma \boldsymbol{\mu} \times \mathbf{B}. \quad (3.6)$$

Equation (3.6) describes the precessional motion of a discrete magnetic moment. In micromagnetism, however, the magnetization is treated as a continuous field rather than discrete moments. Assuming uniform spin density and constant magnitude, this continuous magnetization is expressed as:

$$\mathbf{M}(\mathbf{x}, t) = M_s \mathbf{m}(\mathbf{x}, t), \quad |\mathbf{m}(\mathbf{x}, t)| = 1, \quad (3.7)$$

where $\mathbf{m}$ is a unit vector indicating magnetization direction.

In this continuous framework, the magnetic induction $\mathbf{B}$ is replaced by an effective magnetic field $\mathbf{H}_{\text{eff}}$, incorporating external fields and internal contributions such as exchange, anisotropy, and demagnetization (see Section 3.2). Introducing the rescaled gyromagnetic ratio $\gamma_0 = -\gamma\mu_0$, the pure precession equation for the normalized magnetization becomes:

$$\frac{d\mathbf{m}}{dt} = -\gamma_0 \mathbf{m} \times \mathbf{H}_{\text{eff}}, \quad (3.8)$$

which corresponds to constant-energy precession around the effective field, as illustrated in Figure 3.1 (a). Historically, a time-dependent model describing undamped precession was first introduced by Bloch in 1932 [91].

Experimental observations, however, indicate that magnetization dynamics inherently involve energy dissipation. Although the microscopic mechanisms behind this dissipation remain unclear, they can be modeled effectively through phenomenological damping terms.




Figure 3.1: Magnetization dynamics on the unit sphere. (a) Pure precession: $\mathbf{m}$ rotates around $\mathbf{H}_{\text{eff}}$ at constant polar angle (no damping, $\alpha = 0$). (b) Pure damping: $\mathbf{m}$ relaxes directly toward $\mathbf{H}_{\text{eff}}$ along a meridian (no precession). (c) Combined motion: damped precession spirals toward equilibrium.

In 1935, Landau and Lifshitz [92] introduced a damping term proportional to the component of $\mathbf{H}_{\text{eff}}$ perpendicular to $\mathbf{m}$, leading to the Landau–Lifshitz (LL) equation:

$$\frac{d\mathbf{m}}{dt} = -\gamma_0 \mathbf{m} \times \mathbf{H}_{\text{eff}} - \gamma_0 \lambda \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}}), \quad (3.9)$$

where $\lambda > 0$ is a phenomenological damping parameter. This equation describes the combined effects of precession and damping (see Figure 3.1 (c)).

In 1955, Gilbert [93] proposed an alternative damping formulation by introducing a viscous-type damping proportional to the time derivative of magnetization:

$$\frac{d\mathbf{m}}{dt} = -\gamma_0 \mathbf{m} \times \mathbf{H}_{\text{eff}} + \alpha \mathbf{m} \times \frac{d\mathbf{m}}{dt}, \quad (3.10)$$

Taking the scalar product of Equation (3.10) with $\mathbf{m}$, one obtains the orthogonality condition:

$$\mathbf{m} \cdot \frac{d\mathbf{m}}{dt} = 0, \quad (3.11)$$

ensuring the magnetization magnitude remains constant over time.

Taking the cross product of Equation (3.10) with $\mathbf{m}$ and applying the vector triple product identity:

$$\mathbf{a} \times (\mathbf{b} \times \mathbf{c}) = (\mathbf{a} \cdot \mathbf{c})\mathbf{b} - (\mathbf{a} \cdot \mathbf{b})\mathbf{c}, \quad \forall \mathbf{a}, \mathbf{b}, \mathbf{c} \in \mathbb{R}^3,$$

together with constraints $|\mathbf{m}| = 1$ and $\mathbf{m} \cdot \frac{d\mathbf{m}}{dt} = 0$, Equation (3.10) can be explicitly rewritten in the numerically convenient LLG form:

$$\frac{d\mathbf{m}}{dt} = \underbrace{\frac{-\gamma_0}{1 + \alpha^2} \mathbf{m} \times \mathbf{H}_{\text{eff}}}_{\text{precession}} - \underbrace{\frac{\gamma_0 \alpha}{1 + \alpha^2} \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}})}_{\text{damping}}, \quad (3.12)$$

separating precessional motion from damping.

The energy evolution during magnetization dynamics is obtained by computing the time derivative of the total energy functional:

$$\frac{d}{dt}\mathcal{E}[\mathbf{m}] = \int_{\Omega} \frac{\delta\mathcal{E}}{\delta\mathbf{m}} \cdot \frac{\partial\mathbf{m}}{\partial t} d\mathbf{x} = -\mu_0 M_s \int_{\Omega} \mathbf{H}_{\text{eff}} \cdot \frac{\partial\mathbf{m}}{\partial t} d\mathbf{x}, \quad (3.13)$$

where the second equality follows from the definition of the effective field in Equation (3.17). Substituting Equation (3.12), the precessional term vanishes because

$$\mathbf{H}_{\text{eff}} \cdot (\mathbf{m} \times \mathbf{H}_{\text{eff}}) = 0, \quad (3.14)$$

as the cross product is orthogonal to both of its arguments. Only the damping term contributes. Using the vector identity $\mathbf{H}_{\text{eff}} \cdot [\mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}})] = -\|\mathbf{m} \times \mathbf{H}_{\text{eff}}\|_2^2$, where $\|\cdot\|_2$ denotes the Euclidean norm, which holds for $|\mathbf{m}| = 1$, the energy dissipation relation becomes [94]:

$$\frac{d}{dt}\mathcal{E}[\mathbf{m}] = -\frac{\alpha\gamma_0}{1+\alpha^2} \int_{\Omega} \|\mathbf{m} \times \mathbf{H}_{\text{eff}}\|_2^2 d\mathbf{x} \leq 0. \quad (3.15)$$

The prefactor depends linearly on the damping parameter α , while the factor $(1+\alpha^2)$ in the denominator arises from the conversion between the Gilbert and Landau-Lifshitz forms. This result demonstrates the dissipative nature of Gilbert damping, driving magnetization toward equilibrium alignment with the effective field.

Figure 3.1 summarizes distinct magnetization dynamics: (a) pure precession, (b) pure damping, and (c) combined damped precession. Simulations based on the LLG equation are now essential for designing and optimizing STT-MRAM devices, enabling the prediction of switching behavior, thermal stability, and transient dynamics under realistic operating conditions.

3.2 Effective Magnetic Field

The total energy of a ferromagnetic system with respect to its magnetization consists of several contributions that depend on the material's intrinsic properties. Some of these, such as the demagnetization (or magnetostatic) energy, can be described within the framework of classical magnetostatics. Others, such as exchange energy and magnetocrystalline anisotropy energy, originate from quantum-mechanical interactions.

To derive expressions for the individual contributions to the effective magnetic field used in the LLG equation, one must first define the magnetic free Gibbs energy $\mathcal{E}$ of the ferromagnetic domain. The energy functional comprises several contributions, most notably the exchange energy, the anisotropy energy, the Zeeman energy from an externally applied magnetic field, and the magnetostatic energy. These components will be introduced in more detail in the following sections.

Once the total free Gibbs energy is defined, the equilibrium state of the magnetization corresponds to a configuration that minimizes this energy, leading to the variational problem:

$$\min_{|\mathbf{m}|=1} \mathcal{E}(\mathbf{m}) \quad (3.16)$$

A stable magnetization configuration thus represents a balance among competing energy terms. The effective magnetic field $\mathbf{H}_{\text{eff}}$, which governs the magnetization dynamics in the LLG equation, is defined as the functional derivative of the energy with respect to the magnetization:

$$\mathbf{H}_{\text{eff}} = -\frac{1}{\mu_0 M_s} \frac{\delta \mathcal{E}}{\delta \mathbf{m}} \quad (3.17)$$

From the Euler-Lagrange equations associated with the minimization in Equation (3.16), one obtains the boundary value problem:

$$\mathbf{H}_{\text{eff}} = 0, \quad (3.18a)$$

$$\frac{\delta \mathcal{E}(\mathbf{m})}{\delta \nabla m_i} \cdot \mathbf{n} = 0, \quad \text{for } i = 1, 2, 3 \quad (3.18b)$$

where $\mathbf{n}$ denotes the unit normal vector at the boundary of the ferromagnetic domain. These conditions imply that, at equilibrium, the magnetization vector is locally aligned with the effective field. Consequently, the LLG equation describes a relaxation process toward a state of minimal magnetic Gibbs energy.

In this thesis, the effective field is composed of the following contributions:

$$\mathbf{H}_{\text{eff}} = \mathbf{H}_{\text{exc}} + \mathbf{H}_{\text{ani}} + \mathbf{H}_{\text{ext}} + \mathbf{H}_{\text{demag}} + \mathbf{H}_{\text{curr}} + \mathbf{H}_{\text{th}} + \mathbf{H}_{\text{iec}} \quad (3.19)$$

Here, $\mathbf{H}_{\text{exc}}$ is the exchange field, $\mathbf{H}_{\text{ani}}$ the anisotropy field, $\mathbf{H}_{\text{ext}}$ the externally applied field, $\mathbf{H}_{\text{demag}}$ the demagnetizing (magnetostatic) field, $\mathbf{H}_{\text{curr}}$ the ampère field, $\mathbf{H}_{\text{th}}$ the thermal field, $\mathbf{H}_{\text{iec}}$ the interlayer exchange field. Each of these contributions will be discussed in the subsequent sections.

3.2.1 Exchange Field

A defining property of ferromagnetic materials is their ability to retain a nonzero macroscopic magnetization, known as remanent magnetization, even without an external field. In contrast, systems in which magnetic moments interact only via dipole-dipole coupling do not exhibit such behavior. Their net magnetization vanishes when the external field is removed, as predicted by classical electrodynamics [95].

This distinctly ferromagnetic behavior arises from the quantum mechanical exchange interaction. First proposed by W. Heisenberg in 1926 [96], the exchange interaction favors parallel alignment of neighboring electron spins. Its origin lies in the antisymmetrization requirement imposed on the two-electron wavefunction by the Pauli exclusion principle. Depending on the spin alignment, the symmetry of the total wavefunction affects the spatial part, thereby altering the expected Coulomb energy. A symmetric spatial distribution (corresponding to parallel spins) leads to a larger average separation and consequently a lower Coulomb repulsion. This energetic preference underpins ferromagnetism.

The classical description of this effect is given by the Heisenberg model, in which the exchange energy between two localized unit spins $\mathbf{s}_i$ and $\mathbf{s}_j$ is expressed as:

$$\mathcal{E}_{ij}^{\text{ex}} = -J \mathbf{s}_i \cdot \mathbf{s}_j, \quad (3.20)$$

where J is the exchange integral, a material- and distance-dependent parameter. For a detailed derivation of the Heisenberg model, the reader is referred to standard quantum mechanics textbooks [97].

In a discrete spin system, the total exchange energy at a spin site $\mathbf{x}$ is given by:

$$\mathcal{E}_{\mathbf{x}}^{\text{ex}} = - \sum_i \frac{J_i}{2} \mathbf{m}(\mathbf{x}) \cdot \mathbf{m}(\mathbf{x} - \boldsymbol{\xi}_i), \quad (3.21)$$

where $\boldsymbol{\xi}_i$ denotes the displacement vector to neighboring spins and $\mathbf{m}$ is the normalized magnetization.

Using a Taylor expansion and assuming identical coupling $J_i = J$ and uniform spacing $\boldsymbol{\xi}_i$, this discrete expression can be approximated as:

$$\mathcal{E}_{\mathbf{x}}^{\text{ex}} \approx - \sum_i \frac{J}{2} \left[1 - \frac{1}{2} (\nabla \mathbf{m} \cdot \boldsymbol{\xi}_i)^2 \right] + \mathcal{O}(\boldsymbol{\xi}_i^3). \quad (3.22)$$

Neglecting constant terms and higher-order corrections, and transitioning to the continuum limit under the assumption of a regular lattice, the exchange energy becomes:

$$\mathcal{E}^{\text{ex}} = \int_{\Omega} \sum_{i,j,k} A_{jk} \frac{\partial m_i}{\partial x_j} \frac{\partial m_i}{\partial x_k} d\mathbf{x}, \quad (3.23)$$

where A_{jk} are the components of a symmetric positive-definite matrix that reflects the anisotropy of the exchange interaction and Ω is the volume of the ferromagnetic domain. The constant offset from integrating the zeroth-order term is typically omitted, as it does not influence the magnetization dynamics.

The matrix A_{jk} can always be diagonalized by a suitable choice of coordinate system [98], yielding a simplified form:

$$\mathcal{E}^{\text{ex}} = \int_{\Omega} \sum_{i,j} A \left(\frac{\partial m_i}{\partial x_j} \right)^2 d\mathbf{x}. \quad (3.24)$$

For cubic and isotropic lattice structures, the exchange coefficients further reduce to a scalar exchange stiffness A , giving the widely used micromagnetic expression:

$$\mathcal{E}^{\text{ex}} = A \int_{\Omega} \|\nabla \mathbf{m}(\mathbf{x})\|_F^2 d\mathbf{x}, \quad (3.25)$$

where $\|\nabla \mathbf{m}\|_F^2$ is the Frobenius norm squared of the gradient tensor. Although this form was originally derived for localized spins and isotropic lattices, it has been shown to effectively describe a broad range of materials, including band magnets and anisotropic systems [99].

Equation (3.25) represents the lowest-order phenomenological expression that penalizes spatial inhomogeneities in the magnetization vector field. It plays a central role in micromagnetic simulations. The general relation for the effective field is:

$$\mathbf{H}_{\text{eff}} = - \frac{1}{\mu_0 M_s} \frac{\delta E}{\delta \mathbf{m}}, \quad (3.26)$$

the exchange field contribution becomes:

$$\mathbf{H}_{\text{exc}} = \frac{2A}{\mu_0 M_s} \nabla^2 \mathbf{m}. \quad (3.27)$$

3.2.2 Zeeman Field

An external magnetic field $\mathbf{H}_{\text{ext}}$ couples directly to the magnetization and tends to align it along the field direction. The associated energy contribution, known as the *Zeeman energy*, penalizes deviations of the magnetization from the external field and is expressed as:

$$\mathcal{E}_{\text{ext}} = -\mu_0 M_s \int_{\Omega} \mathbf{H}_{\text{ext}} \cdot \mathbf{m} \, d\mathbf{x}. \quad (3.28)$$

Here, $\mathbf{m}$ denotes the normalized magnetization vector and M_s the saturation magnetization. The resulting field contribution to the effective field follows directly from functional differentiation:

$$\mathbf{H}_{\text{ext}} = -\frac{1}{\mu_0 M_s} \frac{\delta \mathcal{E}_{\text{ext}}}{\delta \mathbf{m}}. \quad (3.29)$$

This term represents the most straightforward contribution to the effective field, as it simply reflects the action of an externally applied magnetic field on the magnetic system.

3.2.3 Anisotropy Field

Another vital contribution to the total magnetic free energy is the anisotropy energy, which favors magnetization alignment along specific directions known as *easy axes*. The energy required to reorient the magnetization in anisotropic materials depends on its direction relative to these preferred axes. This direction-dependent behavior arises from several sources, including:

- (i) anisotropic crystal structures (*magnetocrystalline anisotropy*),
- (ii) lattice deformation or spin-orbit coupling at interfaces (*interface anisotropy*),
- (iii) and the geometry of the magnetic body itself (*shape anisotropy*) [99].

These anisotropies are often described phenomenologically by constructing energy expressions based on symmetry arguments. A key property is that the energy must be invariant under magnetization reversal:

$$\mathcal{E}_{\text{ani}}(\mathbf{m}) = \mathcal{E}_{\text{ani}}(-\mathbf{m}). \quad (3.30)$$

3.2.3.1 Uniaxial Anisotropy

A single easy axis $\mathbf{e}$ exists in the simplest case of uniaxial anisotropy. The anisotropy energy density depends only on the angle between the magnetization $\mathbf{m}$ and the easy axis $\mathbf{e}$. A general expansion in even powers yields:

$$\mathcal{E}_{\text{ani}}^{\text{uni}} = \int_{\Omega} [K_{u1}(\mathbf{m} \cdot \mathbf{e})^2 + K_{u2}(\mathbf{m} \cdot \mathbf{e})^4 + \mathcal{O}((\mathbf{m} \cdot \mathbf{e})^6)] \, d\mathbf{x}, \quad (3.31)$$

where K_{u1} and K_{u2} are the first- and second-order uniaxial anisotropy constants. For many applications, the first term suffices:

$$\mathcal{E}_{\text{ani}}^{\text{uni}} = - \int_{\Omega} K_u (\mathbf{m} \cdot \mathbf{e})^2 d\mathbf{x}. \quad (3.32)$$

Here, $K_u > 0$ favors alignment of $\mathbf{m}$ along $\pm \mathbf{e}$, while $K_u < 0$ leads to *easy-plane anisotropy*, where the magnetization tends to lie perpendicular to the axis.

The corresponding anisotropy field is obtained via functional differentiation:

$$\mathbf{H}_{\text{ani}}^{\text{uni}} = \frac{2K_1}{\mu_0 M_s} (\mathbf{m} \cdot \mathbf{e}) \mathbf{e}. \quad (3.33)$$

Uniaxial anisotropy is typically observed in hexagonal or tetragonal crystal structures, such as cobalt. An example energy landscape for uniaxial anisotropy is shown in Figure 3.2 (a).

3.2.3.2 Planar Anisotropy

Planar anisotropy is a special case of uniaxial anisotropy in which the magnetization prefers to lie in a plane perpendicular to a given axis $\mathbf{e}$. This corresponds to a negative uniaxial anisotropy constant $K_u < 0$. The anisotropy energy density takes the same mathematical form as in Equation (3.32), but with the opposite sign:

$$\mathcal{E}_{\text{ani}}^{\text{planar}} = \int_{\Omega} K_p (\mathbf{m} \cdot \mathbf{e})^2 d\mathbf{x}, \quad (3.34)$$

where $K_p = -K_u > 0$ encourages magnetization to remain orthogonal to $\mathbf{e}$.

The corresponding anisotropy field is given by:

$$\mathbf{H}_{\text{ani}}^{\text{planar}} = - \frac{2K_p}{\mu_0 M_s} (\mathbf{m} \cdot \mathbf{e}) \mathbf{e}. \quad (3.35)$$

Planar anisotropy is common in ultrathin magnetic films or multilayers, where the shape or surface anisotropy causes the magnetization to align in-plane. It is effectively modeled by taking the uniaxial formulation with $K_u < 0$. Figure 3.2 (b) illustrates the planar anisotropy surface, where the minimum energy lies in the plane orthogonal to the easy axis.

3.2.3.3 Cubic Anisotropy

In materials with cubic lattice symmetry, such as body-centered cubic iron or face-centered cubic nickel, the magnetocrystalline anisotropy energy depends on the projections of $\mathbf{m}$ onto three orthogonal easy axes $\mathbf{e}_1$, $\mathbf{e}_2$, and $\mathbf{e}_3$, satisfying $\mathbf{e}_i \cdot \mathbf{e}_j = \delta_{ij}$. The anisotropy energy is then expressed as:

$$\mathcal{E}_{\text{ani}}^{\text{cub}} = \int_{\Omega} [K_{c1} (m_1^2 m_2^2 + m_2^2 m_3^2 + m_3^2 m_1^2) + K_{c2} m_1^2 m_2^2 m_3^2] d\mathbf{x}, \quad (3.36)$$




Figure 3.2: Anisotropy functions visualized as energy surfaces. (a) Uniaxial anisotropy with $K_u = 1$ and $\mathbf{a} = \mathbf{e}_3$. (b) Planar anisotropy with $K_p = 1$ and $\mathbf{a} = \mathbf{e}_3$. (c) Cubic anisotropy with $K_{c1} = K_{c2} = 1$, and $\mathbf{a}_i = \mathbf{e}_i$ for all $1 \leq i \leq 3$.

where $m_i = \mathbf{m} \cdot \mathbf{e}_i$ are the projections of the magnetization along the easy axes. The coefficients K_{c1} and K_{c2} are material-dependent anisotropy constants.

As with uniaxial anisotropy, the sign of K_{c1} and K_{c2} determines whether the magnetization prefers to align with or avoid the easy axes. Figure 3.2 (c) visualizes the corresponding energy surface. For instance, negative constants in nickel (fcc structure) lead to multiple easy axes (four in total).

The effective anisotropy field for cubic symmetry can be written in matrix form:

$$\mathbf{H}_{\text{ani}}^{\text{cub}} = -\frac{2}{\mu_0 M_s} \mathbf{D} \mathbf{m}, \quad (3.37)$$

with the matrix $\mathbf{D}$ defined as:

$$\mathbf{D} = \begin{bmatrix} K_{c1}(m_2^2 + m_3^2) + K_{c2}m_2^2m_3^2 & 0 & 0 \\ 0 & K_{c1}(m_1^2 + m_3^2) + K_{c2}m_1^2m_3^2 & 0 \\ 0 & 0 & K_{c1}(m_1^2 + m_2^2) + K_{c2}m_1^2m_2^2 \end{bmatrix}$$

Anisotropy may originate at the interface in thin films or multilayers due to broken crystal symmetry, lattice strain, or modified electronic structure. This *surface anisotropy* is localized at the boundary $\partial\Omega$ and described similarly to the bulk anisotropy energy, but with surface integrals:

$$\mathcal{E}_{\text{ani}}^{\text{surf}} = \int_{\partial\Omega} f(\mathbf{m}, \mathbf{e}) dS, \quad (3.38)$$

where f represents an anisotropy density similar in form to the volume expressions in Equations (3.32), (3.34), and (3.36).

Although the above formulations are phenomenological, they are well established and accurately capture experimental observations [99]. In many practical cases, the lowest-order terms (e.g., $K_{u2} = 0$, $K_{c2} = 0$) provide a sufficient approximation for modeling anisotropy effects.

3.2.4 Demagnetizing Field

In ferromagnetic materials, the magnetization generates a field contribution known as the *demagnetizing field*, which arises from long-range dipole-dipole interactions. This field opposes the magnetization, tending to reduce the systems total magnetic moment. The associated energy term is called the *demagnetization energy* (also referred to as *magnetostatic energy* or *stray-field energy*), and it plays a key role in the formation of magnetic domains.

Unlike the exchange or anisotropy fields, the demagnetizing field is non-local. It depends on both the shape of the magnetic body and the spatial distribution and orientation of the magnetization, thereby complicating its evaluation. Nonetheless, it can be derived from Maxwell's macroscopic equations under the assumption of no free electric currents ($\mathbf{j}_e = \mathbf{0}$), which simplifies them to:

$$\nabla \cdot \mathbf{B} = 0, \quad (3.39a)$$

$$\nabla \times \mathbf{H}_d = \mathbf{0}, \quad (3.39b)$$

where $\mathbf{B}$ is the magnetic flux density and $\mathbf{H}_d$ is the demagnetizing field. Using the constitutive relation:

$$\mathbf{B} = \mu_0(\mathbf{H}_d + \mathbf{M}), \quad (3.40)$$

and noting from Equation (3.39b) that $\mathbf{H}_d$ is conservative, one can define a scalar magnetostatic potential u such that:

$$\mathbf{H}_d = -\nabla u. \quad (3.41)$$

Substituting into Equation (3.39a) yields:

$$\nabla \cdot (-\nabla u + \mathbf{M}) = 0, \quad (3.42)$$

which simplifies to a Poisson equation for u :

$$\Delta u = \nabla \cdot \mathbf{M}. \quad (3.43)$$

Equation (3.43) is solved over the entire space $\mathbb{R}^3$, with the open boundary condition:

$$u(\mathbf{x}) = \mathcal{O}\left(\frac{1}{|\mathbf{x}|}\right), \quad \text{as } |\mathbf{x}| \rightarrow \infty, \quad (3.44)$$

ensuring that the magnetostatic potential vanishes at infinity [100, 101]. In analogy to electrostatics, one can define effective magnetic charges:

$$\rho_m = -\nabla \cdot \mathbf{M} \quad (\text{volume magnetic charge}), \quad (3.45a)$$

$$\sigma_m = \mathbf{M} \cdot \mathbf{n} \quad (\text{surface magnetic charge}), \quad (3.45b)$$

where $\mathbf{n}$ is the unit normal vector to the boundary $\partial\Omega$ of the magnetic body. Using these definitions, the magnetostatic potential u at position $\mathbf{x}$ can be expressed as:

$$u(\mathbf{x}) = \frac{1}{4\pi} \left(\int_{\Omega} \frac{\rho_m(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d\mathbf{x}' + \int_{\partial\Omega} \frac{\sigma_m(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d\mathbf{s}' \right). \quad (3.46)$$

From Equation (3.41), the demagnetizing field then follows as:

$$\mathbf{H}_d(\mathbf{x}) = -\nabla u(\mathbf{x}). \quad (3.47)$$

The demagnetization (or magnetostatic) energy represents the interaction energy of the magnetization with its stray field and is given by:

$$\mathcal{E}_d = -\frac{\mu_0 M_s}{2} \int_{\Omega} \mathbf{m} \cdot \mathbf{H}_d \, d\mathbf{x}, \quad (3.48)$$

where the factor 1/2 avoids double-counting the self-interaction.

Alternatively, in numerical micromagnetics (e.g., finite difference methods), it is often convenient to express the demagnetizing field as a convolution:

$$\mathbf{H}_d(\mathbf{x}) = - \int_{\Omega} \mathbf{N}(\mathbf{x} - \mathbf{x}') \cdot \mathbf{m}(\mathbf{x}') \, d\mathbf{x}', \quad (3.49)$$

where $\mathbf{N}$ is the demagnetizing tensor kernel that encodes the dipole-dipole interaction [102]. The components N_{ij} of this 3×3 tensor are given by [103]:

$$N_{ij}(\mathbf{x} - \mathbf{x}') = \frac{M_s}{4\pi} \frac{3r_i r_j - \delta_{ij}}{|\mathbf{x} - \mathbf{x}'|^3}, \quad (3.50)$$

where $r_i = (x_i - x'_i)/|\mathbf{x} - \mathbf{x}'|$ and $r_j = (x_j - x'_j)/|\mathbf{x} - \mathbf{x}'|$ are the directional cosines of the displacement vector, and $i, j = x, y, z$ denote Cartesian components. This long-range, shape- and geometry-sensitive contribution is crucial in domain formation and stability analysis in micromagnetic systems.

3.2.5 Ampère Field

When a current flows through the magnetic domain, it generates an additional contribution to the effective magnetic field, usually referred to as the *Ampère field*. This field arises from the interaction between the electric current and the surrounding magnetic material and is governed by the classical Biot-Savart law [97, 104]. For a given current density $\mathbf{J}_C$, the Ampère field $\mathbf{H}_{\text{curr}}$ at position $\mathbf{x}$ is computed as:

$$\mathbf{H}_{\text{curr}}(\mathbf{x}) = \frac{1}{4\pi} \int_{\Omega} \frac{\mathbf{J}_C(\mathbf{x}') \times (\mathbf{x} - \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|^3} \, d\mathbf{x}', \quad (3.51)$$

where the integral is taken over the current-carrying domain Ω . This field contributes to the total effective field used in the LLG equation and becomes particularly relevant in spintronic devices such as MRAM.

3.2.6 Thermal Field

At a nonzero temperature, the magnetization of a ferromagnetic material experiences thermal agitation. These thermal fluctuations can be incorporated into the LLG formalism by introducing an auxiliary, randomly varying term in the effective magnetic

field, commonly referred to as the thermal field. This additional field accounts for the stochastic nature of thermal excitations and is modeled as a Gaussian random vector field with specified statistical properties.

The thermal field $\mathbf{H}_{\text{th}}$ must satisfy the following conditions [105, 106]:

$$\langle H_{\text{th}}^i(t) \rangle = 0, \quad (3.52a)$$

$$\langle H_{\text{th}}^i(t), H_{\text{th}}^j(t') \rangle = D \delta_{ij} \delta(t - t'), \quad (3.52b)$$

where i and j denote Cartesian components, $\langle \cdot \rangle$ indicates the ensemble average over different realizations of the stochastic process, δ_{ij} is the Kronecker delta, and $\delta(t - t')$ is the Dirac delta function. These relations describe a field that is uncorrelated in both time and space.

The prefactor D quantifies the strength of the thermal fluctuations and is derived from the fluctuation-dissipation theorem [107]:

$$D = \frac{2 \alpha k_B T}{\mu_0 M_s \gamma V}, \quad (3.53)$$

where α is the Gilbert damping parameter, k_B is the Boltzmann constant, T is the absolute temperature, M_s is the saturation magnetization, γ is the gyromagnetic ratio, and V is the volume of the computational cell. The thermal field is generated using Gaussian white noise with a zero mean and a variance proportional to D . Its presence leads to a stochastic extension of the LLG equation, often referred to as the stochastic LLG or Langevin LLG model.

3.2.7 Interlayer Exchange Coupling

Magnetic layers within a multilayer stack can experience exchange coupling when separated by a nonmagnetic spacer (NMS), and even across thin TBs. This interaction, initially proposed by Ruderman and Kittel [108], is mediated by conduction electrons of the nonmagnetic spacer. The theory was later extended by Kasuya [109] and Yosida [110], leading to the well-known RKKY interaction, which predicts an oscillatory exchange coupling as a function of spacer thickness, alternating between FM and antiferromagnet (AFM) behavior.

The nature of IEC depends on the electronic properties of the spacer:

- (i) For metallic spacers, the coupling is well-described by the RKKY theory, which explains the indirect interaction between localized spins in the d - or f -orbitals via conduction electrons [111].
- (ii) For semiconducting spacers, exchange interactions are governed by mechanisms such as variable-range hopping or resonant tunneling, mediated by localized states within the band gap [17, 112].
- (iii) In the case of insulating spacers, strong and non-oscillatory IEC is observed, often attributed to spin-dependent tunneling [113, 114].

Two theoretical formulations are commonly used to account for these behaviors across different spacer types: discrete spin models such as the Heisenberg exchange interaction, and effective energy expressions such as the one proposed by Bruno [111]. While the Heisenberg model describes pairwise interactions between localized spins, Bruno's formulation captures the angular dependence of the coupling energy in a continuum framework. It effectively coarse-grains the microscopic exchange by summarizing spin-dependent electron reflections at magnetic interfaces [111].

Bruno proposed the following expression for the interlayer coupling energy, based on the angle $\Delta\phi$ between the magnetization of adjacent magnetic layers (see Figure 3.3):

$$\mathcal{E} = -J_1 \cos(\Delta\phi) - J_2 \cos^2(\Delta\phi), \quad (3.54)$$

where J_1 and J_2 are coupling constants. The term J_1 represents the bilinear coupling (dominant in most systems), responsible for the oscillatory behavior seen in metallic spacers:

- (i) $J_1 > 0$: FM coupling, favoring parallel alignment ($\Delta\phi = 0$)
- (ii) $J_1 < 0$: AFM coupling, favoring anti-parallel alignment ($\Delta\phi = \pi$)

The biquadratic term J_2 is often smaller and associated with structural inhomogeneities such as surface roughness or defects, and can induce non-collinear alignment of the magnetizations [115]. When J_2 is negligible, the expression simplifies to:

$$\mathcal{E} = -J_1 \cos(\Delta\phi), \quad (3.55)$$

highlighting the pure bilinear coupling as in classical RKKY-type interactions.

In the micromagnetic continuum approximation, IEC is modeled as an interaction between the two magnetic interfaces $\partial\Omega_1$ and $\partial\Omega_2$ on either side of the spacer, which are assumed to be of equal area. Integrating the discrete Heisenberg interaction (3.20) over these surfaces yields the following interlayer exchange energy:

$$\mathcal{E}_{\text{icx}} = - \int_{\partial\Omega_1} A \mathbf{m}(\mathbf{x}) \cdot \mathbf{m}(P(\mathbf{x})) \, ds, \quad (3.56)$$

A is the effective interlayer exchange constant, and $P : \partial\Omega_1 \rightarrow \partial\Omega_2$ is a mapping between the corresponding surface points on both interfaces. This continuum representation is consistent with the bilinear term in Bruno's angular model and arises naturally from the coarse-grained limit of the Heisenberg interaction.

This energy formulation underpins the design of SAFs, where the spacer thickness is tuned to promote AFM coupling (i.e., negative A). SAFs are widely implemented in spintronic devices such as MTJs and STT-MRAMs, offering enhanced thermal stability and reduced stray fields.

While Equation (3.56) defines the physical energy contribution, its numerical evaluation in a finite element framework poses a challenge: the non-magnetic spacer is often too thin (< 1 nm) to be explicitly discretized with volume elements without inducing high computational costs. Therefore, in this work, this interaction is implemented as




Figure 3.3: Schematic depiction of a trilayer structure composed of left and right semi-infinite FM regions, separated by a NMS layer. The interface magnetization on the left FM interface points into an arbitrary direction, whereas the interface magnetization in the right FM forms an angle $\Delta\phi$ relative to the magnetization of the left FM. The magnetization vectors originate at the respective interface regions $\partial\Omega_1$ and $\partial\Omega_2$, which are highlighted by hatching. Figure based on the figure originally published in [116].

a specialized boundary condition in the weak form of the effective field. The details of this implementation, including the interface mapping algorithm $P(\mathbf{x})$, are presented in Section 6.4.

In summary, the microscopic (Heisenberg) and mesoscopic (Bruno) representations of IEC are deeply connected. The discrete Heisenberg exchange interaction captures the fundamental spin-spin coupling mechanism at the atomic level. Bruno’s angular model, derived using quantum perturbation theory applied to spin-dependent electron reflection at magnetic interfaces [111], effectively approximates this interaction and expresses the total IEC energy as a function of angular misalignment between magnetizations.

This interaction is subsequently formulated within a micromagnetic framework, bridging theory and computation via a surface integral across magnetic interfaces. This continuum form provides a spatially averaged representation of the bilinear IEC and serves as the theoretical foundation for our numerical approach.

3.3 Modelling Spin-Transfer Torque

A commonly used model for torque in STT-MRAM devices is the macro-spin model proposed by Slonczewski [117]. Here, the dynamics of the FL magnetization are described as behaving like a single spin influenced by the polarization occurring in a polarizing layer. The LLG equation is extended by such a torque term and takes the following form:

$$\frac{\partial \mathbf{m}}{\partial t} = -\gamma_0 \mathbf{m} \times \mathbf{H}_{\text{eff}} + \alpha \mathbf{m} \times \frac{\partial \mathbf{m}}{\partial t} + \frac{1}{M_s} \mathbf{T}_s, \quad (3.57)$$

where $\mathbf{T}_S$ is the spin torque originating from the angular momentum transferred from spin-polarized electrons in the reference layer.

The spin torque $\mathbf{T}_S$ can be expressed as a combination of damping-like and field-like components:

$$\mathbf{T}_S = \underbrace{a\mathbf{m} \times (\mathbf{m} \times \mathbf{p})}_{\text{damping-like}} + \underbrace{b(\mathbf{m} \times \mathbf{p})}_{\text{field-like}}, \quad (3.58)$$

where parameters a and b depend on the current density and $\mathbf{p}$ represents the magnetization direction of the RL. Based on this, the LLG equation becomes:

$$\partial_t \mathbf{m} = \frac{1}{1 + \alpha^2} \left[-\gamma_0 (\mathbf{m} \times \mathbf{H}_{\text{eff}}) - \gamma_0 \alpha (\mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}})) - \frac{\gamma \hbar J_C}{e M_S d_{FL}} \eta(\theta) (\mathbf{m} \times (\mathbf{m} \times \mathbf{p}) - \beta (\mathbf{m} \times \mathbf{p})) \right], \quad (3.59)$$

where J_C is the current density, d_{FL} is the thickness of the free layer, β quantifies the field-like torque strength, and $\eta(\theta)$ is the torque efficiency function that depends on the angular relationship between $\mathbf{m}$ and $\mathbf{p}$ [35, 118], further on whether a spin-valve or an MTJ is being described [119].

While this macro-spin model provides insight into STT dynamics, it assumes spatially uniform magnetization, making it less accurate for larger devices. To improve this, both $\mathbf{m}$ and $\mathbf{p}$ can be treated as spatially dependent fields. However, the model in Equation (3.59) neglects spin diffusion effects, which are crucial for specific materials and configurations [90].

A more general formulation considers the spin-accumulation vector $\mathbf{S}$, representing the non-equilibrium polarization of conducting electrons. The torque due to $\mathbf{S}$ is:

$$\mathbf{T}_S = -D_e \frac{\mathbf{m} \times \mathbf{S}}{\lambda_J^2} - D_e \frac{\mathbf{m} \times (\mathbf{m} \times \mathbf{S})}{\lambda_\varphi^2}, \quad (3.60)$$

where D_e is the electron diffusion coefficient, λ_J is the exchange length, and λ_φ is the spin dephasing length. The first term models the precession of spins in the local exchange field, while the second represents dephasing [32, 120, 121, 122].

The steady-state spin-accumulation satisfies the spin current continuity equation:

$$\partial_t \mathbf{S} = 0 = -\nabla \cdot \mathbf{J}_S - D_e \left(\frac{\mathbf{S}}{\lambda_{sf}^2} + \frac{\mathbf{S} \times \mathbf{m}}{\lambda_J^2} + \frac{\mathbf{m} \times (\mathbf{S} \times \mathbf{m})}{\lambda_\varphi^2} \right), \quad (3.61)$$

where λ_{sf} is the spin-flip length and $\mathbf{J}_S$ is the spin current. The finite element solver implemented in this work solves this equation in both ferromagnetic and nonmagnetic regions [120, 123, 124, 125]. The drift-diffusion formalism introduced here provides the theoretical foundation for computing spin-transfer torques from first principles. However, extending this semi-classical approach to magnetic tunnel junctions requires special treatment of the tunneling process via appropriate boundary conditions at the tunnel-barrier interfaces. The complete derivation of the spin and charge drift-diffusion equations, their finite element implementation, the treatment of tunneling spin currents, and validation against analytical and experimental results are presented in detail in Section 6.1.

Computational Methods

The micromagnetic model introduced in the preceding sections defines a set of nonlinear partial differential equations in space and time that are analytically tractable only for highly simplified cases. In practice, the simulation of static and dynamic micromagnetic phenomena relies on numerical techniques. However, developing efficient numerical methods is nontrivial due to several intrinsic challenges of the underlying equations.

First, the demagnetization field, which represents long-range dipole-dipole interactions, introduces a computational bottleneck. A simple evaluation scales as $\mathcal{O}(n^2)$ quickly becomes restrictive with the number of simulation cells n . To address this, several approaches have been proposed that reduce the complexity to $\mathcal{O}(n \log n)$ [126] or even $\mathcal{O}(n)$ [127].

Second, the exchange interaction imposes a local but highly stiff coupling due to its second-order spatial derivative. This stiffness, combined with the nonlinearity of the governing equations, places strict requirements on the numerical time integration scheme.

Third, the total micromagnetic energy landscape is typically nonconvex and high-dimensional, with multiple local minima. In quasistatic problems such as hysteresis simulations, the system must relax toward a nearby energy minimum. The complexity of the energy landscape necessitates careful selection of minimization algorithms to avoid missing relevant minima.

To solve the micromagnetic equations numerically, spatial and temporal discretizations are typically decoupled. For spatial discretization, the most widely used methods are the finite difference method (FDM) and the finite element method (FEM). Both approaches partition the magnetic domain into cells to compute the local contributions to the effective field $\mathbf{H}_{\text{eff}}$ or the total energy E . While FDM is typically based on a regular Cartesian grid, FEM uses tetrahedral elements to accommodate complex geometries (see Figure 4.1).

The mesh resolution in either method must be fine enough to resolve domain walls, whose characteristic width is given by the exchange length:

$$\ell = \sqrt{\frac{A}{K_{\text{eff}}}}, \quad (4.1)$$

where A is the exchange stiffness and K_{eff} denotes the effective anisotropy constant. The latter includes crystalline anisotropy and shape anisotropy contributions induced by the demagnetization field [99].




Figure 4.1: Discretization of a unit sphere using two common methods in micromagnetics: (a) structured Cartesian mesh used in FDM, and (b) unstructured tetrahedral mesh used in FEM.

Once the system is spatially discretized, the static energy minimum or the dynamic time evolution of the magnetization $\mathbf{m}$ must be computed. This requires a separate class of algorithms, such as energy-minimization techniques or time integration methods for the LLG equation. The effective field must be evaluated only within the magnetic domain in both cases.

The following sections review two spatial discretization techniques: the FDM (Section 4.1) and the FEM (Section 4.2). Particular attention is given to efficient algorithms for computing the demagnetizing field, including the truncation approach and the hybrid finite element-boundary element method (FE-BEM), which are presented in Section 4.3.1. The numerical time integration of the LLG equation is treated separately in Chapter 5.

4.1 Finite Difference Method

The FDM is a widely used numerical technique for solving differential equations and has become a cornerstone of computational science due to its conceptual simplicity and ease of implementation. At its core, FDM seeks to approximate solutions to differential equations, i.e., to find a function (or a discrete approximation) that satisfies the prescribed relationships between various derivatives on a given spatial and/or temporal domain, subject to boundary conditions. This is a nontrivial problem, and analytical solutions are rarely obtained. The finite difference method addresses this challenge by replacing derivatives with finite difference approximations. This approximation yields a large but finite system of algebraic equations that can be solved numerically on a computer [128, 129, 130].

Before solving such systems, one must first understand how to approximate derivatives of a known function using only discrete pointwise evaluations. This foundational step enables the development of higher-order schemes and introduces essential concepts, including the consistency, stability, and accuracy of finite-difference operators. The




Figure 4.2: Mesh points for the FDM grid in one (a) and two (b) spatial dimensions. The grid spacing is uniform with h between adjacent points. In (b), the 2D indexing convention for a 3×3 grid is shown.

discretization typically involves subdividing the spatial domain into an equidistant grid of points. In the 1D case, this corresponds to grid points:

$$x_i = ih, \quad \text{for } i = 0, \dots, N, \quad (4.2)$$

where $h = \frac{1}{N}$ is the uniform grid spacing, and $N \in \mathbb{N}$ is the number of discretization points. This concept naturally extends to higher dimensions. For example, in a 2D setting, the grid points are given by:

$$\mathbf{x}_{i,j} = (ih, jh), \quad \text{for } i, j = 0, \dots, N, \quad (4.3)$$

with the same spacing $h = \frac{1}{N}$ applied along both axes.

Figure 4.2 illustrates the grid layout in both the 1D and 2D cases. In 1D (Figure 4.2 (a)), the points lie along a line with uniform spacing. In 2D (Figure 4.2 (b)), the points form a Cartesian mesh.

In micromagnetics, FDM has proven to be a robust and efficient tool for simulating magnetization dynamics. Its restriction to regular, structured meshes enables the use of highly optimized algorithms, particularly for computing quantities like the demagnetization field, which are computationally expensive due to their long-range nature. However, this same grid regularity limits the applicability of FDM to geometries that conform well to Cartesian meshes, making FDM less suitable for complex or curved domains, where unstructured methods such as FEM are more appropriate.

FDM is a classical numerical technique for approximating the solution of differential equations by discretizing space into a uniform grid and replacing derivatives with finite difference quotients. If one can efficiently and accurately approximate derivatives at discrete points, then in principle, one can also approximate the solution of differential equations. This process begins by approximating derivatives of a real-valued function $u: [a, b] \rightarrow \mathbb{R}$ using difference formulas:

(i) **Forward difference:**

$$\frac{u(x+h) - u(x)}{h} \quad (4.4)$$

(ii) **Backward difference:**

$$\frac{u(x) - u(x - h)}{h} \quad (4.5)$$

(iii) **Central difference:**

$$\frac{u(x + h) - u(x - h)}{2h} \quad (4.6)$$

The central difference formula is generally more accurate, introducing an error term of order $\mathcal{O}(h^2)$, while forward and backward differences have errors of order $\mathcal{O}(h)$. These formulas can be derived from Taylor expansions or from interpolating polynomials evaluated at the mesh points.

This idea extends to second-order derivatives. In one spatial dimension, the second derivative of a function $u(x)$ at a point x_i is approximated by applying the central difference formula twice, once for the first derivative and again for the derivative of that result. This yields the second-order central difference approximation:

$$\left. \frac{d^2 u}{dx^2} \right|_{x_i} \approx \frac{u_{i-1} - 2u_i + u_{i+1}}{h^2} + \mathcal{O}(h^2), \quad (4.7)$$

where h is the grid spacing and u_i denotes the discrete function value at grid point x_i , i.e. $u_i \approx u(x_i)$. This approximation introduces a truncation error of order $\mathcal{O}(h^2)$, making it simple and reasonably accurate for many applications.

To illustrate how FDM works in practice, we consider the Poisson equation, a fundamental equation arising in many scientific and engineering applications:

$$-\nabla^2 u(x) = f(x), \quad x \in \Omega, \quad (4.8)$$

$$u(x) = u_D(x), \quad x \in \partial\Omega. \quad (4.9)$$

In one dimension, the domain $[0, 1]$ is discretized into a uniform mesh of $N \in \mathbb{N}$ grid points, with spacing $h = \frac{1}{N}$. Using the finite difference approximation, the continuous Poisson equation becomes:

$$-\frac{u_{i-1} - 2u_i + u_{i+1}}{h^2} = f_i, \quad i = 1, \dots, N-1. \quad (4.10)$$

This leads to a linear system of equations:

$$Au = r, \quad (4.11)$$

where A is a sparse matrix representing the discrete Laplace operator, u is the vector of unknowns, and r is the discrete right-hand side vector.

Dirichlet boundary conditions can be directly incorporated by assigning values to the boundary nodes u_0 and u_N . To evaluate the difference quotient at the domain boundaries, one might use *ghost* points, with values such as:

$$u_{-1} = u_1, \quad u_{N+1} = u_{N-1}. \quad (4.12)$$

This ensures the scheme can be applied across the full grid. Once the system is assembled, the solution can be computed as:

$$u = A^{-1}r, \quad (4.13)$$

in practice, iterative solvers such as Jacobi, Gauss-Seidel, or successive over-relaxation (SOR) are typically used due to the size and sparsity of A [131, 132].

To extend finite difference methods to higher dimensions, consider the Poisson equation in two dimensions:

$$-\nabla^2 u(x, y) = -\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}\right) = f(x, y), \quad (x, y) \in \Omega = (0, 1) \times (0, 1), \quad (4.14)$$

$$u(x, y) = g(x, y), \quad (x, y) \in \partial\Omega. \quad (4.15)$$

Here, $g(x, y)$ defines the Dirichlet boundary condition on the domain boundary. As in the 1D case, we discretize the domain by defining a mesh that includes the boundary:

$$\mathcal{M}_{M,N} = \left\{ (x_i, y_j) = \left(\frac{i}{M}, \frac{j}{N} \right) \mid 0 \leq i \leq M, 0 \leq j \leq N \right\}, \quad (4.16)$$

where $M, N \in \mathbb{N}$ are the number of intervals along each axis.

To discretize the differential operator, we approximate the second derivatives in each spatial direction using central finite differences:

$$\frac{\partial^2 u}{\partial x^2}(x_i, y_j) \approx \frac{u_{i+1,j} - 2u_{i,j} + u_{i-1,j}}{h^2}, \quad (4.17)$$

$$\frac{\partial^2 u}{\partial y^2}(x_i, y_j) \approx \frac{u_{i,j+1} - 2u_{i,j} + u_{i,j-1}}{h^2}. \quad (4.18)$$

By adding these two approximations, the full Laplacian is approximated as:

$$\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}\right)(x_i, y_j) \approx \frac{u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1} - 4u_{i,j}}{h^2}. \quad (4.19)$$

Substituting this into the 2D Poisson equation results in a discrete equation for interior points:

$$-\frac{u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1} - 4u_{i,j}}{h^2} = f_{i,j}. \quad (4.20)$$

The boundary values are set using the function $g(x, y)$, just like in the 1D case. The resulting linear system again has a sparse matrix structure and can be solved using iterative solvers designed for large-scale problems.

The extension of this method to three dimensions is conceptually straightforward. The 3D Laplacian is defined as:

$$\nabla^2 u(x, y, z) = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2}. \quad (4.21)$$

It can be discretized by independently applying the same finite difference approximations in all three spatial directions. This yields a discrete equation involving the values of u at neighboring grid points along the x , y , and z directions. The resulting system of equations has a similar sparse structure and is solved with the same numerical techniques used in lower dimensions.

4.1.1 Discretization of the Landau-Lifshitz-Gilbert Equation

The simulations in this work are conducted using an in-house micromagnetic solver based on the FDM, originally introduced in [103]. It is applied to the modeling of pMTJ devices. The effective magnetic field is computed directly from the discretized magnetization [133]. Only the dynamics of the FL are simulated, while the magnetostatic stray field of the RL is treated as a static external contribution.

The computational domain is discretized into parallelepiped cells with volume $\Delta x \cdot \Delta y \cdot \Delta z$, where Δx , Δy , and Δz represent the spatial resolution in each Cartesian direction. The number of cells in each direction is denoted by N_x , N_y , and N_z , respectively. A unit magnetization vector $\mathbf{m}$ is assigned to the center of each cell.

4.1.1.1 External Field

The external magnetic field is specified as a uniform, time-independent input:

$$\mathbf{H}_{\text{ext}}(i, j, k) = \mathbf{H}_{\text{ext,input}} \quad (4.22)$$

4.1.1.2 Exchange Field

The exchange field contribution, being proportional to the Laplacian of the magnetization as shown in Equation (3.27), requires knowledge of the magnetization in the computational cell and its six nearest neighbors, two along each Cartesian direction. In a regular FDM grid, these neighbors are located at the adjacent face, sharing positions, specifically at coordinates $(i \pm 1, j, k)$, $(i, j \pm 1, k)$, and $(i, j, k \pm 1)$.

The first-order discretization of the exchange field at cell (i, j, k) yields the following expression [107]:

$$\mathbf{H}_{\text{ex}}(i, j, k) = \frac{2A}{\mu_0 M_s} \sum_{(i', j', k')} \frac{\mathbf{m}(i', j', k') - \mathbf{m}(i, j, k)}{\Delta_{i', j', k'}^2}, \quad (4.23)$$

where A is the exchange stiffness constant, μ_0 the vacuum permeability, M_s the saturation magnetization, and $\Delta_{i', j', k'} \in \{\Delta x, \Delta y, \Delta z\}$ denotes the grid spacing in the direction connecting the central and neighboring cells.

At the boundaries of the computational domain, some neighboring positions may lie outside the simulation grid. To preserve the consistency of the exchange field calculation near these boundaries, so-called *ghost cells* were introduced. These are virtual cells placed just outside the domain, with their magnetization vectors set equal to the nearest valid interior cell, thereby enforcing Neumann-like boundary conditions as detailed in Section 4.1. This approach ensures that the exchange field can be consistently computed throughout the domain, including at the boundaries [129].

4.1.1.3 Anisotropy Field

The computation of the discretized anisotropy field contribution only requires the magnetization vector within the computational cell. Equations (3.33) and (3.37) can

be directly applied for field evaluation. In the present work, the interface-induced perpendicular anisotropy, typical in pMTJs, is modeled as a uniaxial anisotropy contribution. The corresponding discretized expression for the anisotropy field is given by:

$$\mathbf{H}_{\text{ani}}(i, j, k) = \frac{2K_1}{\mu_0 M_s} (\mathbf{m}(i, j, k) \cdot \mathbf{e}) \mathbf{e}, \quad (4.24)$$

where $\mathbf{e}$ is the unit vector defining the easy axis of magnetization, typically oriented perpendicular to the plane of the structure in systems exhibiting PMA.

For materials characterized by cubic anisotropy, the following expression gives the field contribution:

$$\mathbf{H}_{\text{ani}}(i, j, k) = -\frac{2D(i, j, k)}{\mu_0 M_s} \mathbf{m}(i, j, k), \quad (4.25)$$

where $D(i, j, k)$ is a second-rank anisotropy matrix that defines the local anisotropy energy landscape, the construction and role of this matrix are described in detail in Section 3.2.3.3.

4.1.1.4 Demagnetizing Field

The demagnetizing field is the most computationally expensive contribution to the effective field. Originating from dipole-dipole interactions, it is inherently non-local: its value at a given location depends on the magnetization vectors throughout the magnetic structure. Consequently, calculating the demagnetizing field at each cell requires summing contributions from all other cells in the domain.

For the computation of the demagnetizing field in the FL, a discretized form of Equation (3.46) was used. The field at each computational node is calculated as:

$$\mathbf{H}_d(i, j, k) = \frac{M_s}{4\pi} \sum_{i'=1}^{N_x} \sum_{j'=1}^{N_y} \sum_{k'=1}^{N_z} \tilde{G}(i, j, k, i', j', k') \cdot \mathbf{m}(i', j', k'), \quad (4.26)$$

where $\tilde{G}$ is a space-dependent tensor that encodes the dipole-dipole interaction kernel [107, 133]. This tensor is symmetric, and given its structure, only six of its nine components need to be evaluated explicitly. Since it is time-independent and solely geometry-dependent, $\tilde{G}$ is precomputed once at the beginning of the simulation and reused in subsequent time steps.

The magnetostatic coupling field between the RL and the FL shares the exact physical origin as the demagnetizing field in the FL, arising from dipole-dipole interactions. However, because the RL magnetization remains static throughout the simulation, its contribution can be computed once during initialization and treated as a time-invariant external field.

To evaluate this coupling, the RL is discretized into parallelepiped cells, each assigned a unit magnetization vector $\mathbf{p} = \mathbf{M}_{\text{RL}}/M_s^{\text{RL}}$, where $\mathbf{M}_{\text{RL}}$ denotes the local magnetization and M_s^{RL} is the saturation magnetization of the RL.

The resulting magnetostatic field at a given FL cell (i, j, k) is then computed using an expression analogous to Equation (4.26):

$$\mathbf{H}_{\text{ms}}(i, j, k) = \frac{M_s}{4\pi} \sum_{i'=1}^{N_x^{RL}} \sum_{j'=1}^{N_y^{RL}} \sum_{k'=1}^{N_z^{RL}} \tilde{G}(i, j, k, i', j', k') \cdot \mathbf{p}(i', j', k'), \quad (4.27)$$

where N_x^{RL} , N_y^{RL} , and N_z^{RL} represent the number of discretized RL cells along the x -, y -, and z -directions, respectively.

By precomputing this static contribution, the magnetostatic interaction between RL and FL can be incorporated efficiently without increasing the computational burden during the dynamic simulation steps.

4.1.1.5 Ampère Field

The Ampère (or Oersted) field, resulting from the magnetic field generated by the flowing current, is computed using the discretized form of Equation (3.46). The field at each computational cell (i, j, k) is obtained by summing the contributions from the current density in all other cells of the simulation domain:

$$\mathbf{H}_{\text{curr}}(i, j, k) = \frac{1}{4\pi} \sum_{i'} \sum_{j'} \sum_{k'} \mathbf{J}_c(i', j', k') \times \mathbf{G}'(i, j, k, i', j', k'), \quad (4.28)$$

where $\mathbf{J}_c$ is the local current density vector and $\mathbf{G}'$ is a space-dependent interaction vector that encodes the geometric relationship between the source and observation cells. The summation excludes the self-term, i.e., cells for which $(i', j', k') = (i, j, k)$ are omitted.

A detailed derivation of the interaction vector $\mathbf{G}'$, including methods for the efficient numerical evaluation of the required integrals, is provided in [134]. In practice, the Oersted field is approximated at the start of the simulation under the assumption of a constant current-density distribution. It is then treated as a static external field during time integration.

4.1.1.6 Thermal Field

The thermal field is modeled as spatially and temporally uncorrelated Gaussian noise, consistent with the stochastic formulation of the LLG equation presented in [135]. Its implementation must also adhere to the fluctuation-dissipation relation described by Equation (3.52b), ensuring proper thermal agitation behavior across the simulation domain.

The discretized thermal field at each computational cell (i, j, k) is computed as:

$$\mathbf{H}_{\text{th}}(i, j, k) = \sigma(i, j, k) \sqrt{\frac{\alpha}{1 + \alpha^2} \cdot \frac{2k_B T}{\gamma_0 \Delta t \Delta V M_s}}, \quad (4.29)$$

where α is the Gilbert damping factor, k_B the Boltzmann constant, T the temperature, γ_0 the gyromagnetic ratio, Δt the simulation time step, ΔV the volume of the

computational cell, and M_s the saturation magnetization. The term $\sigma(i, j, k)$ represents a random variable drawn from a standard normal distribution, uncorrelated in space and time, with zero mean and unit standard deviation.

The implementation follows the approach in [136], in which each component of the thermal field vector is independently generated from Gaussian noise. This ensures that the thermal fluctuations correctly reflect the statistical properties of a system in thermal equilibrium and preserve the physical fidelity of thermally activated switching processes.

4.1.1.7 Time Integration Scheme

After the initial setup phase, where simulation parameters are loaded from either a configuration file or the command line, the magnetostatic coupling field $\mathbf{H}_{\text{ms}}$ and the Ampère field $\mathbf{H}_{\text{curr}}$ are calculated. Since these contributions remain constant throughout the simulation, they are precomputed and stored as static fields.

The core of the time evolution is handled by an explicit fourth-order Runge-Kutta (RK4) integration scheme, used to solve the LLG equation. Within each time step, the magnetization update proceeds via the evaluation of four successive intermediate stages:

$$\mathbf{k}_1 = \text{LLG}(\mathbf{m}_n, t_n) \Delta t, \quad (4.30a)$$

$$\mathbf{k}_2 = \text{LLG} \left(\mathbf{m}_n + \frac{\mathbf{k}_1}{2}, t_n + \frac{\Delta t}{2} \right) \Delta t, \quad (4.30b)$$

$$\mathbf{k}_3 = \text{LLG} \left(\mathbf{m}_n + \frac{\mathbf{k}_2}{2}, t_n + \frac{\Delta t}{2} \right) \Delta t, \quad (4.30c)$$

$$\mathbf{k}_4 = \text{LLG}(\mathbf{m}_n + \mathbf{k}_3, t_n + \Delta t) \Delta t. \quad (4.30d)$$

Using these coefficients, the magnetization is updated as:

$$\mathbf{m}_{n+1} = \mathbf{m}_n + \frac{\Delta t}{6} (\mathbf{k}_1 + 2\mathbf{k}_2 + 2\mathbf{k}_3 + \mathbf{k}_4). \quad (4.31)$$

This method balances accuracy and computational efficiency, making it well-suited for simulating magnetization dynamics over long time scales. The time loop continues until the desired total simulation time is reached, with field contributions and magnetization consistently updated at each step.

Further emphasis on the mathematical formulation, stability considerations, and numerical performance of time integration methods will be provided in Chapter 5.

Figure 4.4a illustrates the overall flow of this simulation framework.

4.2 Finite Element Method

Similar to the FDM, the FEM is a numerical method for approximating solutions to partial differential equations (PDEs). However, the FEM offers significantly greater flexibility in handling complex computational domains. To solve a problem using the

FEM, the simulation domain is discretized into a mesh composed of smaller, simpler subdomains known as finite elements. These elements are connected at nodes, which are points shared by two or more elements, along with boundary lines or surfaces. This discretization transforms the original PDE into a system of algebraic equations by formulating the problem locally for each element and assembling the local systems into a global system. The following sections will introduce the FEM by outlining the key steps involved in this process. To illustrate how the FEM works in practice, we consider the Poisson equation again:

$$-\nabla^2 u(x) = f(x), \quad x \in \Omega, \quad (4.32)$$

$$u(x) = u_D(x), \quad x \in \Gamma_D, \quad (4.33)$$

$$\frac{\partial u}{\partial n}(x) = g(x), \quad x \in \Gamma_N, \quad (4.34)$$

$$\partial\Omega = \Gamma_D \cup \Gamma_N, \quad \Gamma_D \cap \Gamma_N = \emptyset,$$

where u is the unknown scalar field, f is a known source term, u_D defines the Dirichlet boundary condition on Γ_D , and g prescribes the Neumann condition on Γ_N . This also serves as a blueprint for more complex equations, such as the LLG equation. For general background on the finite element method, see [137, 138, 139, 140].

4.2.1 Weak Formulation and Variational Form

The first step in the FEM procedure is to rewrite the original PDE in its *weak formulation*. This approach relaxes the regularity requirements on the solution and test functions by integrating the governing equation with respect to a set of smooth test functions. We multiply both sides of the equation by a test function v and integrate over the domain Ω :

$$\int_{\Omega} (-\nabla^2 u) v \, dx = \int_{\Omega} f v \, dx. \quad (4.35)$$

Applying integration by parts, Gauss's (divergence) theorem:

$$\int_{\Omega} \nabla u \cdot \nabla v \, dx = - \int_{\Omega} (\Delta u) v \, dx + \int_{\partial\Omega} \frac{\partial u}{\partial n} v \, ds$$

and using the product rule for differentiation, the second-order derivative term is expressed as a sum of first-order derivatives. Here, the left-hand side is a volume integral over the domain Ω , while the right-hand side represents a surface integral over the boundary $\partial\Omega$.

This results in the following weak formulation:

$$\int_{\Omega} \nabla u \cdot \nabla v \, dx - \int_{\partial\Omega} \frac{\partial u}{\partial n} v \, ds = \int_{\Omega} f v \, dx, \quad (4.36)$$

where the term $\frac{\partial u}{\partial n} = \nabla u \cdot \mathbf{n}$ represents the normal derivative of u on the boundary $\partial\Omega$.

We consider the boundary $\partial\Omega$ to be partitioned into two disjoint parts:

$$\partial\Omega = \Gamma_D \cup \Gamma_N, \quad \Gamma_D \cap \Gamma_N = \emptyset,$$

where Γ_D is the portion where Dirichlet boundary conditions are prescribed, and Γ_N is the portion with Neumann boundary conditions. Before applying the Dirichlet condition strongly, the boundary integral over Γ_D remains part of the weak form. It reflects a general formulation suitable for alternative treatments (e.g., Nitsche's method¹). The resulting weak form is:

$$\int_{\Omega} \nabla u \cdot \nabla v \, dx - \int_{\Gamma_D} \frac{\partial u}{\partial n} v \, ds = \int_{\Omega} f v \, dx + \int_{\Gamma_N} g v \, ds \quad (4.37)$$

Assuming the Dirichlet condition $u = u_D$ is imposed strongly (i.e., incorporated directly into the trial function space), the test function v vanishes on Γ_D . The weak form then simplifies to:

$$\int_{\Omega} \nabla u \cdot \nabla v \, dx = \int_{\Omega} f v \, dx + \int_{\Gamma_N} g v \, ds, \quad (4.38)$$

where g denotes the prescribed Neumann boundary portion, such that $\frac{\partial u}{\partial n} = g$ on Γ_N .

This weak formulation is commonly expressed in an abstract variational form:

$$a(u, v) = L(v) \quad \forall v \in V, \quad (4.39)$$

where the bilinear form $a(u, v)$ and the linear form $L(v)$ are defined as:

$$a(u, v) = \int_{\Omega} \nabla u \cdot \nabla v \, dx, \quad (4.40)$$

$$L(v) = \int_{\Omega} f v \, dx + \int_{\Gamma_N} g v \, ds. \quad (4.41)$$

We assume that $f \in L^2(\Omega)$, $g \in L^2(\Gamma_N)$, and $u_D \in H^{1/2}(\Gamma_D)$ to ensure well-posedness of the weak formulation.

The test function and the solution are assumed to belong to *Hilbert* spaces. In the Galerkin method, employed for the results presented in this work, both are chosen from the same *Hilbert* space.

The space V of test functions is a subspace of the *Sobolev* space $H^1(\Omega)$, defined as:

$$V = \{v \in H^1(\Omega) \mid v = 0 \text{ on } \Gamma_D\}, \quad (4.42)$$

and the trial space is given by:

$$U = \{u \in H^1(\Omega) \mid u = u_D \text{ on } \Gamma_D\}. \quad (4.43)$$

¹Nitsche's method enforces Dirichlet boundary conditions weakly by modifying the bilinear form. For details, see [141].

More explicitly, the *Sobolev* space $H^1(\Omega)$ is defined as:

$$H^1(\Omega) := \{v \in L^2(\Omega) \mid \nabla v \in [L^2(\Omega)]^d\}, \quad (4.44)$$

where d is the number of spatial dimensions. The *Lebesgue* space $L^2(\Omega)$ is defined as:

$$L^2(\Omega) = \left\{ v : \Omega \rightarrow \mathbb{R} \mid \int_{\Omega} v^2 dx < \infty \right\}. \quad (4.45)$$

This formulation is referred to as *weak* because it only requires equality in an integral sense, whereas the strong form demands pointwise equality and more regularity.

In practical computations, both the solution and test functions are approximated in finite-dimensional subspaces $U_h \subset U$ and $V_h \subset V$, which typically consist of globally continuous, piecewise affine (i.e., linear on each mesh element) basis functions defined over the computational mesh.

With all components defined, the final variational problem is expressed as follows:

$$\text{Find } u \in U \text{ such that } a(u, v) = L(v) \quad \forall v \in V. \quad (4.46)$$

4.2.2 Finite Element Discretization

To solve the weak form numerically, the domain Ω is subdivided into a finite number of smaller subdomains forming a triangulation $\mathcal{T}_h$. These are typically triangles or quadrilaterals in two dimensions, tetrahedra or hexahedra are common in three dimensions. The set of all mesh nodes is denoted by $\mathcal{N} = \{\mathbf{x}_i\}$, where $\mathbf{x}_i$ are the coordinates of node i .

Based on this triangulation, we define the finite element space V_h , a finite-dimensional subspace of V , consisting of globally continuous, piecewise affine functions:

$$V_h := \{v \in C^0(\Omega) \mid v|_T \text{ is affine } \forall T \in \mathcal{T}_h\}. \quad (4.47)$$

Each function $v_h \in V_h$ is uniquely determined by its nodal values. The set of nodes $\mathcal{N}$ is partitioned into two disjoint subsets: $\mathcal{N}_D$, consisting of nodes on the Dirichlet boundary Γ_D , and $\mathcal{N}_f$, the set of remaining (free) nodes. A nodal basis $\{\phi_i\}$ is chosen such that:

$$\phi_i(\mathbf{x}_j) = \delta_{ij}, \quad (4.48)$$

where δ_{ij} is the Kronecker delta. The space V_h is then given by:

$$V_h = \text{span}\{\phi_1, \phi_2, \dots, \phi_N\}. \quad (4.49)$$

The finite element approximation $u_h \in V_h$ is expressed as a linear combination of basis functions:

$$u_h(x) = \sum_{i=1}^N U_i \phi_i(x), \quad (4.50)$$

here, U_i are the nodal coefficients (degrees of freedom) corresponding to basis functions ϕ_i , and they form the components of the global solution vector $\mathbf{U} = (U_1, U_2, \dots, U_N)^\top$.

Substituting the basis expansion u_h and $v_h = \phi_j$ into the weak form (4.38) and requiring it to hold for all basis functions leads to the component-wise system:

$$\sum_i \left(\int_{\Omega} \nabla \phi_i \cdot \nabla \phi_j \, dx \right) U_i = \int_{\Omega} f \phi_j \, dx + \int_{\Gamma_N} g \phi_j \, ds. \quad (4.51)$$

This can be written in matrix-vector form as:

$$A\mathbf{U} = \mathbf{b}, \quad (4.52)$$

where the matrix and vector entries are defined by:

$$A_{ji} = \int_{\Omega} \nabla \phi_i \cdot \nabla \phi_j \, dx, \quad (4.53)$$

$$b_j = \int_{\Omega} f \phi_j \, dx + \int_{\Gamma_N} g \phi_j \, ds. \quad (4.54)$$

To enforce Dirichlet boundary conditions, described as:

$$U_i = u_D(\mathbf{x}_i), \quad \forall \mathbf{x}_i \in \mathcal{N}_D. \quad (4.55)$$

The system is then partitioned into known (Dirichlet) and unknown (free) degrees of freedom (DOF). The free DOFs correspond to the nodes in $\mathcal{N}_f$, for which the solution values $\mathbf{u}_f$ are not prescribed and must be computed by solving the system. In contrast, the Dirichlet DOFs $\mathbf{u}_D$, associated with nodes in $\mathcal{N}_D$, are imposed directly according to the boundary conditions. Letting $\mathbf{U} = (\mathbf{u}_D, \mathbf{u}_f)^\top$, and partitioning A and $\mathbf{b}$ accordingly, written as:

$$A = \begin{pmatrix} A_{DD} & A_{Df} \\ A_{fD} & A_{ff} \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} b_D \\ b_f \end{pmatrix}. \quad (4.56)$$

Here:

- (i) A_{DD} corresponds to interactions between Dirichlet nodes,
- (ii) A_{Df} (and its transpose A_{fD}) captures the coupling between Dirichlet and free nodes,
- (iii) A_{ff} includes contributions only from the free nodes and defines the part of the system to be solved numerically.

To incorporate the known Dirichlet values into the system, we substitute them directly into the global equation, which yields:

$$\begin{pmatrix} I & 0 \\ A_{fD} & A_{ff} \end{pmatrix} \begin{pmatrix} \mathbf{u}_D \\ \mathbf{u}_f \end{pmatrix} = \begin{pmatrix} \mathbf{u}_D \\ \mathbf{b}_f \end{pmatrix}, \quad (4.57)$$

where I is the identity matrix used to enforce Dirichlet values.




Figure 4.3: Representation of the original solution u and the finite element approximation u_h in one (a) and two (b) spatial dimensions. The basis functions for all nodes are shown at the bottom of the 1D plot. The highlighted basis function $\varphi_2(x)$ and corresponding solution value u_2 at node x_2 are shown in orange. The 2D plot shows the smooth function $u(x, y)$, its piecewise approximation $u_h(x, y)$, and the basis function $\varphi_{5,6}(x, y)$ centered at node (x_5, y_6) .

The final reduced symmetric system for the unknown values $\mathbf{u}_f$ is then:

$$A_{ff}\mathbf{u}_f = \mathbf{b}_f - A_{fD}\mathbf{u}_D. \quad (4.58)$$

This reduced system involves only the unknown free DOFs and forms the core of the numerical solution process. This system can be solved efficiently because of the sparsity of A , which arises from the local support of the basis functions. Moreover, the framework naturally accommodates mixed boundary conditions. It can be extended using higher-order basis functions and adaptive mesh refinement to improve accuracy in regions where the solution varies rapidly.


Figure 4.3 illustrates the representations of the original solution and its finite element approximation in both the 1D and 2D settings. In the 1D case (Figure 4.3 (a)), the nodes are aligned along a single axis, and piecewise linear basis functions approximate the solution. In the 2D case (Figure 4.3 (b)), the domain is discretized into triangles, and the hat function associated with a single node is shown along with a smooth solution surface and projection. The label $\varphi_{5,6}(x_5, y_6)$ denotes the basis function and nodal value associated with the corresponding node. This construction naturally extends to higher-order approximations by increasing the polynomial degree of the basis functions and adding more nodes per element. It also generalizes to 3D problems, where the domain is typically discretized using tetrahedral elements and basis functions are defined over a reference tetrahedron in a similar manner [142].

4.3 Finite Element Micromagnetic Solver

Figure 4.4b illustrates the general workflow of the micromagnetic simulation software *ViennaMag* [143], a finite element (FE)-based tool for the simulation of MRAM devices. Written in C++, the software is built on the open-source FE library *MFEM* [144, 145], which provides essential components for finite element simulations, including mesh handling, solvers, data structures, and result visualization. On top of this foundation, *ViennaMag* adds features specific to micromagnetics, such as the computation of the demagnetizing field, spin-accumulation, IEC, and the solution of the LLG equation.

The simulation begins by loading input parameters, including the computational mesh that describes the device geometry. Data structures are initialized to represent the magnetization in the ferromagnetic layers, and a method for computing the demagnetizing field is selected. Once setup is complete, the main simulation loop is executed.

In each iteration, the demagnetizing field is computed from the current magnetization. The spin-accumulation is calculated in two steps: first, the electric potential



(a) Flowchart of the micromagnetic FDM simulation loop. The LLG solver performs a 4-stage RK integration over each time step Δt to compute the updated magnetization.

(b) Flowchart of the *ViennaMag* FEM simulator. The micromagnetic state is updated iteratively by coupling spin and charge transport with time integration of the LLG equation.

Figure 4.4: Flowcharts of the micromagnetic simulation implementations.

is solved for to determine the charge-current distribution, this result is then used to solve the spin-accumulation equations. These quantities serve as inputs to the LLG equation, which is assembled and solved to update the magnetization. The system evolves over discrete time steps until the final simulation time.

The LLG equation, when coupled with the spin and charge drift-diffusion formalism, enables the modeling of magnetization dynamics in multilayer structures of arbitrary complexity. Since analytical solutions to this system of PDEs are limited to idealized cases, numerical approaches are essential. Section 4.1 presents an FDM-based implementation of the LLG equation, suitable only for thin FL layers with simple elliptical geometries. In contrast, the FEM naturally supports complex geometries and multiple material domains [146, 147], making it the preferred method for implementing a unified solver that models charge transport, spin-accumulation, and magnetization dynamics. This functionality is realized using the C++ library *MFEM*. The weak formulations and discretization strategies for these coupled systems are presented in Section 4.2.

4.3.1 Efficient Algorithms for Demagnetizing Field Computation

As discussed in Section 3.2.4, the demagnetizing field potential satisfies Poisson's equation with open boundary conditions, requiring the problem to be solved over the entire space $\mathbb{R}^3$. This presents a significant computational challenge, as standard FEM methods are restricted to finite domains and cannot directly handle open-boundary problems. In practice, calculating the demagnetizing field is the most computationally demanding task in micromagnetic simulations because it must account for long-range dipole interactions among magnetic moments.

Various numerical approaches have been developed to address this challenge. The *truncation approach* is a common technique in which a sufficiently large finite region replaces the infinite exterior domain. Alternatively, the *hybrid FE-BEM* allows accurate treatment of open boundaries by combining FEM in the magnetized region with boundary element method (BEM) in the exterior domain, significantly reducing computational effort while maintaining accuracy.

In the following, we present and compare these methods for computing the demagnetizing field, focusing on their accuracy, efficiency, and their ability to account for the long-range nature of magnetostatic interactions.

4.3.1.1 Truncation Approach

The most straightforward method for approximating the open boundary condition is the *truncation approach*. Here, the magnetized region Ω_m is embedded in a larger, finite, non-magnetic external region Ω_e , such that the total computational domain is $\Omega = \Omega_m \cup \Omega_e$. Since the magnetic scalar potential decays with distance outside the magnetized region, the open boundary conditions can be approximated by applying homogeneous Dirichlet boundary conditions (i.e., $u = 0$) on the outer surface $\partial\Omega_e$, thereby mimicking the decay to zero at infinity.

In this setting, the demagnetizing field is computed as:

$$\mathbf{H}_{\text{dem}} = -\nabla u, \quad (4.59)$$

where u is the scalar magnetic potential, obtained by solving the Poisson equation derived from Equation (3.43):

$$\nabla \cdot (\nabla u - \mathbf{M}_s) = 0 \quad \text{in } \Omega, \quad (4.60)$$

subject to the boundary condition:

$$u = 0 \quad \text{on } \partial\Omega_e. \quad (4.61)$$

Multiplying the equation by a test function $v \in V_0$ and integrating over the computational domain $\Omega = \Omega_m \cup \Omega_e$, we obtain the weak form:

$$\int_{\Omega} (\nabla u - \mathbf{M}_s) \cdot \nabla v \, dx = 0 \quad \forall v \in V_0. \quad (4.62)$$

Applying integration by parts to the first term (and using the fact that $v = 0$ on $\partial\Omega_e$) yields:

$$\int_{\Omega} \nabla u \cdot \nabla v \, dx = \int_{\Omega} \mathbf{M}_s \cdot \nabla v \, dx \quad \forall v \in V_0. \quad (4.63)$$

To accurately account for the discontinuity of the magnetization $\mathbf{m}$ at the boundary $\partial\Omega_m$, the right-hand side of Equation (4.63) is typically restricted to integration over the magnetic region Ω_m :

$$\int_{\Omega} \nabla u \cdot \nabla v \, dx = \int_{\Omega_m} \mathbf{M}_s \cdot \nabla v \, dx \quad \forall v \in V_0. \quad (4.64)$$

This restriction is significant when the magnetization $\mathbf{m}$ is discretized using piecewise-affine, globally continuous finite-element functions, as it enables a consistent treatment of the discontinuity at $\partial\Omega_m$. The function space V_0 is defined as:

$$V_0 = \{v \in V : v(x) = 0 \, \forall x \in \partial\Omega_e\}. \quad (4.65)$$

Imposing a zero potential at a finite boundary is inherently approximate, since the true potential vanishes only in the limit $r \rightarrow \infty$. Thus, the size of the external domain Ω_e significantly influences the accuracy of the solution. In practice, choosing Ω_e to be approximately *five times* larger than Ω_m in each spatial direction offers a good compromise between computational efficiency and accuracy [148] (see Figure 4.5). To further reduce computational effort without sacrificing precision, the mesh in Ω_e can be progressively coarsened toward the boundary $\partial\Omega_e$.

4.3.1.2 Hybrid Finite Element–Boundary Element Method

To reduce the number of DOFs in the computation of the demagnetizing field, it is advantageous to restrict spatial discretization to the magnetic domain Ω_m .




Figure 4.5: Finite-element mesh used for computing the demagnetization field with the truncation approach. (a) The magnetic region is shown in blue, while the surrounding external region is shown in light blue. The external domain is chosen to be approximately five times larger than the magnetic region in each spatial direction. (b) The magnetic region represents the area of interest and is embedded within the larger external domain used to model open boundary conditions.

This is achieved by a hybrid FE-BEM, initially proposed by Fredkin and Koehler [149]. Within this hybrid approach, the Poisson equation is solved in the magnetic domain using FEM, while the behavior of the potential at infinity is guaranteed by a double-layer potential placed on the boundary of the domain. Similar to the truncation method, the FE-BEM approach uses the scalar magnetic potential u as an intermediate quantity, from which the demagnetizing field is obtained via:

$$\mathbf{H}_{\text{dem}} = -\nabla u. \quad (4.66)$$

The potential is split as:

$$u = u_1 + u_2, \quad (4.67)$$

where u_1 is defined as the solution of Poisson's equation inside the magnetic domain:

$$\nabla^2 u_1(\mathbf{r}) = \nabla \cdot \mathbf{M} \quad \text{for } \mathbf{r} \in \Omega_m, \quad (4.68)$$

$$u_1 = 0 \quad \text{for } \mathbf{r} \notin \Omega_m, \quad (4.69)$$

with the Neumann boundary condition on $\partial\Omega_m$:

$$\frac{\partial u_1}{\partial n} = \mathbf{M} \cdot \mathbf{n}. \quad (4.70)$$

While u_1 satisfies the correct normal jump condition at the boundary, it does not ensure continuity of the potential. Therefore, a correction term u_2 is introduced such that:

$$\nabla^2 u_2 = 0 \quad \text{in } \mathbb{R}^3, \quad (4.71)$$

$$[\nabla u_2 \cdot \mathbf{n}] = 0 \quad \text{on } \partial\Omega_m, \quad (4.72)$$

$$[u_2] = u_1 \quad \text{on } \partial\Omega_m, \quad (4.73)$$

where $[f] = f_{\text{in}} - f_{\text{out}}$ denotes the jump of f across the boundary.

A double-layer potential of the form fulfils these conditions:

$$u_2(\mathbf{r}) = \frac{1}{4\pi} \oint_{\partial\Omega_m} \frac{\partial u_1(\mathbf{r}')}{\partial n'} \cdot \frac{1}{|\mathbf{r} - \mathbf{r}'|} dS'. \quad (4.74)$$

Although u_2 could theoretically be evaluated in the entire domain via this integral, doing so is computationally expensive due to its $\mathcal{O}(n^2)$ complexity. Instead, the integral is only evaluated on the boundary $\partial\Omega_m$ using the BEM, yielding the boundary values of u_2 through a matrix-vector product:

$$\mathbf{u}_2|_{\partial\Omega_m} = \mathbf{B} \mathbf{u}_1|_{\partial\Omega_m}. \quad (4.75)$$

These boundary values then serve as Dirichlet boundary conditions for solving the Laplace equation $\nabla^2 u_2 = 0$ inside Ω_m via the FEM, which provides u_2 throughout the magnetic domain. The total potential $u = u_1 + u_2$ is then used to compute the demagnetizing field.

The resulting matrix $\mathbf{B}$ is dense and scales quadratically with the number of boundary nodes n_B . Hierarchical matrix compression techniques [150, 151] are often employed to mitigate the associated memory and computational costs, reducing complexity to $\mathcal{O}(n_B \log n_B)$.

This hybrid approach significantly reduces the computational domain by discretizing only the magnetic region Ω_m . It remains effective even for disconnected or complex magnetic geometries. The trade-off, however, is the use of dense BEM operators, which differ from the sparse matrices typically encountered in FEM.

4.3.1.3 Verification and Benchmarking

The Micromagnetic Modeling Activity Group² at the NIST Center for Theoretical and Computational Materials Science has defined a set of standard problems to enable consistent comparison between micromagnetic simulation codes and published results. Following these recommendations, we assess the correctness and accuracy of our demagnetizing-field solvers using magnetization states from the μMAG Standard Problem 3 [152] together with geometries that admit analytical solutions.

The benchmark consists of three test cases:

- (i) **Single-domain tests:** A unit cube with three magnetization states (uniform, flower, vortex) to isolate solver accuracy for single-domain ferromagnets.
- (ii) **External-field test:** A two-disc configuration to evaluate the field outside the magnetized body.
- (iii) **Device application:** A realistic STT-MRAM cell with disconnected magnetic domains.

²<https://www.ctcms.nist.gov/~rdm/mumag.org.html>




Figure 4.6: Comparison of demagnetizing energy for three distinct magnetization states: (a) uniform, (b) flower, and (c) vortex. For each case, the top subfigures show the corresponding magnetic potential, and the bottom subfigures show the magnetization distribution.

For a uniformly magnetized cube, the demagnetizing field and thus the demagnetizing energy can be obtained in closed form. The energy is given by:

$$\mathcal{E}_d = -\frac{\mu_0 M_S}{2} \int_{\Omega} \mathbf{m}(\mathbf{r}) \cdot \mathbf{H}_d(\mathbf{r}) \, d\mathbf{r}. \quad (4.76)$$

The scalar magnetic potential reaches its maximum near the cube surfaces, where the divergence of $\mathbf{M}$ is largest.

The flower and vortex initializations follow [153]. The flower state is defined as:

$$m_x(\mathbf{r}) = \frac{1}{a} xz, \quad (4.77)$$

$$m_y(\mathbf{r}) = \frac{1}{c} yz + \frac{1}{b^3} y^3 z^3, \quad (4.78)$$

$$m_z(\mathbf{r}) = 1, \quad (4.79)$$

with parameters $a = c = 1$ and $b = 2$, and the cube centered at $(0, 0, 0)$.




Figure 4.7: Comparison of demagnetizing energy for unit cubes with different magnetization configurations (uniform, flower, vortex), calculated using the truncation approach and the hybrid FE-BEM approach against analytical reference values [153]. The hybrid FE-BEM shows excellent agreement with analytical results, while the truncation approach significantly underestimates the energy. Based on the figure originally published in [154].

The vortex state is defined as:

$$m_x(\mathbf{r}) = -\frac{y}{r} \sqrt{1 - \exp\left(-\frac{4r^2}{r_C^2}\right)}, \quad (4.80)$$

$$m_y(\mathbf{r}) = \frac{x}{r} \sqrt{1 - \exp\left(-\frac{4r^2}{r_C^2}\right)}, \quad (4.81)$$

$$m_z(\mathbf{r}) = \exp\left(-\frac{2r^2}{r_C^2}\right), \quad (4.82)$$

where $r = \sqrt{x^2 + y^2}$ and $r_C = 0.14$ as defined in [153].

Figure 4.6 illustrates the three magnetization configurations, uniform, flower, and vortex, together with their corresponding magnetic scalar potentials. Figure 4.7 presents the demagnetizing energies computed using both the truncation and hybrid FE-BEM approaches, compared against the analytical reference values from [153]. For quantitative comparison, we define the relative error as:

$$\text{err}_{\mathcal{E}} [\%] = 100 \frac{|\mathcal{E}_d^{\text{ref}} - \mathcal{E}_d^{\text{num}}|}{\max(\mathcal{E}_d^{\text{ref}}) + |\min(\mathcal{E}_d^{\text{ref}})|}. \quad (4.83)$$

The hybrid FE-BEM agrees within $< 1\%$ for the uniform and flower states and approximately 5% for the vortex case at identical mesh resolution, whereas the truncation approach consistently exhibits errors greater than 10% . These results confirm that, for a similar mesh in the magnetic domain, the hybrid FE-BEM approach is superior to the truncation approach in both accuracy and computational cost.

To evaluate the demagnetizing field *outside* the magnetized body, we consider a two-disc configuration with known analytical solution [155]. The on-axis field of a




Figure 4.8: Two-disc configuration with parallel magnetizations. Left: demagnetizing field $H_{d,x}$ along the symmetry axis. Right: relative error with respect to the analytical reference (4.84). The FE-BEM maintains accuracy across the axis, while the truncation approach shows significant deviations.




Figure 4.9: Two-disc configuration with anti-parallel magnetizations. Left: demagnetizing field $H_{d,x}$ along the symmetry axis. Right: relative error with respect to the analytical reference. The FE-BEM reproduces the analytical field within a few percent, while the truncation approach yields larger errors.

single uniformly magnetized disc of radius R and thickness L , magnetized along its symmetry axis, is given by:

$$H_{d,x}^{\text{disc}}(x) = \frac{M_S}{2} \left(\frac{x}{\sqrt{x^2 + R^2}} - \frac{x - L}{\sqrt{(x - L)^2 + R^2}} \right), \quad (4.84)$$

which follows from the surface-charge representation $\sigma_m = \mathbf{M} \cdot \mathbf{n}$.

The analytical reference for two identical discs is assembled pointwise from Equation (4.84) by mirroring about the center plane. For axial sampling points $\{x_j\}_{j=0}^{N-1}$:

$$H_{d,x}^{\text{P}}[j] = H_1[j] + H_1[N - 1 - j], \quad (4.85)$$

$$H_{d,x}^{\text{AP}}[j] = H_1[j] - H_1[N - 1 - j], \quad (4.86)$$

where $H_1[j] \equiv H_{d,x}^{\text{disc}}(x_j)$, and P/AP denote parallel and anti-parallel magnetizations, respectively.

To quantify accuracy along the axis, we report the pointwise relative error:

$$\text{err}_H[j] [\%] = 100 \frac{|H_{d,x}^{\text{num}}(x_j) - H_{d,x}^{\text{ref}}(x_j)|}{\max_k |H_{d,x}^{\text{ref}}(x_k)|}, \quad (4.87)$$

which avoids divisions by small numbers near zero crossings (relevant in the anti-parallel case).

For the numerical setup we use discs of thickness $L = 0.1$ nm and radius $R = 1$ nm. To obtain field values between the discs with the hybrid FE-BEM (which solves only on the magnetic domain), we insert a thin probe disc between them with negligible saturation magnetization so as not to perturb the solution. Figure 4.8 and Figure 4.9 show the field profiles and relative errors for both configurations. In both cases, the FE-BEM solution tracks the analytical reference within a few percent across the axis. At the same time, the truncation approach shows larger deviations of up to 10–15%, especially near the disc faces, where gradients are strongest.

Having established accuracy on canonical test cases, we evaluate the demagnetizing field in a realistic STT-MRAM cell. The device has a diameter of 40 nm and consists of two non-magnetic contacts of length 50 nm, a ferromagnetic RL of thickness 1 nm, a tunnel barrier of thickness 1 nm, and a free ferromagnetic layer of thickness 1.7 nm.

The magnetic domain is *disconnected*, consisting of two ferromagnetic layers separated by the tunnel barrier. This geometry provides a thorough test for the solver. Figure 4.10 and Figure 4.11 compare the hybrid FE-BEM and truncation solutions for the P and AP magnetization configurations, respectively. For visualization, the contacts and barrier are hidden, arrows indicate the direction of the demagnetizing field, and the color map encodes the magnetic potential.

Both methods produce very similar fields and potentials, confirming that the FE-BEM implementation remains accurate for disconnected magnetic domains. Given its superior accuracy at comparable mesh resolutions and reduced computational effort relative to large truncated domains, the hybrid FE-BEM is adopted for the complex geometries considered in subsequent chapters, including synthetic antiferromagnets.




Figure 4.10: Demagnetizing field in an STT-MRAM cell with P magnetization configuration. Arrows indicate the demagnetization field direction, while color-coding represents the magnetic scalar potential u . Left: hybrid FE-BEM. Right: truncation approach. Both methods yield similar results, validating the implementation for disconnected magnetic domains. Figure originally published in [154].




Figure 4.11: Demagnetizing field in an STT-MRAM cell with AP magnetization configuration. Arrows indicate the demagnetization field direction, while color-coding represents the magnetic scalar potential u . Left: hybrid FE-BEM. Right: truncation approach. Both methods yield similar results, validating the implementation for disconnected magnetic domains. Figure originally published in [154].

Numerical Time Integration of the LLG Equation

The temporal evolution of magnetization in micromagnetic systems is governed by the LLG equation introduced in Chapter 3. This nonlinear partial differential equation couples precessional dynamics with phenomenological damping. After spatial discretization by finite-element or finite-difference methods (Chapter 4), the LLG equation reduces to a large system of coupled ordinary differential equations (ODEs). The efficient and accurate numerical integration of this system is essential for practical micromagnetic simulations, yet poses several fundamental challenges that distinguish it from generic ODE problems.

This chapter develops the theoretical framework for time integrating the spatially discretized LLG equation and presents a systematic comparison of the integration schemes implemented in this work. Section 5.1 identifies the key numerical challenges as stiffness, the unit-sphere constraint, and geometric structure, that guide the choice of time integrator. Section 5.2 reviews the classical taxonomy of explicit, implicit, and implicit-explicit (IMEX) methods. Section 5.3 introduces the tangent-plane scheme, a reformulation of the LLG equation that transforms the constrained nonlinear problem into a sequence of linear saddle-point systems. Section 5.4 discusses adaptive multistep methods provided by the *SUNDIALS* library. Finally, Section 5.5 presents verification studies using the μ MAG Standard Problem 4 and compares the accuracy of the implemented integrators.

5.1 Challenges in LLG Time Integration

Numerical integration of the LLG equation presents three interrelated challenges that must be addressed by any practical time-stepping scheme: stiffness arising from disparate time scales, preservation of the unit-sphere constraint, and retention of the underlying geometric structure.

5.1.1 Stiffness and Disparate Time Scales

A differential equation is called *stiff* when its solution contains components evolving on vastly different time scales, and when explicit methods require impractically small

time steps to maintain numerical stability rather than accuracy [156]. The LLG equation is inherently stiff due to the exchange interaction.

The exchange field contribution to the effective field takes the form:

$$\mathbf{H}_{\text{exch}} = \frac{2A}{\mu_0 M_s} \nabla^2 \mathbf{m}, \quad (5.1)$$

where A is the exchange stiffness constant. Upon spatial discretization with mesh size h , the Laplacian operator introduces eigenvalues scaling as h^{-2} , leading to an effective exchange field magnitude:

$$|\mathbf{H}_{\text{exch}}| \sim \frac{A}{\mu_0 M_s h^2}. \quad (5.2)$$

For typical ferromagnetic materials ($A \approx 10^{-11}$ J/m, $M_s \approx 10^6$ A/m) and fine meshes ($h \approx 1$ nm), this yields exchange fields of order 10^4 – 10^5 A/m, corresponding to precessional frequencies:

$$\omega_{\text{exch}} = \gamma_0 |\mathbf{H}_{\text{exch}}| \sim 10^{12}$$
– 10^{13} rad/s. (5.3)

The associated time scale $\tau_{\text{exch}} = 2\pi/\omega_{\text{exch}} \sim 10^{-13}$ – 10^{-12} s is orders of magnitude shorter than the physical time scales of interest in device applications, such as magnetization switching times of order 10^{-9} s.

For explicit time-stepping schemes, the Courant–Friedrichs–Lewy (CFL) stability condition requires:

$$\Delta t < \frac{C}{\omega_{\text{exch}}} \sim \frac{C \mu_0 M_s h^2}{\gamma_0 A}, \quad (5.4)$$

where C is an order-unity constant depending on the specific method. This constraint becomes increasingly restrictive as the mesh is refined, rendering explicit schemes computationally impractical for high-resolution simulations [157, 158, 159].

While implicit time-stepping schemes bypass this strict CFL stability limit, allowing significantly larger time steps, stiffness re-emerges as a bottleneck in linear algebra. The large eigenvalues of the exchange operator result in a system matrix $\mathbf{A}$ with a larger condition number $\kappa_2(\mathbf{A})$, defined in terms of the spectral matrix norm $\|\cdot\|_2$ as:

$$\kappa_2(\mathbf{A}) = \|\mathbf{A}\|_2 \cdot \|\mathbf{A}^{-1}\|_2 \approx \frac{|\lambda_{\text{max}}|}{|\lambda_{\text{min}}|}. \quad (5.5)$$

Since the maximum eigenvalue of the discretized exchange operator scales with the inverse square of the mesh size ($\lambda_{\text{max}} \propto h^{-2}$), refining the mesh drastically increases $\kappa_2(\mathbf{A})$. This ill-conditioning causes standard iterative solvers (such as Generalized Minimal Residual (GMRES)) to stagnate, necessitating the use of effective preconditioners.

Figure 5.1 illustrates the separation of time scales in a typical micromagnetic simulation. The exchange-driven dynamics impose a stability limit that is far below the time scales associated with damping, demagnetization, and variations in the external field.




Figure 5.1: Hierarchy of time scales in micromagnetic simulations. The exchange interaction imposes the CFL stability limit ($\tau_{\text{exch}} \sim 10^{-12}$ s for $h = 2$ nm mesh). Demagnetization dynamics occur at $\tau_{\text{demag}} \sim 10^{-11}$ s. Precession due to external fields operates at $\tau_{\text{prec}} \sim 10^{-9}$ s. Damping and switching occur at $10^{-9} - 10^{-8}$ s, far above the explicit stability limit. Material: Permalloy ($A = 1.3 \times 10^{-11}$ J/m, $M_s = 8 \times 10^5$ A/m, $\alpha = 0.02$).

5.1.2 The Unit-Sphere Constraint

A fundamental property of the LLG equation is that the magnetization magnitude is conserved:

$$|\mathbf{m}(\mathbf{x}, t)| = 1 \quad \forall \mathbf{x} \in \Omega, t \geq 0. \quad (5.6)$$

This constraint follows directly from the structure of the LLG equation. Taking the inner product of Equation (3.9) with $\mathbf{m}$ and using the orthogonality of cross products yields:

$$\mathbf{m} \cdot \frac{\partial \mathbf{m}}{\partial t} = \frac{1}{2} \frac{\partial}{\partial t} |\mathbf{m}|^2 = 0. \quad (5.7)$$

This confirms that $|\mathbf{m}|^2$ is a conserved quantity.

Numerical time-stepping schemes do not automatically preserve this constraint. Three strategies are commonly employed:

- (i) **Projection methods:** After each time step, the magnetization is renormalized:

$$\mathbf{m}^{n+1} \leftarrow \frac{\mathbf{m}^{n+1}}{|\mathbf{m}^{n+1}|}. \quad (5.8)$$

This approach is simple but introduces $\mathcal{O}(\Delta t^p)$ perturbations, where p is the order of the underlying scheme.

- (ii) **Constraint-preserving schemes:** Certain schemes, such as the implicit mid-point rule applied to the LLG equation in Gilbert form, preserve the constraint exactly at the discrete level [94].

- (iii) **Tangent-plane formulation:** The equation is reformulated to solve for the velocity $\mathbf{v} = \partial_t \mathbf{m}$ in the tangent space $T_{\mathbf{m}} S^2 = \{\mathbf{v} : \mathbf{v} \cdot \mathbf{m} = 0\}$, enforcing orthogonality either exactly via Lagrange multipliers or approximately via projection [135, 160].

The tangent-plane approach, developed in detail in Section 5.3, is employed throughout this work.

5.1.3 Structural Properties

The LLG equation possesses a rich geometric structure that certain numerical schemes can exploit or preserve. In the absence of damping ($\alpha = 0$), the equation reduces to:

$$\frac{\partial \mathbf{m}}{\partial t} = -\gamma_0 \mathbf{m} \times \mathbf{H}_{\text{eff}}, \quad (5.9)$$

which is a Hamiltonian system with the total micromagnetic energy as the Hamiltonian. The magnetization precesses on surfaces of constant energy, and the dynamics preserve both the energy and the phase-space volume (Liouville's theorem).

When damping is present ($\alpha > 0$), the system becomes dissipative. As derived in Chapter 3 (cf. Equation (3.15)), the energy decreases monotonically [94]:

$$\frac{d\mathcal{E}}{dt} = -\frac{\alpha\gamma_0}{1+\alpha^2} \int_{\Omega} \|\mathbf{m} \times \mathbf{H}_{\text{eff}}\|_2^2 d\mathbf{x} \leq 0. \quad (5.10)$$

A necessary property of numerical schemes is to preserve this dissipative character at the discrete level, ensuring that artificial energy growth does not occur.

5.2 Classification of Time Integrators

Time integration schemes for ODEs are commonly classified as explicit or implicit, and as one-step or multistep methods. This section reviews the classes most relevant for integrating the LLG equation.

5.2.1 Explicit Runge–Kutta Methods

The LLG equation can be written as a general initial value problem $\dot{\mathbf{m}} = f(t, \mathbf{m})$, where f denotes the right-hand side (RHS) of Equation (3.12) including all effective field contributions and, where applicable, spin-transfer torque terms. Explicit Runge–Kutta methods advance this system from t_n to $t_{n+1} = t_n + \Delta t$ via:

$$\mathbf{m}_{n+1} = \mathbf{m}_n + \Delta t \sum_{i=1}^s b_i \mathbf{k}_i, \quad (5.11)$$

where the s stages are computed sequentially:

$$\mathbf{k}_i = f \left(t_n + c_i \Delta t, \mathbf{m}_n + \Delta t \sum_{j=1}^{i-1} a_{ij} \mathbf{k}_j \right), \quad i = 1, \dots, s. \quad (5.12)$$

The coefficients a_{ij} , b_i , and c_i are conveniently represented in a Butcher tableau [156]:

$$\begin{array}{c|c} \mathbf{c} & A \\ \hline & \mathbf{b}^\top \end{array} \quad (5.13)$$

where $A = (a_{ij}) \in \mathbb{R}^{s \times s}$ is the Runge–Kutta matrix, $\mathbf{b} = (b_1, \dots, b_s)^\top$ the quadrature weights, and $\mathbf{c} = (c_1, \dots, c_s)^\top$ the stage nodes at which the right-hand side is evaluated. For an explicit method, A is strictly lower triangular.

Classic examples include the forward Euler method ($s = 1$, first-order), Heun’s method ($s = 2$, second-order), and the classical fourth-order Runge–Kutta (RK4) method ($s = 4$). For stiff problems, explicit methods are subject to severe stability restrictions. The stability region of the forward Euler method, for instance, is the unit disk centered at $(-1, 0)$ in the complex plane. Eigenvalues of the linearized LLG operator corresponding to exchange modes lie far to the left of this region, necessitating time steps satisfying Equation (5.4) (see Figure 5.2). The classical RK4 method has an extended stability region, extending to approximately $|\lambda \Delta t| < 2.8$ along the negative real axis. For the exchange-dominated eigenvalues $|\lambda_{\max}| \sim 10^{13} \text{ rad s}^{-1}$ typical of fine micromagnetic meshes, stability requires $\Delta t \lesssim 10^{-13} \text{ s}$. While this constraint necessitates many time steps (e.g., 20000 steps for a 2 ns simulation), RK4 remains a viable option when the computational cost per step is low and fourth-order accuracy in constraint preservation is desired.

5.2.2 Implicit Methods and A-Stability

Implicit methods evaluate the RHS at the unknown solution value $\mathbf{m}_{n+1}$, thereby enlarging the stability region. The simplest example is the backward Euler method:

$$\mathbf{m}_{n+1} = \mathbf{m}_n + \Delta t f(t_{n+1}, \mathbf{m}_{n+1}). \quad (5.14)$$

This method is *A-stable*: its stability region includes the entire left half of the complex plane, meaning that all stable modes of the continuous problem remain stable regardless of the time-step size [156].

The implicit midpoint rule is defined as:

$$\mathbf{m}_{n+1} = \mathbf{m}_n + \Delta t f \left(t_{n+1/2}, \frac{\mathbf{m}_n + \mathbf{m}_{n+1}}{2} \right). \quad (5.15)$$

It is also A-stable and additionally preserves the unit-sphere constraint for the LLG equation when combined with appropriate spatial discretization [90].

The computational cost of implicit methods arises from solving the nonlinear system that occurs at each time step. For the LLG equation, Newton iteration requires the Jacobian $\partial f / \partial \mathbf{m}$, which is dense due to the demagnetization field. However, the tangent-plane reformulation (Section 5.3) transforms the problem into a *linear* system for each time step, dramatically reducing the computational overhead.




Figure 5.2: Stability regions in the complex plane $z = \lambda\Delta t$ with physical eigenvalues. The *Physics of Interest* (Precession, Damping) lies near the origin and could be stable for explicit methods. The exchange eigenvalue position depends on the time step: for $\Delta t = 10^{-13}$ s, $\lambda_{\text{exch}}\Delta t \approx -1$ lies within the RK4 stability region (black dot). Increasing Δt moves it left (gray dot), eventually exiting explicit stability regions. Explicit solvers (Red/Blue) are forced to take tiny time steps ($\Delta t \rightarrow 0$) just to keep the black dot inside their stability boundary. The Implicit solver (Green) remains stable for any Δt , allowing us to step over the stiff exchange modes. IMEX methods combine both approaches: exchange modes are treated implicitly (Green), while precession and damping remain explicit (Red/Blue).

5.2.3 Multistep Methods: BDF Schemes

Backward differentiation formulas (BDFs) are implicit linear multistep methods that approximate the time derivative using values from multiple previous time steps. The k -step BDF method takes the form:

$$\sum_{j=0}^k \alpha_j \mathbf{m}_{n+1-j} = \Delta t f(t_{n+1}, \mathbf{m}_{n+1}), \quad (5.16)$$

where the coefficients α_j are chosen to maximize the order of accuracy. For example, BDF2 reads:

$$\frac{3}{2} \mathbf{m}_{n+1} - 2 \mathbf{m}_n + \frac{1}{2} \mathbf{m}_{n-1} = \Delta t f(t_{n+1}, \mathbf{m}_{n+1}). \quad (5.17)$$

This achieves second-order accuracy. BDF methods up to order 6 exist, but only BDF1 and BDF2 are A-stable. The higher-order methods BDF3–BDF5 are $A(\alpha)$ -stable, meaning their stability region contains a wedge $|\arg(-\lambda)| < \alpha$ for some angle $\alpha < \pi/2$

rather than the full left half-plane [156]. This weaker condition, also referred to as stiff stability, is sufficient for problems where all stiff eigenvalues lie near the negative real axis, as is the case for exchange-dominated micromagnetic systems. BDF6 is not even $A(\alpha)$ -stable and is therefore not used in practice.

The C-language Variable-coefficient Ordinary Differential Equation solver (CVODE) solver from the *SUNDIALS* library [161] implements variable-order, variable-step BDF methods with automatic step-size and order selection based on local error estimates. This adaptive capability is particularly valuable for micromagnetic simulations, where the dynamics may alternate between rapidly varying transients and slowly evolving quasi-static phases.

5.2.4 Implicit-Explicit Method

When the RHS naturally splits into stiff and non-stiff components, it can be written as:

$$\frac{d\mathbf{m}}{dt} = f_E(t, \mathbf{m}) + f_I(t, \mathbf{m}). \quad (5.18)$$

IMEX methods treat f_I implicitly for stability while evaluating f_E explicitly for efficiency [162]. For the LLG equation, we assign the stiff exchange-damping contribution to the implicit partition:

$$f_I = -\gamma_L \alpha \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{exch}}), \quad (5.19)$$

$$f_E = -\gamma_L \mathbf{m} \times \mathbf{H}_{\text{eff}} - \gamma_L \alpha \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}}^{(\text{non-exch})}), \quad (5.20)$$

where $\mathbf{H}_{\text{eff}}^{(\text{non-exch})} = \mathbf{H}_{\text{eff}} - \mathbf{H}_{\text{exch}}$. Crucially, only the *damping* component of exchange enters f_I , while the precessional term $-\gamma_L \mathbf{m} \times \mathbf{H}_{\text{exch}}$ remains in f_E . This choice ensures that the implicit Jacobian $\partial f_I / \partial \mathbf{m}$ is symmetric positive-definite, enabling efficient algebraic multigrid (AMG) or preconditioned conjugate gradient (PCG) solvers (cf. Section 5.4.2).

The ARKStep solver from *SUNDIALS* implements additive Runge–Kutta (ARK) methods for IMEX problems [161]. This approach is investigated in Section 5.5.

5.3 The Tangent-Plane Scheme

The tangent-plane scheme (TPS), introduced by Alouges [135] and further developed by Abert *et al.* [90] and Praetorius *et al.* [160], reformulates the LLG equation to exploit its geometric structure. Rather than solving directly for $\mathbf{m}_{n+1}$, the scheme solves for the velocity $\mathbf{v} = \partial_t \mathbf{m}$ constrained to lie in the tangent plane at the current magnetization. This reformulation offers two key advantages: it transforms the nonlinear problem into a linear one and enforces the unit-sphere constraint to $\mathcal{O}(\Delta t^2)$ via the tangent-plane projection, with exact restoration achieved through post-step renormalization.

5.3.1 Reformulation of the LLG Equation

The Gilbert form of the LLG Equation (3.10) is:

$$\frac{\partial \mathbf{m}}{\partial t} = -\gamma_0 \mathbf{m} \times \mathbf{H}_{\text{eff}} + \alpha \mathbf{m} \times \frac{\partial \mathbf{m}}{\partial t}, \quad (5.21)$$

where $\mathbf{H}_{\text{eff}}$ is the effective field defined in Equation (3.19), can be rewritten by taking the cross product with $\mathbf{m}$ and using the vector triple product identity. With $\mathbf{v} = \partial_t \mathbf{m}$, one obtains:

$$\alpha \mathbf{v} + \mathbf{m} \times \mathbf{v} = \gamma_0 (\mathbf{H}_{\text{eff}} - (\mathbf{m} \cdot \mathbf{H}_{\text{eff}}) \mathbf{m}). \quad (5.22)$$

The RHS is the projection of $\mathbf{H}_{\text{eff}}$ onto the tangent plane at $\mathbf{m}$. Crucially, if $\mathbf{v}$ satisfies the constraint $\mathbf{v} \cdot \mathbf{m} = 0$, then the RHS automatically lies in the tangent plane, and the equation is consistent.

5.3.2 Weak Formulation

The finite-element discretization proceeds by multiplying Equation (5.22) by a test function $\mathbf{w}$ and integrating over the domain Ω . In weak form:

$$\int_{\Omega} (\alpha \mathbf{v} + \mathbf{m} \times \mathbf{v}) \cdot \mathbf{w} \, d\mathbf{x} = \gamma_0 \int_{\Omega} \mathbf{H}_{\text{eff}} \cdot \mathbf{w} \, d\mathbf{x}, \quad (5.23)$$

where the term $(\mathbf{m} \cdot \mathbf{H}_{\text{eff}}) \mathbf{m} \cdot \mathbf{w}$ vanishes when $\mathbf{w}$ lies in the tangent space.

The effective field contributions are treated as follows. For the exchange field, integration by parts yields:

$$\int_{\Omega} \mathbf{H}_{\text{exch}} \cdot \mathbf{w} \, d\mathbf{x} = -\frac{2A}{\mu_0 M_s} \int_{\Omega} \nabla \mathbf{m} : \nabla \mathbf{w} \, d\mathbf{x}, \quad (5.24)$$

where $\nabla \mathbf{m} : \nabla \mathbf{w} = \sum_{i,j} \partial_j m_i \partial_j w_i$ denotes the Frobenius inner product, and homogeneous Neumann boundary conditions have been assumed. Other field contributions (external, anisotropy, demagnetization) enter as volume integrals against $\mathbf{w}$.

5.3.3 Complete Weak Form with Spin-Transfer Torque

For simulations of current-induced magnetization dynamics, spin-transfer torque contributions must be incorporated. The spin-accumulation $\mathbf{S}$, computed from the drift-diffusion equations, enters via the torque expression derived in Section 6.1. The complete weak form reads:

$$\begin{aligned} \int_{\Omega_m} (\alpha \mathbf{v} + \mathbf{m}^n \times \mathbf{v}) \cdot \mathbf{w} \, d\mathbf{x} = & \frac{2A\gamma_0}{\mu_0 M_s} \int_{\Omega_m} \nabla \mathbf{m}^n : \nabla \mathbf{w} \, d\mathbf{x} \\ & - \gamma_0 \int_{\Omega_m} (\mathbf{H}_{\text{ext}} + \mathbf{H}_{\text{ani}}(\mathbf{m}^n) + \mathbf{H}_{\text{dem}}(\mathbf{m}^n)) \cdot \mathbf{w} \, d\mathbf{x} \\ & + \frac{D_e}{M_s} \int_{\Omega_m} \left(\frac{\mathbf{S}^n}{\lambda_J^2} + \frac{\mathbf{m}^n \times \mathbf{S}^n}{\lambda_{\varphi}^2} \right) \cdot \mathbf{w} \, d\mathbf{x}, \end{aligned} \quad (5.25)$$

where D_e is the electron diffusion coefficient, and λ_J , λ_{φ} are the exchange and dephasing lengths characterizing the spin-transfer torque. The integrals are evaluated only over the magnetic subdomain Ω_m .

5.3.4 Time Discretization: The θ -Method

A θ -method time discretization treats the effective field at a weighted average of the current and next time levels:

$$\mathbf{H}_{\text{eff}}^\theta = (1 - \theta)\mathbf{H}_{\text{eff}}(\mathbf{m}^n) + \theta\mathbf{H}_{\text{eff}}(\mathbf{m}^{n+1}). \quad (5.26)$$

For the exchange field, which depends linearly on $\mathbf{m}$, this introduces implicit dependence on $\mathbf{m}^{n+1} = \mathbf{m}^n + \Delta t \mathbf{v}$. The key observation is that:

$$\nabla \mathbf{m}^{n+1} = \nabla \mathbf{m}^n + \Delta t \nabla \mathbf{v}. \quad (5.27)$$

Defining the θ -weighted magnetization as $\mathbf{m}^\theta = (1 - \theta)\mathbf{m}^n + \theta\mathbf{m}^{n+1} = \mathbf{m}^n + \theta\Delta t \mathbf{v}$, the exchange contribution becomes:

$$-\frac{2A\gamma_0}{\mu_0 M_s} \int_{\Omega} \nabla \mathbf{m}^\theta : \nabla \mathbf{w} \, d\mathbf{x} = -\frac{2A\gamma_0}{\mu_0 M_s} \int_{\Omega} \nabla \mathbf{m}^n : \nabla \mathbf{w} \, d\mathbf{x} - \theta\Delta t \frac{2A\gamma_0}{\mu_0 M_s} \int_{\Omega} \nabla \mathbf{v} : \nabla \mathbf{w} \, d\mathbf{x}. \quad (5.28)$$

In the full weak formulation (Equation (5.25)), the first term is known from $\mathbf{m}^n$ and moves to the right-hand side, while the second term depends on the unknown $\mathbf{v}$ and contributes to the left-hand side.

5.3.5 The Saddle-Point System

Enforcement of the tangent-plane constraint $\mathbf{v} \cdot \mathbf{m} = 0$ via Lagrange multipliers leads to a block-structured linear system. Let $\mathbf{v}_h$ and λ_h denote the finite-element approximations to the velocity and Lagrange multiplier, expanded in the basis $\{\phi_i\}$ defined in Section 4.2.2:

$$\mathbf{v}_h = \sum_i v_i \phi_i, \quad \lambda_h = \sum_i \lambda_i \phi_i, \quad (5.29)$$

where ϕ_i are scalar basis functions (with ϕ_i the corresponding vector-valued functions). The constraint is imposed weakly:

$$\int_{\Omega} \mathbf{v}_h \cdot \mathbf{m}^n \psi \, d\mathbf{x} = 0 \quad \forall \psi \in V_h, \quad (5.30)$$

where V_h is the scalar finite-element space.

Assembling the discrete equations yields the block system:

$$\begin{pmatrix} A & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} \mathbf{v} \\ \boldsymbol{\lambda} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{0} \end{pmatrix}, \quad (5.31)$$

where:

- (i) A is the $(1, 1)$ -block with entries $A_{ij} = \int_{\Omega} (\alpha \phi_j + \mathbf{m}^n \times \phi_j) \cdot \phi_i \, d\mathbf{x} + \theta \Delta t K_{ij}$,
- (ii) $K_{ij} = \frac{2A\gamma_0}{\mu_0 M_s} \int_{\Omega} \nabla \phi_j : \nabla \phi_i \, d\mathbf{x}$ is the exchange stiffness matrix,

- (iii) B is the constraint matrix with entries $B_{ij} = \int_{\Omega} (\mathbf{m}^n \cdot \phi_j) \phi_i d\mathbf{x}$,
- (iv) $\mathbf{f}$ is the RHS incorporating explicit field contributions.

The zero block in the $(2, 2)$ position renders the overall system indefinite, giving rise to a saddle-point problem [163]: the solution corresponds to a minimum with respect to $\mathbf{v}$ and a maximum with respect to $\boldsymbol{\lambda}$.

The $(1, 1)$ -block A has the structure:

$$A = \alpha M + [\mathbf{m}^n]_{\times} + \theta \Delta t K, \quad (5.32)$$

where M is the mass matrix, $[\mathbf{m}^n]_{\times}$ is the skew-symmetric matrix representing the cross product with $\mathbf{m}^n$, and K is the stiffness matrix (see Figure 5.3). In addition to the indefiniteness of the overall saddle-point structure, the skew-symmetric contribution renders A itself non-symmetric, necessitating GMRES [164] or other Krylov methods for iterative solution.

The linear algebra routines are provided by the *MFEM* library [144, 145].

5.3.6 Constraint Enforcement and Magnetization Update

The tangent-plane constraint $\mathbf{v} \cdot \mathbf{m}^n = 0$ is enforced via Lagrange multipliers by solving the full saddle-point system (5.31). This yields a velocity $\mathbf{v}$ that satisfies the constraint in the finite-element sense, ensuring that the magnetization update remains tangent to the unit sphere.

After solving for the velocity $\mathbf{v}$, the magnetization is updated:

$$\mathbf{m}^{n+1} = \mathbf{m}^n + \Delta t \mathbf{v}. \quad (5.33)$$

Even with exact tangent-plane enforcement, this first-order update does not preserve $|\mathbf{m}| = 1$ exactly and therefore the deviation is $\mathcal{O}(\Delta t^2)$ (as shown in Section 5.1.2).

Renormalization is therefore applied:

$$\mathbf{m}^{n+1} \leftarrow \frac{\mathbf{m}^{n+1}}{|\mathbf{m}^{n+1}|}. \quad (5.34)$$

5.4 Adaptive Multistep Methods via *SUNDIALS*

The *SUNDIALS* (SUite of Nonlinear and Differential/ALgebraic equation Solvers) library [161] provides implementations of adaptive time-stepping methods. Two solvers from this suite are integrated into the present framework: CVODE for implicit multistep methods and ARKStep for IMEX Runge–Kutta methods.

Both solvers operate on the ODE form $d\mathbf{m}/dt = f(\mathbf{m})$. The FEM discretization developed in Section 4.2.2 transforms the LLG Equation (3.12) into the semi-discrete system (discrete in space, continuous in time):

$$M \frac{d\mathbf{m}}{dt} = F(\mathbf{m}), \quad (5.35)$$




Figure 5.3: Block structure of the tangent-plane scheme saddle-point system corresponding to (5.31). The (1,1)-block A combines the damping, precession, and implicit exchange terms. The off-diagonal blocks B and B^T enforce the tangent-plane constraint $\mathbf{v} \cdot \mathbf{m} = 0$, while the solution vector contains both the velocity $\mathbf{v}$ and Lagrange multipliers $\boldsymbol{\lambda}$.

where M is the mass matrix with entries $M_{ij} = \int_{\Omega} \phi_i \phi_j d\mathbf{x}$, and $F(\mathbf{m})$ is the weak-form assembly of the LLG RHS:

$$F_i(\mathbf{m}) = -\gamma_L \int_{\Omega} [\mathbf{m} \times \mathbf{H}_{\text{eff}} + \alpha \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}})] \cdot \phi_i d\mathbf{x}, \quad (5.36)$$

with $\mathbf{H}_{\text{eff}}$ as defined in (3.19). To obtain the explicit ODE form required by *SUNDIALS*, we employ row-sum mass lumping, replacing M with a diagonal matrix M^L having entries $M_{ii}^L = \sum_j M_{ij}$. This approximation preserves the unit-length constraint at each node and has been proven to converge to a weak solution of the LLG equation [165, 166]. This yields the nodewise system:

$$\frac{d\mathbf{m}}{dt} = f(\mathbf{m}) = (M^L)^{-1} F(\mathbf{m}), \quad (5.37)$$

enabling efficient explicit evaluation and nodewise constraint enforcement $|\mathbf{m}_i| = 1$ at each mesh node [90].

5.4.1 CVODE: Variable-Order BDF

CVODE implements variable-order, variable-step BDF (cf. Section 5.2.3) with automatic error control. At each step, the solver performs a predictor-corrector sequence:

- (i) **Prediction:** An initial guess $\mathbf{m}_{n+1}^{(0)}$ is constructed via polynomial extrapolation from the history array.
- (ii) **Correction:** The predicted solution is refined by solving the nonlinear residual equation $G(\mathbf{m}_{n+1}) = 0$ using Newton iteration. The residual is defined as:

$$G(\mathbf{m}) = \sum_{j=0}^q \alpha_j \mathbf{m}_{n+1-j} - \Delta t \beta_0 f(\mathbf{m}), \quad (5.38)$$

with coefficients α_j, β_0 as defined in (5.16).

- (iii) **Adaptation:** The local truncation error is estimated to dynamically adjust both the step size Δt and the method order q (typically $1 \leq q \leq 5$). To ensure unconditional stability for the purely imaginary eigenvalues characteristic of the exchange interaction, the maximum BDF order is restricted to $q = 2$, ensuring A-stability (similar to the strategy in [158]). This order restriction is essential for the oscillatory dynamics of the LLG equation. The precessional term generates eigenvalues $\lambda \approx \pm i\omega$ lying close to the imaginary axis, where $\omega = \gamma_0 |\mathbf{H}_{\text{eff}}|$ is the Larmor precession frequency. While BDF1 and BDF2 are A-stable (their stability regions include the entire left half-plane), BDF3 and higher orders exclude wedge-shaped regions near the imaginary axis [156]. When these precessional eigenvalues fall outside the stability region, the adaptive solver is forced to reduce the time step to maintain stability rather than accuracy, thereby negating the efficiency gains of implicit integration. Restricting to order $q \leq 2$ guarantees unconditional stability for all eigenvalue configurations encountered in micromagnetic simulations.

The core computational challenge lies in the Newton correction. Let $\mathbf{m}^{(k)}$ denote the k -th iteration. The update $\delta \mathbf{m} = \mathbf{m}^{(k+1)} - \mathbf{m}^{(k)}$ is obtained by solving the linear system:

$$(I - \gamma J) \delta \mathbf{m} = -G(\mathbf{m}^{(k)}), \quad (5.39)$$

where $\gamma = \Delta t \beta_0$ and $J = \partial f / \partial \mathbf{m}$ is the system Jacobian. For the discretized LLG equation, J is a dense, non-symmetric matrix of dimension $3N \times 3N$. Direct factorization of $I - \gamma J$ is computationally prohibitive ($\mathcal{O}(N^3)$). Instead, CVODE employs preconditioned Krylov subspace methods (GMRES), which approximate the solution using only matrix-vector products.

The system Jacobian $J = \partial f / \partial \mathbf{m}$ inherits contributions from each term in the effective field (3.19):

$$J = J_{\text{exch}} + J_{\text{aniso}} + J_{\text{demag}} + J_{\text{STT}} + \dots \quad (5.40)$$

The exchange contribution J_{exch} , arising from the Laplacian in $\mathbf{H}_{\text{exch}}$ (3.27), dominates the spectral radius with eigenvalues scaling as $\mathcal{O}(h^{-2})$. The demagnetization Jacobian J_{demag} is dense due to the nonlocal nature of $\mathbf{H}_{\text{demag}}$, while J_{aniso} is sparse and bounded. This structure motivates a preconditioner that captures the stiff exchange modes while neglecting the expensive demagnetization contribution.

Preconditioning Strategy: The spectral distribution of the iteration matrix governs the convergence rate of GMRES. Without preconditioning, the stiffness of the exchange interaction leads to eigenvalues scaling as $\mathcal{O}(h^{-2})$ distributed along both the negative real axis (damping) and the imaginary axis (precession). This results in a significant condition number, causing the linear solver to stagnate.

Suess *et al.* [158] proposed a physics-based preconditioner within a finite-element framework employing the *box scheme* [167] for spatial discretization on tetrahedral meshes. One starts from the exchange contribution to f :

$$f_{\text{exch}} = -\gamma_L \mathbf{m} \times \nabla^2 \mathbf{m} - \gamma_L \alpha \mathbf{m} \times (\mathbf{m} \times \nabla^2 \mathbf{m}), \quad (5.41)$$

where γ_L denotes the gyromagnetic ratio. Linearizing around the current magnetization $\mathbf{m}$ and using $|\mathbf{m}| = 1$ yields:

$$\frac{\partial f_{\text{exch}}}{\partial \mathbf{m}} \approx -\gamma_L [\mathbf{m}]_{\times} \nabla^2 + \gamma_L \alpha \nabla^2, \quad (5.42)$$

where $[\mathbf{m}]_{\times}$ denotes the skew-symmetric matrix representing the cross product with $\mathbf{m}$:

$$[\mathbf{m}]_{\times} = \begin{pmatrix} 0 & -m_z & m_y \\ m_z & 0 & -m_x \\ -m_y & m_x & 0 \end{pmatrix}, \quad \text{such that } [\mathbf{m}]_{\times} \mathbf{v} = \mathbf{m} \times \mathbf{v}. \quad (5.43)$$

The original Suess preconditioner takes the form:

$$P_{\text{Suess}} = M + \gamma \alpha C_{\text{exch}} K_{\text{diff}} + \gamma C_{\text{exch}} K_{\text{skew}}, \quad (5.44)$$

where M is the mass matrix, $C_{\text{exch}} = 2\gamma_L A_{\text{ex}}/(\mu_0 M_s)$ is the exchange coefficient, and:

- (i) K_{diff} : Symmetric positive-definite matrix arising from the damping term, mathematically equivalent to a vector diffusion operator.
- (ii) K_{skew} : Skew-symmetric matrix arising from the precessional term $-\mathbf{m} \times \nabla^2 \mathbf{m}$, capturing the rotation of the magnetization.

In the *box scheme*, the effective field at each node i is computed as $H_{i,\text{eff}}^k \approx -m_i^{-1} \partial E_t / \partial u_i^k$, where $m_i = \int_{V_i} J_s(\mathbf{x}) dV$ is the magnetic moment associated with the control volume V_i surrounding node i . This yields a naturally diagonal weighting with entries m_i , enabling the second derivatives of the total energy to be computed *analytically* on an element-by-element basis, from which the system Jacobian can be constructed explicitly [158]. The stray field contribution is omitted from the preconditioner to preserve sparsity. The resulting preconditioner system is solved using Biconjugate Gradient Stabilized (BiCGSTAB) method [168] with relaxed incomplete LU (RILU) factorization [169].

Adaptation from Box Scheme to Galerkin Finite Elements: In contrast to the *box scheme* employed by Suess *et al.*, the present work uses a standard Galerkin finite-element discretization with row-sum mass lumping (cf. Section 5.5.4). This change in spatial discretization requires adapting the preconditioner construction. The semi-discrete system (5.35) involves the consistent mass matrix M , which is sparse but non-diagonal. Following standard practice [90, 165], we employ row-sum mass lumping to obtain a diagonal approximation M^L with entries $M_{ii}^L = \sum_j M_{ij}$. The preconditioner is then constructed in the dimensionless form appropriate for the lumped system:

$$P = I + \gamma (M^L)^{-1} (\alpha C_{\text{exch}} K_{\text{diff}} + C_{\text{exch}} K_{\text{skew}}). \quad (5.45)$$

In both the *box scheme* and the lumped Galerkin formulation, the diagonal entries (m_i or M_{ii}^L , respectively) vary across nodes on non-uniform meshes. However, in the *box scheme*, the m_i^{-1} weighting is intrinsic to the effective field computation, and

the Jacobian is assembled directly in this weighted form. In the Galerkin approach, the weighting arises as an *a posteriori* left-multiplication $(M^L)^{-1}K_{\text{skew}}$ that destroys the skew-symmetric structure of the assembled stiffness matrix. Specifically, letting $D = (M^L)^{-1}$ denote the non-uniform diagonal scaling, the product DK_{skew} is no longer skew-symmetric:

$$(DK_{\text{skew}})^T = K_{\text{skew}}^T D = -K_{\text{skew}} D \neq -DK_{\text{skew}}, \quad (5.46)$$

unless D is a scalar multiple of the identity. The resulting matrix is neither symmetric nor skew-symmetric, with eigenvalues scattered in the complex plane. This causes BiCGSTAB to stagnate or diverge, particularly at larger time steps, where the skew contribution dominates [170]. For this reason, the full Suess preconditioner with the skew-symmetric term is not used in this work.

Practical Implementation: Instead, we adopt a simplified preconditioner that omits the problematic skew-symmetric term entirely. The original Suess preconditioner [158] was developed using the *box scheme* for spatial discretization, in which the m_i^{-1} weighting is incorporated during Jacobian assembly rather than applied as a post-hoc scaling. In the present Galerkin formulation, the analogous scaling $(M^L)^{-1}$ is applied after assembly of the global stiffness matrices, which disrupts the skew-symmetric structure of the precessional contribution. Standard Galerkin finite-element discretization with mass lumping produces a diagonal matrix M^L whose entries vary significantly on unstructured meshes. Following the same approach as Suess *et al.* [158], who validated the original preconditioner against the μMAG standard problems, the simplified preconditioner adopted in this work reads:

$$P = I + \gamma\alpha C_{\text{exch}}(M^L)^{-1}K_{\text{diff}}. \quad (5.47)$$

This matrix is symmetric positive-definite, enabling the use of the PCG method with AMG preconditioning via BoomerAMG from the *HYPRE* library [171]. The AMG solver is configured for 3D vector problems using the systems interface with nodal coarsening. A loose relative tolerance of 10^{-1} is employed, as the preconditioner needs only to provide an approximate solution as the outer GMRES iteration in CVODE corrects any residual error.

While this preconditioner does not capture the precessional dynamics, it effectively clusters the stiff exchange eigenvalues, reducing the GMRES iteration count from $\mathcal{O}(100)$ (unpreconditioned) to $\mathcal{O}(20)$ per Newton step. The trade-off, approximately 30–40% more linear iterations compared to the full Suess preconditioner, is acceptable given the guaranteed convergence and robustness on meshes of arbitrary quality. The correctness of this approach is validated against the μMAG Standard Problem 4 in Section 5.5.2, where all solvers using this preconditioner reproduce the published reference solutions. Table 5.1 summarizes the preconditioner variants and their applicability.

Alternative: Incomplete LU Factorization. An alternative approach employs incomplete LU (ILU) factorization, which directly factorizes the non-symmetric preconditioner matrix without requiring structural modifications. Unlike AMG, which relies

¹Non-symmetric due to $J_{\text{FL}} = c_{\text{FL}}[\mathbf{S}]_{\times}$ (skew-symmetric).

Table 5.1: Comparison of preconditioner strategies for the CVODE Newton system.

Preconditioner	Structure	Solver	Mesh
$P = M + \gamma(\alpha K + K_{\text{skew}})$	Indefinite	BiCGSTAB+RILU	Box scheme
$P = I + \gamma\alpha C_{\text{ex}}(M^L)^{-1}K_{\text{diff}}$	SPD	PCG+AMG	Any
$P = I + \gamma\alpha C_{\text{ex}}(M^L)^{-1}K_{\text{diff}} - \gamma(M^L)^{-1}J_{\text{STT}}$	Non-sym. ¹	GMRES+ILU	Any

on smoothing properties that degrade for skew-symmetric operators, ILU factorization handles indefinite systems naturally by computing approximate L and U factors. The Euclid parallel ILU implementation in *HYPRE* [171] provides a scalable option for moderate problem sizes. With ILU preconditioning, the full preconditioner including K_{skew} and J_{STT} can be employed:

$$P_{\text{ILU}} = I + \gamma\alpha C_{\text{exch}}(M^L)^{-1}K_{\text{diff}} + \gamma C_{\text{exch}}(M^L)^{-1}K_{\text{skew}} + \gamma(M^L)^{-1}J_{\text{STT}}. \quad (5.48)$$

This approach recovers the full Jacobian approximation at the cost of increased fill-in compared to AMG. For problems where STT dominates, the improved Jacobian approximation can reduce the number of outer GMRES iterations sufficiently to offset the higher preconditioner cost.

Extension for Spin-Transfer Torque: The standard preconditioner assumes the system is driven towards equilibrium by damping. However, current-induced switching introduces a non-conservative STT that acts as *negative damping*, pumping energy into the system and destabilizing the initial state.

In the switching regime, the STT terms dominate the dynamics. If the preconditioner ignores these terms, the linear solver will operate with an incorrect model of the system stiffness, leading to slow convergence. The total derivative of the STT with respect to magnetization is:

$$\frac{d\tau_{\text{STT}}}{d\mathbf{m}} = \left. \frac{\partial\tau_{\text{STT}}}{\partial\mathbf{m}} \right|_{\mathbf{S}} + \frac{\partial\tau_{\text{STT}}}{\partial\mathbf{S}} \cdot \frac{\partial\mathbf{S}}{\partial\mathbf{m}}. \quad (5.49)$$

The second term represents the feedback of magnetization changes on the spin-accumulation through the drift-diffusion system. Computing this term exactly would require differentiating through the solution of the spin transport PDE, effectively solving an adjoint problem for each magnetization degree of freedom, rendering it computationally prohibitive.

We therefore adopt the *frozen-S* approximation:

$$J_{\text{STT}} \approx \left. \frac{\partial\tau_{\text{STT}}}{\partial\mathbf{m}} \right|_{\mathbf{S}=\text{const}}. \quad (5.50)$$

Explicitly, for the damping-like torque $\boldsymbol{\tau}_{\text{DL}} = c_{\text{DL}} \mathbf{m} \times (\mathbf{m} \times \mathbf{S})$ and field-like torque $\boldsymbol{\tau}_{\text{FL}} = c_{\text{FL}} \mathbf{m} \times \mathbf{S}$:

$$\left. \frac{\partial \boldsymbol{\tau}_{\text{DL}}}{\partial \mathbf{m}} \right|_{\mathbf{S}} = -c_{\text{DL}} [\mathbf{m} \otimes \mathbf{S} + (\mathbf{m} \cdot \mathbf{S}) \mathbf{I}], \quad (5.51)$$

$$\left. \frac{\partial \boldsymbol{\tau}_{\text{FL}}}{\partial \mathbf{m}} \right|_{\mathbf{S}} = c_{\text{FL}} [\mathbf{S}]_{\times}. \quad (5.52)$$

Where $[\mathbf{S}]_{\times}$ denotes the skew-symmetric cross-product matrix of $\mathbf{S}$. This approximation neglects the feedback term $(\partial \boldsymbol{\tau} / \partial \mathbf{S})(\partial \mathbf{S} / \partial \mathbf{m})$ in the Jacobian. Physically, this relies on the same time-scale separation that motivates the steady-state assumption $\partial_t \mathbf{S} = 0$ in the drift-diffusion formulation (see Section 6.1): the spin-flip relaxation time τ_{sf} is of the order of picoseconds, whereas magnetization dynamics evolve on nanosecond time scales [35]. Consequently, $\mathbf{S}$ equilibrates quasi-instantaneously with respect to $\mathbf{m}$ and can be recomputed from $\mathbf{m}$ at the beginning of each time step and held constant during the Newton iteration that advances $\mathbf{m}$.

Both the damping-like torque contribution (5.51) and the field-like term (5.52) are incorporated into the preconditioner. While the damping-like term J_{DL} is symmetric, the field-like contribution $J_{\text{FL}} = c_{\text{FL}} [\mathbf{S}]_{\times}$ is *skew-symmetric*. This skew-symmetry introduces the same structural problem encountered with K_{skew} : AMG smoothers, designed for elliptic operators, degrade when applied to matrices with significant skew-symmetric components.

Consequently, when spin-transfer torque is active in the BDF solver, the preconditioner employs ILU factorization rather than AMG. Unlike algebraic multigrid, ILU directly factorizes the non-symmetric matrix without relying on smoothing properties, handling the skew-symmetric J_{FL} naturally. The Euclid parallel ILU implementation from *HYPRE* [171] is employed with fill level $k = 1$.

The extended preconditioner for BDF with spin-transfer torque reads:

$$P_{\text{BDF}} = \mathbf{I} + \gamma \alpha C_{\text{exch}} (M^L)^{-1} K_{\text{diff}} - \gamma (M^L)^{-1} J_{\text{STT}}, \quad (5.53)$$

where $J_{\text{STT}} = J_{\text{DL}} + J_{\text{FL}}$ comprises the damping-like (symmetric) and field-like (skew-symmetric) Jacobian contributions defined in Equations (5.51) and (5.52). Note that K_{skew} is *omitted* from the preconditioner despite appearing in the true Jacobian: numerical experiments showed that including K_{skew} after mass lumping degraded convergence rather than improving it (cf. the discussion following Equation (5.46)). The outer linear system is solved using GMRES preconditioned by ILU.

The numerical validation of this extended preconditioner for spin-transfer torque switching is presented in Section 6.8, following the derivation of the spin-accumulation model.

5.4.2 ARKStep: IMEX Runge–Kutta

ARKStep implements Additive Runge–Kutta (ARK) methods for the IMEX splitting introduced in Section 5.2.4. The effectiveness of IMEX strategies for micromagnetic

simulations, treating exchange implicitly while evaluating lower-order contributions explicitly, has been analyzed rigorously by Mauser *et al.* [166]. The formulation applies diagonally implicit tableaux to f_I and explicit tableaux to f_E simultaneously:

$$\mathbf{k}_i^E = f_E \left(t_n + c_i \Delta t, \mathbf{m}_n + \Delta t \sum_{j=1}^{i-1} a_{ij}^E \mathbf{k}_j^E + \Delta t \sum_{j=1}^{i-1} a_{ij}^I \mathbf{k}_j^I \right), \quad (5.54)$$

$$\mathbf{k}_i^I = f_I \left(t_n + c_i \Delta t, \mathbf{m}_n + \Delta t \sum_{j=1}^{i-1} a_{ij}^E \mathbf{k}_j^E + \Delta t \sum_{j=1}^i a_{ij}^I \mathbf{k}_j^I \right), \quad (5.55)$$

with the solution advanced via:

$$\mathbf{m}_{n+1} = \mathbf{m}_n + \Delta t \sum_{i=1}^s (b_i^E \mathbf{k}_i^E + b_i^I \mathbf{k}_i^I), \quad (5.56)$$

where b_i^E, b_i^I are the quadrature weights from the respective Butcher tableaux.

Implicit Preconditioning: The computational bottleneck in IMEX schemes is the solution of the nonlinear systems arising in the implicit stages (5.55). For the LLG equation, the implicit partition f_I (5.19) isolates the exchange damping term. Because the exchange field depends linearly on $\mathbf{m}$, the Jacobian of the implicit function in weak form is:

$$J_{\text{IMEX}} = M - \gamma \frac{\partial f_I}{\partial \mathbf{m}} = M + \gamma \alpha C_{\text{exch}} K_{\text{stiff}}, \quad (5.57)$$

where M is the mass matrix and K_{stiff} is the stiffness matrix associated with the vector Laplacian of the magnetization. Applying mass lumping ($M \rightarrow M^L$, $M_{ii}^L = \sum_j M_{ij}$) and left-multiplying by $(M^L)^{-1}$ yields the dimensionless form suitable for iterative solution:

$$(M^L)^{-1} J_{\text{IMEX}} \approx I + \gamma \alpha C_{\text{exch}} (M^L)^{-1} K_{\text{stiff}}. \quad (5.58)$$

Unlike the full Newton Jacobian in CVODE (5.53), this operator is symmetric positive-definite (SPD) when the precessional cross-product and spin-transfer torque are moved to f_E . This structural property enables highly efficient linear solvers such as PCG with AMG preconditioning, which are optimal for elliptic diffusion problems. The extension to spin-transfer torque simulations preserves this advantage: since $\mathbf{J}_{\text{STT}}$ does not appear in the implicit partition, it is omitted from the IMEX preconditioner, maintaining the SPD structure regardless of spin-torque magnitude.

However, a fundamental limitation arises from the explicit treatment of precession. The partition f_E contains the full precession term $-\gamma_L \mathbf{m} \times \mathbf{H}_{\text{eff}}$, which oscillates at the Larmor frequency $\omega = \gamma_0 |\mathbf{H}_{\text{eff}}| \sim 10^{10} - 10^{11} \text{ rad s}^{-1}$. While the IMEX splitting successfully removes the exchange-driven CFL stability restriction, the time step remains constrained by the *truncation error* of this explicit term. To accurately resolve the rapid oscillations, the local error must remain small, requiring $\Delta t \lesssim C/\omega \sim 0.1 - 1 \text{ ps}$, comparable to the explicit stability limit. This accuracy constraint, rather than stability, becomes the limiting factor.

Evidence for this behavior appears in the ARKStep solver statistics: throughout our Standard Problem 4 simulations, `num_exp_limited_stepfails` = 0 (no stability-induced rejections) while `num_acc_limited_steps` equals the total step count (all

Table 5.2: Summary of implemented time integrators.

Method	Type	Key Features
TPS	Implicit	Exact constraint via Lagrange multipliers
RK4	Explicit	Fourth-order, fixed step
CVODE BDF	Implicit	Variable-order, adaptive step, ILU preconditioner ²
ARKStep IMEX	IMEX	Exchange implicit, SPD Jacobian, adaptive step

steps accuracy-limited). The solver never encounters the explicit stability boundary, while the truncation error of the explicit precession term determines every time step.

A further consequence is oscillatory time-step behavior. As the magnetization precesses, the effective field magnitude $|\mathbf{H}_{\text{eff}}|$ varies periodically, causing the truncation error and hence the allowed time step to fluctuate, see Figure 5.7 (c). The adaptive controller continuously adjusts the step size, *chasing* the oscillating error estimate. This contrasts sharply with the smooth adaptation observed in fully implicit BDF methods, where the implicit treatment of precession filters the high frequency dynamics, allowing the error controller to respond to the slowly varying solution. Table 5.2 summarizes the time integrators implemented in this work.

5.5 Verification and Validation

Verification confirms that the equations are solved correctly (solving the equations right), while validation assesses whether the equations adequately represent the physical system (solving the right equations) [172]. This section presents verification studies using analytical solutions and the μMAG Standard Problem 4.

5.5.1 The μMAG Standard Problem 4

Standard Problem 4, defined by the Micromagnetic Modeling Activity Group (μMAG) [173], is a widely used benchmark for validating micromagnetic simulation codes. It considers a thin Permalloy film with dimensions $500 \times 125 \times 3 \text{ nm}^3$ and material parameters: Exchange stiffness $A = 1.3 \times 10^{-11} \text{ J m}^{-1}$, Saturation magnetization $M_s = 8.0 \times 10^5 \text{ A m}^{-1}$, Gilbert damping $\alpha = 0.02$, Anisotropy $K = 0$ (no crystalline anisotropy). The initial magnetization is an equilibrium *S-state*, prepared according to the μMAG specifications by applying and then slowly reducing a saturating field along the $(1, 1, 1)$ diagonal. To achieve this, a saturating external field of 1.0 T is applied along the $(1, 1, 1)$ direction and held constant for 0.5 ns to align the magnetization. This field is then linearly ramped down to zero over a period of 3 ns, after which the system is allowed to relax in zero field for an additional 2.5 ns. This procedure eliminates metastable states, yielding the characteristic "S"-shaped domain pattern required for the benchmark. Following the relaxation, a field of sufficient magnitude to reverse the

²Preconditioner adapts to physics: PCG+AMG when SPD (no STT), GMRES+ILU when non-symmetric (STT active), see Section 5.4.1.




Figure 5.4: Setup for the μ MAG Standard Problem 4. (a) Permalloy film geometry ($500 \times 125 \times 3 \text{ nm}^3$) discretized by an unstructured tetrahedral mesh with 45391 elements. (b) Initial *S*-state magnetization configuration showing the characteristic *S* pattern in the *xy*-plane, colored by the *x*-component of magnetization.

magnetization is applied. The simulation geometry, discretization, and equilibrium *S*-state magnetization configuration are shown in Figure 5.4. The field is applied instantaneously at $t = 0$ to the equilibrium *S*-state:

(i) **Field 1:** $\mathbf{H}_{\text{ext}} = (-24.6, 4.3, 0) \text{ mT}$.

The time evolution of the spatially averaged magnetization components, $\langle m_x \rangle$, $\langle m_y \rangle$, and $\langle m_z \rangle$, is then examined as the system approaches equilibrium in the new field. Reference solutions are available from multiple independent sources [173].

5.5.2 Solver Comparison on Standard Problem 4

Figure 5.5 compares the magnetization dynamics obtained with different solvers for Field 1 of Standard Problem 4. The results are validated against reference solutions from d'Aquino *et al.* [174] and Albuquerque *et al.* [175], both submitted to the μ MAG repository. All solvers produce trajectories in excellent agreement with these reference solutions, validating the correctness of the implementation. The *S*-state relaxed to an average magnetization of $\mathbf{m} = (m_x, m_y, m_z) = (0.96849366, 0.12573718, -0.00088421495)$. The slight deviation from FDM reference values ($m_x \approx 0.966$, $m_y \approx 0.127$, $m_z \approx 0.00000078$) is attributed to differences in




Figure 5.5: Magnetization dynamics for Standard Problem 4 (Field 1) computed with different time integrators. All solvers agree with the μ MAG reference solutions [174, 175].

spatial discretization: the FEM provides a more accurate representation of the magnetization configuration near edges and corners, whereas FDM solvers employ a staircase approximation that affects the local demagnetizing field calculation in these regions.

Table 5.3 provides quantitative comparison metrics. The maximum deviation from the reference solutions is computed as:

$$e_\infty = \max_{t \in [0, T]} \|\langle \mathbf{m} \rangle(t) - \langle \mathbf{m} \rangle_{\text{ref}}(t)\|_2, \quad (5.59)$$

where $\langle \mathbf{m} \rangle_{\text{ref}}$ denotes the reference solutions of d'Aquino *et al.* [174] and Albuquerque *et al.* [175], with the corresponding deviations shown in Figure 5.6.

The computational efficiency results reveal that both adaptive implicit methods significantly outperform the fixed-step explicit schemes. The BDF method achieves a $2.4\times$ speedup over the TPS baseline despite requiring approximately 25 linear solver iterations per Newton step. By taking time steps $10\text{--}50\times$ larger than the explicit CFL limit ($1\text{--}5$ ps versus 0.1 ps), the total step count reduces from 20000 to just 1099, more than compensating for the increased per-step cost.




Figure 5.6: Deviation of computed magnetization trajectories from the μ MAG reference solutions: (left) d'Aquino *et al.* [174], (right) Albuquerque *et al.* [175]. All solvers exhibit comparable deviation from both references, with differences attributable to the finite element discretization rather than time integration errors.

Table 5.3: Quantitative comparison of solvers on μ MAG Standard Problem 4. Errors are computed against both reference solutions.

Solver	Configuration	e_∞		Steps	CPU [s]	Speedup
		d'Aquino	Albuquerque			
TPS	$\Delta t = 10^{-13}$ s	0.10899498	0.12411681	20000	96752	1.00×
TPS	$\Delta t = 10^{-12}$ s	0.31432466	0.25772080	2000	11962	8.09×
RK4	$\Delta t = 10^{-13}$ s	0.09044942	0.12825486	20000	78813	1.23×
BDF	rtol= 10^{-2} , atol= 10^{-4}	0.11185953	0.16341087	1099	40014	2.42×
IMEX	rtol= 10^{-2} , atol= 10^{-4}	0.08920000	0.12800000	3818	34510	2.80×

The IMEX method achieves the best overall performance, completing the simulation in 34510s, which results in a $2.8\times$ speedup over TPS and $2.3\times$ over RK4. Although the time step remains accuracy-limited to 0.3–0.5ps due to explicit precession (cf. Section 5.4.2), the SPD structure of the implicit Jacobian enables highly efficient linear solves via AMG. This per-step efficiency advantage compensates for the larger step count (3818 versus 1099 for BDF), demonstrating that IMEX methods can be competitive when the implicit partition is carefully chosen to yield favorable algebraic structure.

To investigate whether the TPS could achieve comparable efficiency by simply increasing the time step, we tested the scheme with $\Delta t = 1$ ps, the same order of magnitude used by the adaptive BDF solver. While the simulation remains stable (the TPS is unconditionally A-stable for $\theta = 1$), the solution exhibits significant phase drift, with $e_\infty = 0.31$ against d'Aquino and $e_\infty = 0.26$ against Albuquerque, more than double the error of any other solver. This demonstrates the critical distinction between *stability* and *accuracy*: an A-stable method guarantees bounded solutions regardless

of time step, but does not guarantee accurate solutions. The first-order truncation error of the TPS update (5.33) accumulates as $\mathcal{O}(N\Delta t) = \mathcal{O}(T)$, where $N = T/\Delta t$ is the number of steps. Increasing Δt by a factor of 10 increases the per-step error by the same factor, resulting in the observed phase lag visible in Figure 5.5.

This finding underscores the value of adaptive time-stepping methods. The BDF solver achieves large time steps (1–5 ps) while maintaining accuracy because the embedded error estimator automatically reduces the step size when truncation error would exceed the specified tolerance. In contrast, a fixed-step method such as TPS requires the user to balance efficiency and accuracy manually, a balance that depends on the specific problem and is difficult to determine *a priori*.

These results demonstrate that both fully implicit (BDF) and IMEX methods achieve significant speedup over explicit schemes for precession-dominated problems. The IMEX splitting following Di Fratta *et al.* [176] and Mauser [166] successfully addresses exchange stiffness while maintaining SPD structure in the implicit partition, enabling efficient AMG solves that offset the accuracy-limited time step.

5.5.3 Constraint Preservation and Solver Accuracy

A key aspect of numerical time integration for the LLG equation is the preservation of the unit-sphere constraint $|\mathbf{m}| = 1$. Figure 5.7 (a) shows the constraint deviation $\| |\mathbf{m}| - 1 \|_\infty$ (maximum over all mesh nodes) before renormalization.

For explicit single-step methods, theoretical results establish that a scheme of order p preserves the magnetization magnitude within $\mathcal{O}(\Delta t^{p+1})$ error per step [177, 178]. This favorable behavior exploits the cross-product structure of the LLG equation, which guarantees $\partial_t \mathbf{m} \perp \mathbf{m}$. Our RK4 results confirm this: with $\Delta t = 10^{-13}$ s, the fourth-order method yields per-step drift of $\mathcal{O}(\Delta t^5) \sim 10^{-65}$, accumulating to $\sim 10^{-9}$ over the simulation. The TPS method, being first-order in its tangent-plane update, produces $\mathcal{O}(\Delta t^2) \sim 10^{-26}$ drift per step, accumulating to $\sim 10^{-5}$.

The implicit multistep methods (BDF and IMEX) exhibit fundamentally different drift characteristics. The $\mathcal{O}(\Delta t^{p+1})$ bound does not apply because the multistep formulation combines solution values from previous steps that each carry minor constraint violations, and the Newton iteration required for the implicit solve explores intermediate states that may lie far from the unit sphere. Unlike the RK4 method, the adaptive BDF and IMEX schemes do not preserve the constraint order-by-order. Consequently, they exhibit larger oscillating deviations that correlate with the step size (Figure 5.7 (c)). Crucially, however, this norm deviation acts primarily as a scaling artifact rather than a physical error. In the unconstrained formulation, the effective field scales linearly with magnetization magnitude ($|\mathbf{H}| \propto |\mathbf{m}|$). A deviation in norm (e.g., $|\mathbf{m}| > 1$) scales the update vector but preserves its direction. Since the projection step $\mathbf{m} \leftarrow \mathbf{m}/|\mathbf{m}|$ maps the solution back to the unit sphere radially, this *length error* is discarded without corrupting the *angle error* (physical precession). This geometric property explains why the BDF and IMEX solvers maintain high trajectory accuracy (Figure 5.7 (b)) even when the intermediate norm deviation is significant.




Figure 5.7: (a) Deviation from the unit-sphere constraint $|\mathbf{m}| = 1$ before renormalization. (b) Difference between solver trajectories and the TPS reference solution. (c) Time step evolution for adaptive solvers (BDF, IMEX) compared to the fixed time step used by TPS and RK4.

All solvers apply post-step renormalization to restore $|\mathbf{m}| = 1$, ensuring the physical constraint is satisfied regardless of the intermediate drift. The constraint deviation generally decreases as the simulation progresses, reflecting the reduction in magnetization velocity $|\partial_t \mathbf{m}|$ as the system relaxes toward equilibrium.

The decreasing constraint deviation for fixed-step methods reflects the diminishing magnetization velocity $|\partial_t \mathbf{m}| \rightarrow 0$ as the system relaxes, thereby reducing the per-step truncation error. In contrast, the adaptive solvers maintain approximately constant deviation by increasing the time step to consume their full error budget, as evidenced by the growing Δt in Figure 5.7 (c).

Figure 5.7 (b) quantifies the deviation of each solver from the TPS reference solution. All solvers remain within acceptable error bounds, with differences arising primarily from the time-discretization approach rather than from fundamental inconsistencies in the underlying physics. The deviation from the TPS solution confirms that all implemented solvers converge to the same physical trajectory within the specified tolerances. Interestingly, the RK4, BDF, and IMEX solvers exhibit nearly identical deviations from TPS ($e_\infty \approx 5 \times 10^{-2}$), while differing from each other by only $\mathcal{O}(10^{-6})$. This reflects the different underlying formulations: TPS solves for the velocity in the tangent plane with exact constraint enforcement, whereas the other methods integrate the explicit LLG equation in $\mathbb{R}^3$ and then renormalize. Both approaches produce physically correct results that match the μMAG reference solutions within spatial discretization error.

5.5.4 Adaptive Time Stepping

A key advantage of the *SUNDIALS*-based solvers is their automatic time-step adaptation. Figure 5.7 (c) shows the time step history during the Standard Problem 4 simulation, with the fixed time step of the TPS and RK4 solvers shown for reference.

The BDF solver achieves time steps of 1–5 ps during the post-switching relaxation phase, representing a 10–50 $\times$ improvement over the explicit CFL limit. During the rapid switching transient ($t \approx 0.1$ – 0.3 ns), the solver automatically reduces the time

step to capture the fast dynamics, then smoothly increases it as the system relaxes. This smooth adaptation profile reflects the fully implicit treatment: by solving for the magnetization at the future time level, the method effectively filters the high-frequency precession, allowing the error controller to respond to the slowly varying envelope.

The IMEX solver exhibits markedly different behavior, with oscillatory time steps that correlate with the magnetization dynamics. As $|\mathbf{H}_{\text{eff}}|$ varies during precession, the local truncation error of the explicit term fluctuates, causing the adaptive controller to adjust the step size continuously. This results in $\Delta t \approx 0.15\text{--}0.3\text{ ps}$ with significant variation, and a step-rejection rate of approximately 16%. The solver effectively *chases* the oscillating error estimate rather than settling into a stable large time step.

Note that the tolerances employed here ($\text{reitol}=10^{-2}$, $\text{abstol}=10^{-4}$) are stricter than those commonly used in micromagnetic simulations. Relaxing to $\text{reitol}=0$ and $\text{abstol}=10^{-3}$ [158] permits larger time steps at the expense of increased constraint violation, which may be acceptable when only qualitative switching behavior is of interest.

5.6 Summary

This chapter developed a theoretical and computational framework for temporal discretization of the LLG equation in micromagnetic simulations. The key contributions are:

- (i) **Identification of challenges:** The stiffness arising from the exchange interaction, the unit-sphere constraint, and the geometric structure of the LLG equation impose specific requirements on time-stepping schemes.
- (ii) **Tangent-plane formulation:** The TPS reformulates the nonlinear LLG equation as a linear saddle-point system at each time step, with exact constraint enforcement via Lagrange multipliers.
- (iii) **Stability vs. accuracy:** The TPS with $\Delta t = 1\text{ ps}$ demonstrates that unconditional stability does not imply accuracy. While the A-stable scheme remains bounded, the first-order truncation error accumulates to produce significant phase drift ($e_\infty > 0.25$). This motivates the use of adaptive methods that automatically balance step size against accuracy requirements.
- (iv) **Explicit integration:** The classical RK4 method provides a straightforward alternative with excellent constraint preservation ($\mathcal{O}(\Delta t^5)$ drift per step) when the CFL stability limit is acceptable.
- (v) **Adaptive BDF methods:** The CVODE solver with order-2 BDF achieves a $2\times$ wall-clock speedup over explicit methods by taking time steps $10\text{--}50\times$ larger than the CFL limit. The order restriction to $q \leq 2$ ensures A-stability for the oscillatory eigenvalues arising from LLG precession.

- (vi) **IMEX performance:** The IMEX method achieves a $2.8\times$ speedup over TPS, comparable to BDF. Although the time step remains accuracy-limited by explicit precession ($\Delta t \approx 0.3\text{--}0.5$ ps), the SPD structure of the implicit Jacobian enables highly efficient AMG solves, compensating for the larger step count.

The verification studies reveal that all implemented solvers, TPS, RK4, BDF, and IMEX, produce nearly identical magnetization trajectories that agree well with the μMAG reference solutions. The explicit RK4 method exhibits the smallest constraint deviation due to its fourth-order accuracy, while the TPS provides a geometrically motivated first-order scheme with exact tangent-plane enforcement. For problems involving rapid transients or widely varying time scales, the adaptive CVODE and ARKStep integrators from *SUNDIALS* provide automatic step-size control at the cost of increased complexity in preconditioner design.

The central finding of this chapter is that adaptive implicit methods provide the best combination of efficiency and accuracy for micromagnetic simulations. Fixed step methods face a difficult trade-off: small time steps (e.g., $\Delta t = 0.1$ ps for TPS) ensure accuracy but require many steps, while large time steps ($\Delta t = 1$ ps) improve efficiency but sacrifice accuracy. Adaptive methods like BDF and IMEX resolve this trade-off automatically, achieving $2.5\text{--}2.8\times$ speedup over the TPS baseline while maintaining solution accuracy within user-specified tolerances. The BDF method takes fewer, larger time steps (1099 steps at 1–5 ps), while the IMEX method takes more, smaller steps (3818 steps at 0.3–0.5 ps) but benefits from simpler SPD linear systems. The TPS, although it does not offer adaptive time stepping, enforces the tangent-plane constraint exactly and is therefore suitable for applications where constraint preservation is paramount.

A key contribution of this chapter is the adaptation of established time integration techniques to a standard Galerkin FEM framework for micromagnetics. While the individual ingredients exist in the literature, namely BDF integration via CVODE [161] and the Suess preconditioner [158] (originally developed for the box scheme), their combination with Galerkin FEM required non-trivial modifications. In particular, the analysis in Section 5.4.1 showed that the standard Galerkin mass-lumping procedure destroys the skew-symmetric structure exploited by the original preconditioner, necessitating the simplified SPD variant of Equation (5.47). Similarly, the IMEX splitting proposed by Di Fratta *et al.* [176] was implemented here using the adaptive ARKStep solver from the *SUNDIALS* library, providing automatic step-size control that is not available in existing tangent-plane implementations. For general-purpose micromagnetic simulations, either BDF or IMEX represents a practical choice: BDF requires a non-symmetric preconditioner (e.g., ILU) but takes fewer, larger steps, while IMEX maintains SPD structure in the implicit partition, enabling efficient AMG preconditioning at the cost of more, smaller steps.

Spin-Transfer Torques in Ultra-Scaled MRAM Cells

As introduced in Section 3.3, the torque generated by the polarization of electrons passing through the ferromagnetic layers in an MTJ stack drives the writing process in modern MRAM cells. While the standard Slonczewski macro-spin models provide fundamental insight into STT dynamics, they typically assume spatially uniform magnetization or neglect the complex interplay of spin diffusion in ultra-scaled geometries.

This chapter presents a comprehensive FEM-based modeling and simulation approach for computing spin-transfer torques in ultra-scaled MRAM cells. This approach builds directly upon the computational framework established in Chapter 4, utilizing the hybrid FE-BEM method for magnetostatics and employing the time integration schemes detailed in Chapter 5 to resolve the stiff dynamics of the coupled system.

The methodology relies on the coupled spin and charge drift-diffusion formalism, which provides a unified framework for computing both the interface-induced Slonczewski torques and the bulk Zhang–Li torques arising from magnetization textures. The drift-diffusion approach, which accounts only for semiclassical transport properties, is extended to incorporate tunneling across an MTJ via appropriate boundary conditions at the tunnel-barrier interfaces. Furthermore, to address the complexity of modern SAFs and composite free layers, the solver is extended to include a numerical implementation of IEC.

The presented approach reproduces both the angular and the voltage dependencies expected in MTJs. It also predicts an interdependence between the torque acting in the presence of magnetization gradients and the tunneling spin-current. The solver is then applied to simulate the magnetization dynamics of ultra-scaled multi-layered structures with composite ferromagnetic layers.

The switching process can be simulated by solving the LLG equation for magnetization dynamics, including a term describing the torque acting on the magnetization. A more complete description of the process, beyond the Slonczewski approximation, can be obtained by computing the non-equilibrium spin-accumulation $\mathbf{S}$ across the whole structure. The general expression for the torque term is:

$$\mathbf{T}_S = -\frac{D_e}{\lambda_J^2} \mathbf{m} \times \mathbf{S} - \frac{D_e}{\lambda_\varphi^2} \mathbf{m} \times (\mathbf{m} \times \mathbf{S}), \quad (6.1)$$

where D_e is the electron diffusion coefficient. The first term describes precession around the exchange field and is characterized by the exchange length λ_J . The second term describes the dephasing process of the spins of the conducting electrons, characterized by the dephasing length λ_φ .

The spin-accumulation $\mathbf{S}$ describes the deviation of the polarization of the conducting electrons from the equilibrium configuration created by a charge-current density $\mathbf{J}_C$, expressed in units of the transported magnetic moment (A m^{-1}). By definition, $\mathbf{S}$ is non-zero only when an electric current flows through the system. A solution for $\mathbf{S}$ in all non-magnetic and ferromagnetic layers of an MRAM cell can be obtained by means of the spin and charge drift-diffusion formalism.

6.1 Spin and Charge Drift-Diffusion Formalism

6.1.1 Derivation of the Transport Equations

Valet and Fert [179] first derived the spin and charge drift-diffusion equations for a spin-valve structure with collinear (parallel or anti-parallel) magnetization in the ferromagnetic layers. Zhang, Levy, and Fert [35] later reported a formulation of the equations to compute $\mathbf{S}$ in the presence of arbitrary magnetization orientations.

In a magnetic multilayer with the current perpendicular to the plane of the layer, along the x -direction, the linear response of the current to the electric field can be written in spinor form as:

$$\hat{j} = \hat{\sigma} \hat{E} + \hat{D} \frac{\partial \hat{n}}{\partial x}, \quad (6.2)$$

where $\hat{E}$ is the electric field, which can be expressed in terms of the electric potential as $\hat{E} = -\partial \hat{V} / \partial x$, and $\hat{j}$, $\hat{\sigma}$, $\hat{D}$, and $\hat{n}$ are the spinor matrices representing the current density, the conductivity, the diffusion coefficient, and the accumulation at a given position, respectively. For a metal, the diffusion constant and the conductivity are connected via the Einstein relation $\hat{D} = \hat{\sigma} / e^2 \hat{N}$, with $\hat{N}$ the spin-dependent density of states at the Fermi level.

These matrices can be expressed in terms of the vector of Pauli matrices $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \sigma_3)$:

$$\hat{n} = nI + \mathbf{S} \cdot \boldsymbol{\sigma}, \quad (6.3)$$

$$\hat{\sigma} = \sigma I + \boldsymbol{\sigma}_\sigma \cdot \boldsymbol{\sigma}, \quad (6.4)$$

$$\hat{D} = D_e I + \mathbf{D} \cdot \boldsymbol{\sigma}, \quad (6.5)$$

where n is the charge accumulation, σ is related to the conductivity, and D_e is related to the electron diffusion coefficient. The vectors $\boldsymbol{\sigma}_\sigma$ and $\mathbf{D}$ correspond to the different conductivities and diffusion constants for the majority and minority electrons, denoted σ_+ , σ_- , D_+ , and D_- , respectively. This difference is characterized by the conductivity polarization parameter $\beta_\sigma = (\sigma_+ - \sigma_-) / (\sigma_+ + \sigma_-)$ and the diffusion polarization parameter $\beta_D = (D_+ - D_-) / (D_+ + D_-)$.

6.1.2 Three-Dimensional Formulation

Generalizing the one-dimensional formulation to three dimensions yields the following expressions for the charge-current density vector $\mathbf{J}_C$ and the spin-current tensor $\tilde{\mathbf{J}}_S$:

$$\mathbf{J}_C = \sigma \mathbf{E} + \frac{e}{\mu_B} \beta_D D_e (\nabla \mathbf{S})^T \mathbf{m}, \quad (6.6)$$

$$\tilde{\mathbf{J}}_S = -\frac{\mu_B}{e} \beta_\sigma (\mathbf{m} \otimes \sigma \mathbf{E}) - D_e \nabla \mathbf{S}, \quad (6.7)$$

where $\otimes$ denotes the outer product, $\mathbf{m}$ is the unit magnetization vector, and the components of $\tilde{\mathbf{J}}_S$ indicate the flow of the i -th component of spin polarization in the j -th direction. The term $(\nabla \mathbf{S})^T \mathbf{m}$ is a vector with components $\sum_i m_i \partial S_i / \partial x_j$. The factor μ_B / e converts from the units of electric charge to the units of the spin polarization current density.

6.1.3 Spin-Accumulation Equation

The equation of motion for the spin-accumulation in a ferromagnet is determined by the interaction of the local magnetic moments and the spin of the itinerant electrons. This interaction is described by the Hamiltonian term:

$$H_{sd} = -J_{sd} \mathbf{m} \cdot \boldsymbol{\sigma}, \quad (6.8)$$

where J_{sd} is the coupling strength between local moments and itinerant electrons. From this interaction, the equation of motion for the spin-accumulation is derived as follows:

$$\frac{d\mathbf{S}}{dt} = -\frac{\mathbf{S}}{\tau_{sf}} - \frac{D_e}{\lambda_J^2} \mathbf{m} \times \mathbf{S}, \quad (6.9)$$

where τ_{sf} is the spin-flip relaxation time of the conduction electrons. The second term describes the precessional motion of the accumulation due to exchange interaction, which becomes important when the non-equilibrium spin-accumulation and the local magnetization are not parallel.

The typical time-scales for the spin-accumulation and the magnetization differ significantly: while the former is of the order of picoseconds, the latter is of the order of nanoseconds [35]. For the computation of the spin torque to be added to the LLG equation, it is thus sufficient to consider a steady-state expression for the spin-accumulation. With $\partial \mathbf{S} / \partial t = 0$, the equation describing the spin-accumulation becomes:

$$-\nabla \cdot \tilde{\mathbf{J}}_S - D_e \frac{\mathbf{S}}{\lambda_{sf}^2} + D_e \frac{\mathbf{m} \times \mathbf{S}}{\lambda_J^2} = 0, \quad (6.10)$$

where the spin-flip length λ_{sf} and the exchange length λ_J are defined as $\lambda_{sf} = \sqrt{D_e \tau_{sf}}$ and $\lambda_J = \sqrt{D_e \hbar / (2J_{sd})}$.

6.1.4 Spin Dephasing and Complete Transport Equations

Dephasing arises when, after propagating over a certain distance, different spins have undergone unequal amounts of precession, causing their transverse components to partially cancel. In the context of spin-transfer torque, this effect originates either from variations in electron velocities across the Fermi surface or from spins that precess at the same rate but reach a given position at different times because of scattering. The combined influence of precession and dephasing on spin-accumulation has been formulated within the Continuous Random Matrix Theory (CRMT) framework, and the equivalence between CRMT and the spin- and charge-drift-diffusion formalism has been established.

This equivalence between CRMT and the spin and charge drift-diffusion formalism yields the following relations between the characteristic length scales [180, 181]:

$$\sigma = \frac{1}{\rho_*(1 - \beta^2)}, \quad (6.11)$$

$$\lambda_{\text{sf}}^2 = \frac{l_{\text{sf}}^2}{1 - \beta^2}, \quad (6.12)$$

$$\lambda_J^2 = \frac{l_* l_L}{1 - \beta^2}, \quad (6.13)$$

$$\lambda_\varphi^2 = \frac{l_\perp}{l_L} \lambda_J^2, \quad (6.14)$$

$$\lambda^2 = \frac{l_*^2}{1 - \beta^2}, \quad (6.15)$$

where l_{sf} is the spin-flip diffusion length, l_* is the average mean free path of majority and minority electrons defined as $1/l_* = 1/l_\uparrow + 1/l_\downarrow$, l_L is the Larmor spin precession length, $l_\perp$ is the transverse spin coherence length, ρ_* is the spin-averaged resistivity, λ_φ is the spin dephasing length, $\lambda = \sqrt{D_e \tau}$ is the momentum relaxation length, τ is the time for momentum relaxation, and $\beta = (l_\uparrow - l_\downarrow)/(l_\uparrow + l_\downarrow)$ is the spin asymmetry parameter with $l_\uparrow$ and $l_\downarrow$ the mean free paths for majority and minority electrons, respectively.

The complete set of three-dimensional spin and charge drift-diffusion equations, including the dephasing term, takes the form [32]:

$$\mathbf{J}_C = \sigma \mathbf{E} + \frac{e}{\mu_B} \beta_D D_e (\nabla \mathbf{S})^T \mathbf{m}, \quad (6.16)$$

$$\tilde{\mathbf{J}}_S - \frac{\lambda^2}{\lambda_J^2} [\mathbf{m}]_\times \tilde{\mathbf{J}}_S - \frac{\lambda^2}{\lambda_\varphi^2} [\mathbf{m}]_\times \times \tilde{\mathbf{J}}_S = -\frac{\mu_B}{e} \beta_\sigma \mathbf{m} \otimes (\sigma \mathbf{E}) - D_e \nabla \mathbf{S}, \quad (6.17)$$

$$\nabla \cdot \mathbf{J}_C = 0, \quad (6.18)$$

$$-\nabla \cdot \tilde{\mathbf{J}}_S - D_e \frac{\mathbf{S}}{\lambda_{\text{sf}}^2} + D_e \frac{\mathbf{m} \times \mathbf{S}}{\lambda_J^2} + D_e \frac{\mathbf{m} \times (\mathbf{m} \times \mathbf{S})}{\lambda_\varphi^2} = 0. \quad (6.19)$$

where $[\mathbf{m}]_{\times}$ and $[\mathbf{m}]_{\times\times}$ denote the matrices associated with the cross-product and double cross-product with unit magnetization vector $\mathbf{m}$:

$$[\mathbf{m}]_{\times} = \begin{pmatrix} 0 & -m_z & m_y \\ m_z & 0 & -m_x \\ -m_y & m_x & 0 \end{pmatrix}, \quad [\mathbf{m}]_{\times\times} = \begin{pmatrix} m_x^2 - 1 & m_x m_y & m_x m_z \\ m_x m_y & m_y^2 - 1 & m_y m_z \\ m_x m_z & m_y m_z & m_z^2 - 1 \end{pmatrix}. \quad (6.20)$$

The result $\nabla \cdot \mathbf{J}_C = 0$ stems from the absence of electric current sources within the metallic layers, due to the rapid redistribution of any charge imbalance. Equation (6.19) differs from the formulation without dephasing by the inclusion of the spin dephasing term dependent on λ_φ .

The present work focuses on a finite element implementation that excludes the additional terms on the left-hand side of Equation (6.17), which mix the orthogonal spin-current components depending on the local magnetization orientation. These terms originate from the underlying ballistic nature of the transverse spin precession or dephasing and rely mainly on the ratio between the momentum relaxation length and the transverse absorption lengths λ_J and λ_φ . In transition-metal ferromagnets, these length scales are of the same order of magnitude. This implementation demonstrates that an appropriate treatment of the tunneling layer and tuning of the system parameters reproduce the most crucial torque properties expected in MTJs, while retaining the ability to obtain all torque contributions across several ferromagnetic layers from a unified formalism.

6.2 Finite Element Implementation

6.2.1 Coupled LLG and Drift-Diffusion Framework

The LLG equation, coupled with the spin and charge drift-diffusion formalism for computing torques, enables the description of the magnetization dynamics of structures containing an arbitrary number of layers. The LLG equation is given by:

$$\frac{\partial \mathbf{m}}{\partial t} = -\gamma \mathbf{m} \times \mathbf{H}_{\text{eff}} + \alpha \mathbf{m} \times \frac{\partial \mathbf{m}}{\partial t} + \frac{1}{M_S} \mathbf{T}_S, \quad (6.21)$$

where γ is the gyromagnetic ratio, α is the Gilbert damping parameter, M_S is the saturation magnetization, and $\mathbf{H}_{\text{eff}}$ is the effective field. The effective field includes the magnetic anisotropy field, the exchange field, and the demagnetization field. As detailed in Chapter 4, the demagnetization field contribution is evaluated only on the disconnected magnetic domains by using the hybrid approach FE-BEM. The TPS described in Chapter 5 handles the temporal evolution, ensuring numerical stability even with the stiff exchange terms.

Analytical solutions to these sets of partial differential equations are possible only in simplified scenarios, thus numerical methods are necessary to resolve the dynamics in general settings. The FEM naturally handles meshes with complex geometries and multiple material domains. Therefore, it serves to implement a solver capable of simulating charge, spin-accumulation, and magnetization dynamics. The implementation utilizes the open-source C++ FE library MFEM [145].

6.2.2 Treatment of the Tunnel Barrier




Figure 6.1: (Left) Current density distribution in an MTJ biased under a constant voltage for a non-uniform magnetization configuration of the FL. The magnetization varies from parallel to the reference layer in the center to anti-parallel on the sides. (Right) The current density is larger in the center, where the FL magnetization is parallel to that of the RL and the magnetization-dependent conductivity is highest. Based on the figure originally published in [182].

A key challenge in applying the drift-diffusion formalism to MTJs is the treatment of the TB. While the TMR effect can be reproduced by describing the TB conductivity through an angle-dependent expression, using effective material parameters in the spin diffusion equation is insufficient to capture all expected torque properties. The TB is modeled as a poor conductor with a local resistance that depends on the relative orientation of the magnetization [123]. Accordingly, this dependence is introduced phenomenologically by expressing the TB conductivity as an explicit function of the relative magnetization angle, given by:

$$\sigma(\theta) = \frac{\sigma_P + \sigma_{AP}}{2} \left(1 + \frac{\text{TMR}}{2 + \text{TMR}} \cos \theta \right), \quad (6.22)$$

where σ_P and σ_{AP} are the conductivities in the parallel and anti-parallel configurations, respectively, TMR is the tunneling magnetoresistance ratio, and θ is the angle between the RL and FL magnetizations. This is a manifestation of Ohm's law relating the voltage and the charge-current through a structure with many transverse modes [32]. Figure 6.1 shows the charge-current density computed in an MTJ structure with non-uniform magnetization in the FL. The magnetization varies from parallel to the RL in the center to anti-parallel at the edges. Due to the angular dependence of the TB conductivity, the current density is highest in the center, where the magnetizations are parallel, demonstrating the expected TMR-induced current redistribution.

6.3 Tunneling Spin-Current Boundary Conditions

The drift-diffusion approach accounts only for semiclassical transport properties and must be adapted to describe tunneling across an MTJ. The Nonequilibrium Green

Functions (NEGF) formalism is frequently employed to compute expressions for the charge and spin-current flowing through the TB of an MTJ. However, performing such calculations at every iteration of the FE solver would be computationally prohibitive. Using a circuit-theory approach and analytical solutions of the NEGFs equations, the expressions for the tunneling spin and charge-currents can be simplified for practically thick TBs to capture the most prominent transport characteristics by using only a few key parameters.

6.3.1 Tunneling Current Expressions

With the magnetization in the RL pointing in the x -direction and the one in the FL lying in the xz -plane at an angle θ from the RL, the tunneling charge- and spin-currents can be expressed as [118]:

$$J_C^{TB} = J_0(V) (1 + P_{RL}P_{FL} \cos \theta), \quad (6.23)$$

$$J_{S,x}^{TB} = -\frac{a_{RL}P_{RL} + a_{FL}P_{FL} \cos \theta}{1 + P_{RL}P_{FL} \cos \theta} \frac{\hbar}{2e} J_C^{TB}, \quad (6.24)$$

$$J_{S,y}^{TB} = -\frac{\frac{1}{2}(P_{RL}P_{RL}^\eta - P_{FL}P_{FL}^\eta) \sin \theta}{1 + P_{RL}P_{FL} \cos \theta} \frac{\hbar}{2e} J_C^{TB}, \quad (6.25)$$

$$J_{S,z}^{TB} = -\frac{a_{FL}P_{FL} \sin \theta}{1 + P_{RL}P_{FL} \cos \theta} \frac{\hbar}{2e} J_C^{TB}, \quad (6.26)$$

where $J_0(V)$ contains the voltage-dependent portion of the current density, P_{RL} and P_{FL} are the in-plane Slonczewski polarization parameters, P_{RL}^η and P_{FL}^η are out-of-plane polarization parameters, and a_{RL} and a_{FL} describe the influence of the interface spin-mixing conductance on the transmitted in-plane spin-current. The spin-current is expressed in units of $J m^{-2}$.

Equation (6.23), describing the dependence of the charge-current on the RL and FL polarization parameters and on the relative angle between their magnetization vectors, is already included in the model through Equation (6.22). The spin-current component must also be accounted for. The FE approach to the drift-diffusion equations, based on the Galerkin method, enforces the continuity of the spin-accumulation and current across all interfaces. The tunneling expressions must therefore be included while preserving the continuous nature of both quantities.

6.3.2 Finite Element Boundary Terms

To obtain the solution through the FE solver, the following boundary terms must be added on the RHS of the weak formulation (5.25):

$$-\int_{RL|TB} \mathbf{J}_S^{TB} \cdot \mathbf{w} \, d\mathbf{x} + \int_{TB|FL} \mathbf{J}_S^{TB} \cdot \mathbf{w} \, d\mathbf{x}, \quad (6.27)$$

$$\begin{aligned} \mathbf{J}_S^{TB} = & -\frac{\mu_B}{e} \frac{\mathbf{J}_C \cdot \mathbf{n}}{1 + P_{RL}P_{FL} \mathbf{m}_{RL} \cdot \mathbf{m}_{FL}} \left[a_{RL}P_{RL} \mathbf{m}_{RL} + a_{FL}P_{FL} \mathbf{m}_{FL} \right. \\ & \left. + \frac{1}{2} (P_{RL}P_{RL}^\eta - P_{FL}P_{FL}^\eta) \mathbf{m}_{RL} \times \mathbf{m}_{FL} \right], \end{aligned} \quad (6.28)$$

where RL|TB and TB|FL are the interfaces between the TB and the RL and FL, respectively, $\mathbf{n}$ is the interface normal (which for the employed meshes points in the positive x -direction for both interfaces), and $\mathbf{m}_{\text{RL(FL)}}$ is the unit magnetization vector of the RL (FL) at the interface. This formulation is consistent with physics-based factorization approaches for MTJ modeling [183]. The diffusion coefficient of the TB D_S is also taken to be low, proportional to the conductivity from Equation (6.22), as no diffusive transport occurs across the barrier.

By solving the weak formulation with the inclusion of Equation (6.27) and the low value of D_S , the spin-current in the TB is fixed to the value prescribed by Equation (6.28) when $\mathbf{J}_C$ flows through the MTJ. This is crucial for accurately describing the spin-currents and the spin-accumulations in the RL and FL of an MTJ.

6.3.3 Implementation of Interface Mapping

The computation of the boundary term associated with Equation (6.27) requires knowledge of the magnetization vector on the opposite interface. To access these values, the coefficient describing the boundary integral is initialized as follows:

- (i) For each integration point on the RL|TB interface requiring the computation of the tunneling spin-current, the solver iterates through the integration points of the TB|FL interface.
- (ii) The TB|FL point with coordinates closest to the current RL|TB one is selected.
- (iii) The integration point index and element index associated with the identified TB|FL point are mapped to the coordinates of the RL|TB point.
- (iv) The mapping procedure is repeated for the TB|FL interface.

In a transient simulation, the search is performed only during solver initialization. At every time step, the data necessary for the computation of Equation (6.28) can be accessed through the generated maps, without the need to repeat the search procedure.

6.4 Numerical Implementation of Interlayer Exchange Coupling

As discussed in Section 6.3.3, the IEC is crucial for SAFs. While Chapter 3 defined the physics, this section details the numerical implementation in the FEM solver, which follows the same logic as the TB mapping described above.

6.4.1 IEC Energy and Effective Field

The interlayer exchange coupling energy density for two ferromagnetic layers separated by a thin non-magnetic spacer (such as Ru or Ta) is given by:

$$\mathcal{E}_{\text{iec}} = -J_{\text{iec}}(\mathbf{m}_1 \cdot \mathbf{m}_2), \quad (6.29)$$

where J_{iec} is the interlayer exchange coupling constant (in J m^{-2}), and $\mathbf{m}_1$ and $\mathbf{m}_2$ are the unit magnetization vectors at the two interfaces. The sign of J_{iec} determines the preferred alignment: positive values favor parallel alignment (ferromagnetic coupling), while negative values favor anti-parallel alignment (antiferromagnetic coupling), as required for SAF structures.

In dynamic micromagnetics, representing energy contributions in terms of effective fields rather than boundary conditions is often more convenient, as it facilitates direct integration into the LLG equation. The effective field contribution from the IEC acting on a layer at the interface is obtained via $\mathbf{H}_{\text{eff}} = -\frac{1}{\mu_0 M_s} \frac{\partial \mathcal{E}}{\partial \mathbf{m}}$, yielding:

$$\mathbf{H}_{\text{iec}} = \frac{J_{\text{iec}}}{\mu_0 M_s t} \mathbf{m}_{\text{neighbor}}, \quad (6.30)$$

where t is the thickness of the ferromagnetic layer.

6.4.2 IEC Boundary Term in Weak Formulation

In the FE weak formulation of the LLG equation, the IEC contribution appears as boundary integrals at the ferromagnet-spacer interfaces. For a structure with left (L) and right (R) ferromagnetic layers separated by a spacer (see Figure 3.3), the following boundary terms must be added to the RHS of the weak formulation (5.25):

$$\frac{J_{\text{iec}} \gamma}{\mu_0 M_{s,L}} \int_{\Gamma_{\text{R}|\text{spacer}}} \mathbf{m}_L \cdot \mathbf{w} \, d\Gamma + \frac{J_{\text{iec}} \gamma}{\mu_0 M_{s,R}} \int_{\Gamma_{\text{spacer}|\text{L}}} \mathbf{m}_R \cdot \mathbf{w} \, d\Gamma, \quad (6.31)$$

where $\Gamma_{\text{R}|\text{spacer}}$ and $\Gamma_{\text{spacer}|\text{L}}$ denote the interfaces between the spacer and the right and left ferromagnetic layers, respectively, and $\mathbf{w}$ is the test function.

6.4.3 Mapping Strategy for IEC

To evaluate Equation (6.31), the solver employs the same mapping strategy as described in Section 6.3.3 for the tunneling spin-current:

- (i) During initialization, for each integration point on one side of the IEC interface, the solver identifies the nearest integration point on the opposite side.
- (ii) The element ID and local integration point index are stored in a mapping table.
- (iii) During time-stepping, the magnetization values $\mathbf{m}_{\text{neighbor}}$ are accessed through these pre-computed maps with $O(1)$ complexity.

This unified approach allows the simulation of arbitrary stacks (e.g., Co/Ru/Co SAFs).

6.5 Validation Results

6.5.1 Spin-Accumulation and Spin Torque Solutions

The spin-accumulation and torque obtained from the parameters of Table 6.1 are presented in Figure 6.2 (a) and (b). The structure employed consists of semi-infinite ferromagnetic leads and a TB of 1 nm thickness. The magnetization of the RL points in the x -direction, while that of the FL points in the z -direction. A comparison between the FE solution and the analytical solutions demonstrates that the additional boundary terms in the FE implementation yield perfect agreement.

The inclusion of the tunneling spin-current terms creates a discontinuity jump between the values of the spin-accumulation components parallel to the magnetization at the left and right interfaces of the TB. This is a manifestation of the MTJ polarization effects on the spin-current.

Figure 6.2 (c) shows the angular dependence of the damping-like torque with the inclusion of the spin-current boundary conditions, computed in the same structure. The typical sinusoidal dependence of the torque acting on the FL in an MTJ is reproduced exactly, for various values of the RL|TB interface spin polarization. The structure is biased by a fixed voltage of $V = 1.3$ V, so that the torque is independent of the TB|FL polarization, and linearly depends only on the RL|TB polarization.

6.5.2 Effect of Dephasing Length on Torque Localization

While employing Equation (6.28) fixes the spin current in the TB to the value expected in MTJs, the length parameters entering the spin-accumulation equation still determine the region over which the transverse spin-accumulation components are absorbed. They also determine the behavior of the torque in the bulk of the FM layers.




Figure 6.2: Validation of the spin drift-diffusion model with tunneling spin-current boundary conditions for semi-infinite ferromagnetic layers. (a) spin-accumulation and (b) spin-transfer torque components computed along the stack direction, comparing the analytical solution with the FE implementation. The magnetization is oriented along x in the RL and along z in the FL. (c) Angular dependence of the damping-like torque for several values of the RL spin polarization ($P_{RL} = 0.4, 0.7, 0.9$), reproducing the expected sinusoidal behavior characteristic of MTJs. Figure (c) adapted from [182].

Table 6.1: Parameters used in the drift-diffusion simulations with the TB boundary terms.

Parameter	Value
Conductivity polarization, β_σ	0.7
Diffusion polarization, β_D	0.8
NM diffusion coefficient, $D_{e,NM}$	$10^{-2} \text{ m}^2 \text{ s}^{-1}$
FM diffusion coefficient, $D_{e,FM}$	$2.0 \times 10^{-3} \text{ m}^2 \text{ s}^{-1}$
TB diffusion coefficient, D_S	$2.0 \times 10^{-8} \text{ m}^2 \text{ s}^{-1}$
NM conductivity, σ_{NM}	$5.0 \times 10^6 \text{ S m}^{-1}$
FM conductivity, σ_{FM}	$1.0 \times 10^6 \text{ S m}^{-1}$
TB conductance, G_0	$4.76 \times 10^{-5} \text{ S}$
Polarization factors, $P_{RL} = P_{FL}$	0.707
In-plane torque reduction, $a_{RL} = a_{FL}$	1.0
Out-of-plane polarization, $P_{RL}^\eta = P_{FL}^\eta$	0.2
Spin-flip length, λ_{sf}	10 nm
Spin exchange length, λ_J	1 nm
Spin dephasing length, λ_φ	0.4 nm

Figure 6.3 shows the spin-accumulation and torque obtained in a symmetrical MTJ structure including 50 nm thick non-magnetic (NM) layers, where the FM layers are 20 nm thick and the TB is 1 nm thick. For $\lambda_\varphi = 5 \text{ nm}$, the torque components are not completely absorbed in the FL, contrary to what is usually expected in strong ferromagnets, and the spin-accumulation components transverse to the local magnetization penetrate inside the NM contacts. Setting an effective dephasing length of $\lambda_\varphi = 0.4 \text{ nm}$, it guarantees a faster decay of the transverse spin-accumulation components, ensuring that the torque acts only in the proximity of the TB interface.

6.6 GMR Effect in Spin-Valves

The approach detailed in the previous sections can compute both the TMR and the torques in an MTJ. Modern MRAM devices can, however, also include conductive spacer layers between the FM layers, either as part of the SAF or to incorporate additional RLs and improve switching performance. In the presence of such a metallic spacer layer, the GMR effect becomes significant.

In the drift-diffusion formalism, the dependence of resistance on the relative angle between the magnetization vectors in the presence of a metallic spacer is accounted for by including the second RHS term of Equation (6.6) when computing the electrical potential V . By imposing $\nabla \cdot \mathbf{J}_C = 0$ and expressing the electric field as $\mathbf{E} = -\nabla V$ in Equation (6.6), the equation for V yields:

$$\nabla \cdot (\sigma \nabla V) = \nabla \cdot \left(\frac{e}{\mu_B} \beta_D D_e (\nabla \mathbf{S})^T \mathbf{m} \right). \quad (6.32)$$




Figure 6.3: Spin-accumulation and torque in a symmetric MTJ structure including NM contacts. (a) and (b) present results for $\lambda_\varphi = 0.4$ nm. The shorter dephasing length guarantees faster decay of the transverse spin-accumulation components, so that the torque acts only in the proximity of the TB interface. Adapted from the figure [182].

In this formulation, the spin-accumulation and the electrical potential are coupled, yielding a coupled system of equations. To account for the interdependence of $\mathbf{S}$ and V , the FE solver is updated to iteratively solve the system until a convergence threshold is reached:

$$\frac{\|\mathbf{S}_n\|_{L^2} - \|\mathbf{S}_{n-1}\|_{L^2}}{\|\mathbf{S}_n\|_{L^2}} < \epsilon, \quad (6.33)$$

where $\|\cdot\|_{L^2} = \sqrt{\int_{\Omega} |\cdot|^2 dx}$ denotes the $L^2(\Omega)$ norm.

The iterative solution is computed by first obtaining an initial estimate for the spin-accumulation $\mathbf{S}_0$ by solving the standard equations, including the RHS term of Equation (6.6) in the spin equation. This estimate is then used to compute a solution to Equation (6.32), which in turn yields an updated spin-accumulation estimate $\mathbf{S}_1$. These two steps are repeated until the solver converges. The dependence of the total charge-current on the relative angle between the magnetization vectors in the FL and RL can be fitted using [117]:

$$I(\theta) = \frac{V}{R_P} \frac{1 + \chi \cos^2(\theta/2)}{1 + \text{GMR} + (\chi - \text{GMR}) \cos^2(\theta/2)}, \quad (6.34)$$

where V is the applied bias voltage, R_P is the resistance in the parallel state, and χ and GMR are fitting parameters.

Figure 6.4 compares the angular dependence of the charge current for an MTJ with an MgO TB and a spin valve with a metallic spacer. Both structures employ symmetric 2 nm FM layers separated by a 1 nm spacer, with 50 nm NM contacts and a diameter of 40 nm. The MTJ is biased at -1.3 V and simulated for two resistance configurations: $R_P = 14$ k Ω , $R_{AP} = 42$ k Ω (TMR = 200 %) and $R_P = 14$ k Ω , $R_{AP} = 35$ k Ω (TMR = 150 %). Both are well described by the Slonczewski conductance model




Figure 6.4: Angular dependence of the charge current for (a) an MTJ with MgO tunnel barrier at two TMR ratios ($V = -1.3\text{ V}$) and (b) a spin valve with metallic spacer at two GMR ratios ($V = -0.2\text{ V}$). Markers represent simulation data and solid/dashed lines show analytical fits using Equation (2.7) and Equation (6.34), respectively.

(Equation (2.7)). The spin valve is biased at -0.2 V , and the iterative coupling between the spin-accumulation and the electrical potential produces a GMR of approximately 11 % for a conductivity polarization $\beta_\sigma = 0.7$ and 7.5 % for $\beta_\sigma = 0.58$, both fitted by Equation (6.34).

Since the voltage drop in an MTJ is localized at the TB, the iterative solution is not necessary for the correct computation of the current. In this case, the iterative solver always converges for $n = 1$.

6.7 Torques in Elongated Structures

In the presence of elongated FLs, simultaneous switching of the entire layer is not guaranteed: domain walls or magnetization textures may form, and their propagation through the FL can affect the switching behavior. In such cases, the additional spin torques arising from magnetization gradients in the bulk of the ferromagnetic layers must be taken into account. These torques are modeled by the Zhang and Li (ZL) equation [184]:

$$\mathbf{T}_{\text{ZL}} = -\frac{\mu_{\text{B}}}{e} \frac{\beta_\sigma}{1 + \epsilon^2} (\mathbf{m} \times [\mathbf{m} \times (\mathbf{J}_{\text{C}} \cdot \nabla) \mathbf{m}] - \epsilon \mathbf{m} \times (\mathbf{J}_{\text{C}} \cdot \nabla) \mathbf{m}), \quad (6.35)$$

where $\epsilon = (\lambda_{\text{J}}/\lambda_{\text{sf}})^2$. This expression is derived from Equation (6.10) by setting $\nabla \mathbf{S} = \mathbf{0}$ and taking the cross product with $\mathbf{m}$ twice. This approximation is strictly valid only when the changes in $\mathbf{S}$ are relatively slow, namely, when they occur over length scales longer than λ_{sf} .




Figure 6.5: (a) Non-uniform magnetization texture with the magnetization orientation changing from x to $-z$. (b) Comparison of the spin torque $\mathbf{T}_S$ to the ZL torque $\mathbf{T}_{ZL}$ for an approximately 100 nm long magnetization texture, with the parameters of Table 6.1. (c) Comparison for an approximately 3 nm long magnetization texture with $\lambda_\varphi = 0.4$ nm. The shorter dephasing length quickly dissipates the transverse spin-accumulation, ensuring agreement between the ZL approximation and the drift-diffusion solution. The two approaches show good agreement in both cases.

By applying the same procedure to Equation (6.19), the ZL expression can be generalized to include λ_φ :

$$\mathbf{T}_{ZL} = -\frac{\mu_B}{e} \frac{\beta_\sigma}{1 + (\epsilon + \epsilon')^2} \left[(1 + \epsilon'(\epsilon + \epsilon')) \mathbf{m} \times [\mathbf{m} \times (\mathbf{J}_C \cdot \nabla) \mathbf{m}] - \epsilon \mathbf{m} \times (\mathbf{J}_C \cdot \nabla) \mathbf{m} \right], \quad (6.36)$$

where $\epsilon' = (\lambda_J/\lambda_\varphi)^2$ accounts for the additional contributions of the term involving λ_φ .

6.7.1 Comparison with Drift-Diffusion Solution

To test the ZL approximation, a magnetization profile with the magnetization orientation changing from $+z$ to $-x$ (shown in Figure 6.5 (a)) serves to compute both $\mathbf{T}_S$ and $\mathbf{T}_{ZL}$, with the parameters of Table 6.1. The structure used in the simulations consists of a single ferromagnetic layer 450 nm thick. A constant current density value of $J_{C,x} = -1.33 \times 10^{12}$ A m $^{-2}$, flowing in the x -direction, is employed for both approaches.

Figure 6.5 (b) demonstrates that, for the given magnetization profile with a width of approximately 100 nm, $\mathbf{T}_S$ is well reproduced with Equation (6.35). Figure 6.5 (c) reports the torques obtained using the same parameters and a shorter magnetization texture of approximately 3 nm, in an FM layer of 70 nm thickness. With the standard dephasing length, the spin-accumulation gradients neglected in Equation (6.35) become significant, leading to a large deviation of $\mathbf{T}_S$ from $\mathbf{T}_{ZL}$, particularly for the field-like torque (y -component). Employing a shorter spin dephasing length of $\lambda_\varphi = 0.4$ nm ensures the rapid absorption of transverse spin-accumulation components, restoring good agreement between the $\mathbf{T}_{ZL}$ approximation and the drift-diffusion solution.

6.7.2 Interaction Between Tunneling and Bulk Torques

In MRAM cells with elongated FLs, magnetization textures and domain walls can form along the layer during switching. In this case, the polarized tunneling spin-current can interact with such magnetization textures, generating torques that further deviate from the ZL approximation.

To verify this, the torque is computed in an experimental MTJ structure with a 5 nm RL, 0.9 nm TB, and an elongated FL of 15 nm with a magnetization profile similar to that shown in Figure 6.5 (a), with the magnetization vector transitioning from the z -direction to the $-x$ -direction over the length of the layer. The magnetization in the RL is pointing toward the x -direction and 50 nm NM contacts are included.

The torque acting in the FL for this magnetization profile is shown in Figure 6.6 (a). Near the TB the tunneling spin-current generates an observable Slonczewski torque contribution, while the magnetization texture causes the torque contribution in the bulk. Figure 6.6 (b) shows a close-up of the bulk portion of $\mathbf{T}_S$, compared with the result obtained by applying the ZL expression (6.36) to the magnetization configuration of the FL, for the exact value of the current density. The comparison reveals a substantial difference between the torques obtained using the presented model and those obtained with the ZL approximation. This implies that the Slonczewski and ZL torque contributions cannot simply be added together, but influence and interact with each other.




Figure 6.6: (a) Torque computed for an MRAM cell with elongated RL and FL and a magnetization profile in the FL similar to that of Figure 6.5 (a), with a width of approximately 3 nm. The brown vectors indicate the magnetization direction in the RL and in two parts of the FL. (b) Close-up of the spin torque $\mathbf{T}_S$ compared to the ZL torque $\mathbf{T}_{ZL}$. The presence of the MTJ also influences the bulk torque component, making the unified approach the most suitable for ultra-scaled MTJs with elongated ferromagnetic layers. Figure (a) originally published in [182].

6.8 Preconditioner Evaluation

As established in Section 5.4.1, the inclusion of spin-transfer torque requires a preconditioner that captures the dominant Jacobian contributions. For the BDF method, all terms, including spin-transfer torque, are treated implicitly. Consequently, the preconditioner must incorporate the exchange-damping contribution and the linearized spin-torque terms to remain effective during switching transients.

BDF Preconditioner: Following the mass-lumped formulation of Section 5.4.1, the BDF preconditioner takes the form:

$$\mathbf{P}_{\text{BDF}} = \mathbf{I} + \gamma\alpha C_{\text{ex}}\mathbf{M}_L^{-1}\mathbf{K}_{\text{diff}} - \gamma\mathbf{M}_L^{-1}\mathbf{J}_{\text{STT}}, \quad (6.37)$$

where $\gamma = \Delta t\beta_0$ denotes the BDF scaling factor, α the Gilbert damping constant, $C_{\text{ex}} = 2\gamma_L A_{\text{ex}}/(\mu_0 M_s)$ the exchange coefficient, $\mathbf{M}_L$ the lumped mass matrix, $\mathbf{K}_{\text{diff}}$ the symmetric exchange stiffness matrix arising from damping, and $\mathbf{J}_{\text{STT}} = \mathbf{J}_{\text{DL}} + \mathbf{J}_{\text{FL}}$ comprises the damping-like and field-like Jacobian contributions defined in Equations (5.51) and (5.52).

While the damping-like contribution $\mathbf{J}_{\text{DL}}$ is symmetric, the field-like term $\mathbf{J}_{\text{FL}} = c_{\text{FL}}[\mathbf{S}]_{\times}$ is skew-symmetric. This renders the preconditioner non-symmetric, preventing the use of PCG. Furthermore, classical AMG components (in particular smoothers) are primarily designed for SPD or nearly SPD operators and can lose effectiveness for strongly non-symmetric matrices with large skew-symmetric parts [170]. Incomplete LU factorization handles such systems naturally by directly computing approximate $\mathbf{L}$ and $\mathbf{U}$ factors without relying on smoothing properties. The Euclid parallel ILU implementation from *HYPRE* [171] is employed, and the Newton linear system is solved using GMRES preconditioned by ILU. The spin-accumulation $\mathbf{S}$ entering the STT Jacobian is held constant (*frozen*) between preconditioner rebuilds. During STT-driven switching, both $\mathbf{m}$ and $\mathbf{S}$ evolve rapidly, so the *frozen-S* approximation becomes less accurate and the number of GMRES iterations increases. To limit this effect, the preconditioner is rebuilt periodically rather than only on demand. This is controlled via `CVodeSetLSetupFrequency`, set to 3 in the implementation, balancing ILU setup cost against the overhead of Krylov iterations with a stale preconditioner.

Note that the exchange-precession contribution, $\mathbf{K}_{\text{skew}}$, is omitted from the preconditioner, although it appears in the full Jacobian. Numerically, adding $\mathbf{K}_{\text{skew}}$ in the mass-lumped form degraded Krylov convergence (cf. Section 5.4.1). A likely reason is that the lumped scaling $\mathbf{M}_L^{-1}$ destroys the exact skew-symmetry of the cross-product operator under the relevant inner product, producing a more difficult non-symmetric preconditioned system.

IMEX Preconditioner: In contrast, the IMEX scheme assigns only the stiff exchange damping to the implicit partition f_I , while all other contributions, including spin-transfer torque, are evaluated explicitly in f_E . The implicit Jacobian (5.57) involves only the symmetric exchange operator:

$$\mathbf{P}_{\text{IMEX}} = \mathbf{I} + \gamma\alpha C_{\text{ex}}\mathbf{M}_L^{-1}\mathbf{K}_{\text{diff}}. \quad (6.38)$$

This matrix is SPD, enabling efficient solution via PCG with AMG preconditioning. Since $\mathbf{J}_{\text{STT}}$ does not appear in the implicit partition, it is correctly omitted from the preconditioner, maintaining the SPD structure regardless of spin-torque magnitude.

The performance is evaluated similarly to the Standard Problem 4. For the switching simulation, a simplified MTJ structure with 1 nm thick RL, 1.7 nm thick FL, and a 1 nm thick TB is employed. The magnetization of the RL points in the positive x -direction, while the FL is initialized in the negative x -direction, with a small tilt to break symmetry. A bias voltage of -3 V is applied to the left contact to induce switching. The structure is sized down to a diameter of 10 nm to reduce the computational cost of the tests.

Although this configuration is not physically realistic, the applied bias voltage exceeds typical operating conditions for such a thin tunnel barrier and would likely cause dielectric breakdown in practice. This setup was chosen intentionally to create a STT dominated regime. This allows for a meticulous evaluation of how the different time integrators and preconditioner strategies handle the stiff, non-symmetric dynamics arising from strong spin-transfer torque.




Figure 6.7: Spatially-averaged magnetization dynamics during STT-driven switching computed with different time integrators. The TPS solution with $\Delta t = 0.01$ ps serves as the reference.

Figure 6.7 presents the spatially-averaged magnetization dynamics computed with the four time integration methods. The reference solution is obtained using the TPS method with a fixed time step of $\Delta t = 0.01$ ps. To enable a meaningful comparison at similar temporal resolution, an additional TPS run with $\Delta t = 0.1$ ps (10^{-13} s) is performed, matching the order of magnitude of the adaptive time steps achieved by the BDF solver. All implicit methods (TPS, BDF, IMEX) capture the qualitative switching behavior, with the magnetization reversing from the initial anti-parallel to the parallel configuration under the applied STT. Visible phase differences appear between the adaptive solvers and the reference, particularly in m_y around $t \approx 0.1$ ns. These deviations reflect the inherent trade-off between accuracy and efficiency: both solvers satisfy the prescribed tolerances while taking larger time steps than the reference. The IMEX scheme exhibits approximately twice the e_∞ trajectory error of BDF (0.398 vs. 0.240, cf. Table 6.2), attributable to the splitting error from explicit STT treatment and the accumulation over ~ 22000 steps compared to only ~ 488 for BDF. Despite these deviations, both methods capture the correct switching dynamics and converge to the same equilibrium state. The coarser TPS run ($\Delta t = 0.1$ ps) achieves a $9.7\times$ speedup but exhibits the largest trajectory error ($e_\infty = 0.733$), with noticeable phase drift visible in all three components. This highlights a key limitation of fixed-step methods without error control: unconditional stability allows arbitrarily large time steps, but accuracy degradation goes undetected. In contrast, RK4 shows a comparable phase deviation, but, as discussed below, this apparent agreement is misleading, as the method becomes numerically unstable during switching.

Figure 6.8 provides a quantitative analysis of solver accuracy and efficiency. Panel (a) shows the deviation from the unit-sphere constraint $|\mathbf{m}| = 1$ before renormalization. For explicit single-step methods, theoretical results establish that a scheme of order p preserves the magnetization magnitude within $\mathcal{O}(\Delta t^{p+1})$ error per step, exploiting the cross-product structure of the LLG equation, which guarantees $\partial_t \mathbf{m} \perp \mathbf{m}$. The TPS method, being first-order in its tangent-plane update, produces $\mathcal{O}(\Delta t^2)$ drift per step, accumulating to deviations below 10^{-7} that further decrease as the simulation progresses toward equilibrium.




Figure 6.8: Quantitative solver analysis: (a) deviation from the unit-sphere constraint before renormalization; (b) trajectory deviation from the TPS reference solution; (c) time step evolution for the adaptive solvers compared to the fixed $\Delta t = 0.01$ ps baseline.

In contrast, the explicit RK4 method, despite its fourth-order accuracy (and the associated $\mathcal{O}(\Delta t^5)$ local constraint drift expected for smooth problems), becomes unstable in the switching regime, with $\| |\mathbf{m}| - 1 \|_\infty$ growing to $\mathcal{O}(1)$. The rapid reorientation driven by strong exchange and STT yields a stiff discretized LLG operator whose eigenvalues leave the stability region of explicit Runge–Kutta schemes at the time steps considered, leading to numerical blow-up. Notably, the RK4 trajectory in Figure 6.7 appears to follow the qualitative switching behavior despite this failure. This misleading agreement arises because the instability primarily corrupts the magnetization *magnitude* while its *direction* still roughly follows the torque-driven dynamics. The constraint plot in Figure 6.8 (a) reveals the true extent of the breakdown, underscoring that trajectory agreement alone is insufficient to validate numerical accuracy, the unit-sphere constraint must be monitored explicitly. Consequently, RK4 is not suitable for STT-driven switching unless prohibitively small time steps are used.

The implicit multistep methods (BDF and IMEX) exhibit fundamentally different drift characteristics, maintaining deviations at around 10^{-1} . The $\mathcal{O}(\Delta t^{p+1})$ bound of single-step methods does not apply here: the multistep formulation combines solution values from previous steps that each carry minor constraint violations, and the Newton iteration required for the implicit solve explores intermediate states that may lie far from the unit sphere. Consequently, these methods exhibit larger oscillating deviations that correlate with the adaptive step size, as visible in Figure 6.8 (c). Nevertheless, periodic renormalization maintains the constraint within acceptable bounds throughout the simulation.

Panel (b) reports the deviation from the reference trajectory using the maximum-in-time error metric e_∞ defined in (5.59). The BDF method maintains a deviation below 10^{-1} throughout the simulation, whereas the IMEX scheme accumulates approximately twice the trajectory error. This discrepancy stems from two sources: (i) the splitting error inherent to IMEX methods, where the explicit treatment of STT introduces an additional $\mathcal{O}(\Delta t)$ error term, and (ii) the vastly different step counts, the BDF solver requires only ~ 440 steps compared to ~ 22000 for IMEX, leading to different error accumulation characteristics despite identical local tolerances.

Table 6.2: Quantitative performance comparison for STT-driven switching. The trajectory error is measured by e_∞ relative to the TPS reference (baseline row). Speedup is computed as $\text{Speedup} = \text{CPU}_{\text{TPS,ref}} / \text{CPU}_{\text{method}}$.

Solver	Configuration	e_∞	Steps	CPU [s]	Speedup
TPS _{ref}	$\Delta t = 10^{-14}$ s	—	40000	42705	1.00×
TPS	$\Delta t = 10^{-13}$ s	0.733	4000	4393	9.27×
RK4	$\Delta t = 10^{-14}$ s	0.654 (unstable)	40,000	62623	0.68×
BDF	rtol= 10^{-2} , atol= 10^{-4}	0.240	488	12870	3.32×
IMEX	rtol= 10^{-2} , atol= 10^{-4}	0.398	22088	65345	0.65×

Panel (c) reveals the key efficiency difference between the adaptive solvers. The BDF method, equipped with the extended preconditioner (5.53) that incorporates the STT Jacobian, achieves time steps of $\mathcal{O}(1 \text{ ps})$, two orders of magnitude larger than the IMEX scheme. The latter remains constrained to $\mathcal{O}(0.01 \text{ ps})$ due to the explicit treatment of both precession and STT, which imposes accuracy limitations from the rapidly oscillating terms in f_E . This confirms that for STT-dominated dynamics, the efficiency gains from implicit exchange treatment alone are insufficient, and the spin-torque contributions must also be treated implicitly to unlock larger time steps.

Table 6.2 summarizes the performance metrics. The BDF method achieves an $82\times$ reduction in time steps compared to the fixed-step TPS reference while maintaining acceptable trajectory accuracy. Despite requiring more RHS evaluations per step (due to Newton iterations and the coupled drift-diffusion solve), the overall computational cost is significantly reduced. The IMEX scheme provides no improvement ($0.67\times$) over the baseline, as the stability advantage from implicit exchange is negated by the accuracy constraints imposed by explicit STT and precession treatment. Both adaptive solvers were configured with identical tolerances ($\text{rtol} = 10^{-2}$, $\text{atol} = 10^{-4}$).

Linear Solver Performance: The effectiveness of the preconditioner is measured by the number of GMRES iterations required to solve each Newton linear system and, more critically, by the achievable time step size. With the extended preconditioner (6.37) incorporating $\mathbf{J}_{\text{STT}}$, the average GMRES iteration count remains around 40 per Newton step, and the solver achieves time steps of $\mathcal{O}(1 \text{ ps})$. Without the STT Jacobian contribution, i.e., using only the exchange-damping preconditioner (6.38) with AMG, the solver behavior degrades catastrophically: GMRES iterations exceed 100 during the switching transient, additional Newton convergence failures occur, and the adaptive time step collapses to $\mathcal{O}(0.01 \text{ ps})$, negating any efficiency advantage of implicit integration. This demonstrates that for STT-dominated dynamics, incorporating the spin-torque Jacobian into the preconditioner is not merely an optimization, but a necessity for practical simulation.

In contrast, the IMEX scheme does not suffer from this preconditioner mismatch. Since the spin-transfer torque is assigned to the explicit partition f_E , the implicit Jacobian naturally contains only the exchange-damping contribution, and the SPD preconditioner (6.38) remains appropriate. However, this comes at the cost of explicit stability constraints on the STT terms, limiting the achievable time step to $\mathcal{O}(0.01 \text{ ps})$ regardless of preconditioner quality.

These results demonstrate that for strongly STT-driven dynamics, the fully implicit BDF approach with an appropriately constructed preconditioner offers substantial efficiency gains while maintaining physical fidelity, making it the preferred choice for large-scale or long-duration simulations. However, the TPS method remains attractive for its simplicity, robust constraint preservation, and straightforward implementation, it requires no preconditioner tuning. For the device simulations presented in the following sections, the TPS method is employed: its computational overhead is acceptable for the moderate problem sizes considered, and its unconditional stability and predictable behavior simplify the analysis of STT-driven switching phenomena.

6.9 Conclusion

This chapter has presented a comprehensive three-dimensional finite-element framework for computing spin-transfer torques and simulating magnetization dynamics in ultra-scaled MRAM cells. The approach is founded on the coupled spin and charge drift-diffusion formalism, which provides a unified description of both interface-induced Slonczewski torques and bulk Zhang–Li torques arising from magnetization textures within a single computational framework, eliminating the need for separate analytical expressions or ad hoc superposition of torque contributions.

A central achievement of this work is the extension of the drift-diffusion formalism to MTJs through a physically motivated treatment of the tunnel barrier. By modeling the TB as a poor conductor with magnetization-dependent conductivity and incorporating spin-current boundary conditions derived from NEGFs and circuit theory, the framework accurately reproduces the characteristic angular and voltage dependencies expected in MTJs. The implementation employs a unified interface mapping strategy that enables consistent treatment of both tunneling spin-currents at TB interfaces and interlayer exchange coupling at SAF interfaces, facilitating the simulation of complex multi-layered stacks representative of modern MRAM architectures.

Validation against analytical solutions confirms that the finite-element implementation with the tunneling boundary conditions yields quantitative agreement for the spin-accumulation profile, the damping-like and field-like torque components, and the sinusoidal angular dependence characteristic of MTJs. The investigation of elongated ferromagnetic structures with magnetization textures reveals that the Slonczewski and Zhang–Li torque contributions are not simply additive: the polarized tunneling spin-current modifies the bulk spin-accumulation in a manner that cannot be captured by superposing independently computed torque terms. This finding underscores the necessity of a self-consistent treatment when simulating STT-driven dynamics in devices where domain walls or non-uniform magnetization states may develop during switching.

The systematic evaluation of time integration strategies demonstrates that the fully implicit BDF method, equipped with a preconditioner incorporating the linearized spin-torque Jacobian, achieves time steps that are two orders of magnitude larger than those of explicit or IMEX schemes, offering substantial computational savings for large-scale simulations. Nevertheless, the TPS method remains attractive due to its algorithmic simplicity, robust preservation of the unit-sphere constraint, and freedom from the need to tune a preconditioner. For the device simulations presented in Chapter 7, the TPS method is employed: its computational overhead is acceptable for the moderate problem sizes considered, and its unconditional stability simplifies the analysis of STT-driven switching phenomena without introducing numerical artifacts that could obscure the underlying physics.

Advanced Switching Dynamics and Novel Device Structures

7.1 Introduction

This chapter applies the simulation framework developed in the preceding chapters to investigate switching dynamics and novel device structures in next-generation STT-MRAM devices. As device dimensions shrink toward the single-digit nanometer regime, phenomena that were negligible in larger cells become dominant, necessitating advanced structures that exploit interlayer exchange coupling and composite FL designs.

Section 7.3 addresses ultra-scaled composite FL switching at a 2.3 nm diameter, where elongated structures fail via domain wall formation and composite segmentation enables sequential reversal. Section 7.4 characterizes the deterministic back-hopping phenomenon in these composite structures and demonstrates its exploitation for four-state multi-level cell (MLC) operation. Section 7.5 introduces IEC as a systematic design parameter and examines its role across three stack families. Section 7.6 then advances to hybrid FLs with metallic spacers, SAF-enhanced structures, and double-spin torque magnetic tunnel junction (ds-MTJ) structures with dual reference layers, optimizing IEC for switching speed, write efficiency, and reference layer stability.

A unifying thread emerges across this progression: the initial studies (Sections 7.3 and 7.4) are conducted without IEC ($J_{\text{iec}} = 0$) to isolate the purely spin-transport-driven switching and back-hopping dynamics. IEC is then introduced as a systematic design parameter in Sections 7.5 and 7.6, where it evolves from a modifier of switching dynamics to an engineering parameter that must be quantitatively optimized. Section 7.7 consolidates the findings into cross-architecture design guidelines.

7.2 Computational Setup

The simulations in this chapter employ the coupled micromagnetic and spin-charge drift-diffusion framework developed in Chapter 6, integrated with the FE discretization described in Section 4.2. The spin-transfer torque density at each spatial point is computed using the drift-diffusion formalism (Section 6.1) [32], with tunneling boundary conditions at each MgO barrier derived from the NEGFs analysis in Section 6.3 [118].

Table 7.1: Standard simulation parameters for the switching studies presented in this chapter. Values are representative of CoFeB|MgO-based magnetic tunnel junctions and are sourced from the Boris Computational Spintronics material database and published experimental literature.

Parameter	Symbol	Value
Gilbert damping constant	α	0.015–0.02
Gyromagnetic ratio	γ	$1.76 \cdot 10^{11} \text{ rad s}^{-1} \text{ T}^{-1}$
Saturation magnetization (CoFeB)	M_S	$0.81 \cdot 10^6 \text{ A m}^{-1}$ to $1.2 \cdot 10^6 \text{ A m}^{-1}$
Exchange stiffness constant (CoFeB)	A	$1 \cdot 10^{-11} \text{ J m}^{-1}$ to $2 \cdot 10^{-11} \text{ J m}^{-1}$
Interface anisotropy	K_i	$1.53 \cdot 10^{-3} \text{ J m}^{-2}$
Current spin polarization	β_σ	0.52–0.7
Diffusion spin polarization	β_D	0.7–1.0
Electron diffusion coefficient (CoFeB)	$D_{e,\text{FM}}$	$1 \cdot 10^{-2} \text{ m}^2 \text{ s}^{-1}$ to $2 \cdot 10^{-2} \text{ m}^2 \text{ s}^{-1}$
Exchange length	λ_J	0.8 nm to 1 nm
Spin dephasing length	λ_φ	0.4 nm
Spin-flip length (CoFeB)	$\lambda_{sf,\text{CoFeB}}$	10 nm
Spin-flip length (Ta)	$\lambda_{sf,\text{Ta}}$	1.9 nm
Spin-flip length (Ru)	$\lambda_{sf,\text{Ru}}$	4 nm
Resistance (P state)	R_P	410 k Ω (ultra-scaled)
Resistance (AP state)	R_{AP}	750 k Ω (ultra-scaled)

Interlayer exchange coupling between non-adjacent magnetic layers is incorporated via the weak-form interface mapping described in Section 6.4. Magnetization dynamics are solved using the tangent-plane scheme (Section 5.3). Demagnetization fields are computed with the hybrid FE-BEM approach from Chapter 4.

The standard simulation parameters employed throughout this chapter are summarized in Table 7.1. Material set differences are detailed in Table 7.2.

The simulated resistance values yield tunnel magnetoresistance ratios consistent with experimental CoFeB|MgO devices: $\text{TMR} = (R_{AP} - R_P)/R_P = 83\%$ for the ultra-scaled composite FL ($R_P = 410 \text{ k}\Omega$, $R_{AP} = 750 \text{ k}\Omega$), and for the 70 nm SAF-enhanced stacks the input $\text{RA} = 6.5 \Omega \mu\text{m}^2$ with $\text{TMR} = 150\%$ lies within the experimentally reported range ($\text{RA} \approx 3.5 \Omega \mu\text{m}^2$ to $7 \Omega \mu\text{m}^2$, $\text{TMR} \approx 100\text{--}250\%$) for optimized MgO barriers at room-temperature [185, 186, 187]. The Gilbert damping constant $\alpha = 0.015\text{--}0.02$ is representative of sputtered CoFeB thin films [53]. In the macrospin limit, the critical switching current scales linearly with damping ($I_{c0} \propto \alpha$) [31], so the absolute switching times scale accordingly. However, the qualitative features identified in this chapter, namely the sequential switching order, the back-hopping threshold, and the IEC role progression, are governed by the torque hierarchy and coupling geometry rather than the damping magnitude.

A critical insight emerges from investigating multi-layer composite FL structures: the spin-transfer torques in a multi-segment stack are inherently non-additive. The coupled partial differential equation system governing spin transport exhibits essential non-linearity: the spin-accumulation profile in each magnetic layer is modified by the magnetization state of neighboring layers through spin-transport coupling across the

shared interfaces, creating strong feedback that cannot be decomposed into independent contributions. This non-additivity has direct practical implications: superposing single-layer simulation results will produce systematically incorrect switching times and potentially miss qualitative features such as the sequential switching mechanism [154, 188].

All simulations presented in this chapter are performed at zero temperature ($T = 0\text{ K}$). This deliberate choice isolates the deterministic switching dynamics from thermally activated processes, enabling distinct identification of torque-driven phenomena. In particular, the back-hopping mechanism characterized in Section 7.4 is identified as a deterministic consequence of the composite FL torque hierarchy precisely because stochastic thermal fluctuations are excluded from the simulations. The approximation is justified quantitatively: the characteristic switching times of 1 ns to 3 ns and back-hopping cycling periods of approximately 0.5 ns are many orders of magnitude shorter than the thermally activated reversal time $\tau_{\text{th}} = \tau_0 \exp(\Delta)$, which for $\Delta > 60$ exceeds 10^{17} s . Thermal fluctuations would add stochastic jitter to the trajectories without altering the qualitative switching sequence or the back-hopping mechanism.

7.3 Ultra-Scaled MRAM Cell Switching

The non-additivity of spin-transfer torques in multi-segment stacks demonstrated above raises a fundamental question: how do composite FL structures actually switch under these coupled conditions? This section addresses this by systematically simulating ultra-scaled MRAM cells [123, 182], in which the inter-segment spin-transfer torques through TB_2 govern the sequential magnetization reversal. All simulations in this section are performed without interlayer exchange coupling ($J_{\text{iec}} = 0$), isolating the purely spin-transport-driven dynamics.

As the cell diameter shrinks, the thermal stability factor $\Delta = \mathcal{E}_{\text{B}}/k_{\text{B}}T$ (requiring $\Delta > 60$ for 10-year retention, cf. Chapter 2) demands compensating increases in layer thickness or anisotropy, both facing practical limits. Experiments from Tohoku University [189, 190] demonstrated that elongating the FL perpendicular to the film plane, leveraging shape anisotropy to supplement interface-induced PMA, enables ultra-scaled cells at single-digit nanometer diameters. However, the elongated geometry fundamentally alters the switching dynamics, as the following analysis reveals.

7.3.1 Device Architecture and Geometry

Figure 7.1(a, b) illustrates conventional MRAM cells with a lateral diameter of approximately 40 nm: (a) a standard MTJ stack consisting of a 1 nm RL, a 1.7 nm FL, and a 1 nm MgO TB, and (b) a variant employing a NMS, a 2 nm metal film. Figure 7.1(c–e) show scaling concepts for ultra-scaled diameters: (c) an elongated FL with a total length of 15 nm to exploit perpendicular shape anisotropy, (d) a segmented FL with three MgO layers (TB_1 – TB_3) to increase the number of CoFeB/MgO interfaces and thus the interfacial PMA, and (e) a split FL separated by a NMS.




Figure 7.1: Three-dimensional renderings of (a)–(b) conventional large-diameter MTJ cells and (c)–(e) ultra-scaled 2.3 nm diameter pillar geometries with different FL configurations: elongated single FL, composite two-segment FL with middle TB, and composite FL with non-magnetic metallic spacer [182].

The investigated ultra-scaled MRAM cell (Figure 7.1) consists of a 2.3 nm diameter cylindrical pillar with a 5 nm CoFeB RL (magnetization fixed in $+\hat{x}$), a primary MgO barrier TB_1 (0.9 nm), a first FL segment FL_1 (5 nm), a middle MgO barrier TB_2 (0.9 nm) that boosts interfacial PMA and generates spin torque critical for sequential switching, a second FL segment FL_2 (5 nm), a capping MgO layer TB_3 (0.9 nm), and 50 nm metallic contacts.

In conventional perpendicular MTJs, the anisotropy is primarily determined by the CoFeB|MgO interface adjacent to the TB. A common strategy to enhance the effective perpendicular anisotropy is therefore to introduce additional CoFeB|MgO interfaces within the FL. For ultra-scaled junctions, however, interfacial PMA alone is insufficient to maintain thermal stability. Instead, the FL can be elongated along the pillar axis so that the shape anisotropy provides a positive contribution to the perpendicular easy axis once the magnetic thickness becomes comparable to or larger than the pillar diameter. This concept of perpendicular shape anisotropy has been experimentally demonstrated for single-digit nanometer MTJs and enables current-driven switching while preserving large thermal stability factors at very small diameters [189]. Further scaling down to a 2.3 nm diameter was achieved using a composite FL separated by MgO, which yields very high thermal stability and low switching voltage due to the combined action of interfacial anisotropy, magnetostatic coupling, and spin-transfer torque [190]. Recent experiments have extended this approach to single-nanometer MTJs with multi-interface FL designs [86].

The present device combines these approaches. The segmentation of the FL by MgO barriers creates multiple CoFeB|MgO interfaces and thus enhances the interfacial contribution to the perpendicular anisotropy. At the same time, the increased total FL length places the device in the regime where perpendicular shape anisotropy becomes substantial, which is essential for maintaining sufficient thermal stability at a pillar diameter of 2.3 nm. The intermediate MgO barrier TB₂ magnetically partitions the FL and enables inter-segment spin-transfer torque, resulting in a sequential switching mechanism. The top MgO layer TB₃ acts as an MgO cap that provides an additional CoFeB|MgO interface for anisotropy engineering [191, 192]. Since it is not part of the primary RL–FL tunneling conduction path, it does not significantly increase the resistance-area product of the junction. TB₃ also does not contribute significantly to STT, because the spin current polarized by the RL is attenuated by FL₁ and TB₂ before reaching FL₂, and any residual polarization at the FL₂|TB₃ interface is weak ($P_3 = 0.2$ in the simulations). The spin polarization at each MgO barrier (P_1, P_2, P_3) can be independently controlled through fabrication parameters: MgO thickness governs the coherent tunneling efficiency, and annealing conditions control CoFeB|MgO interface crystallinity [53, 69, 191]. This provides the experimental basis for treating P_1, P_2 , and P_3 as independent design parameters in the simulations.

The continuum micromagnetic approach remains valid at this scale because the exchange length $\lambda_{\text{ex}} = \sqrt{2A/(\mu_0 M_s^2)} \approx 5$ nm exceeds the 2.3 nm device diameter, so each FL segment behaves approximately as a macrospin.

7.3.2 Elongated FL: Domain Wall Failure

The most direct approach to thermal stability is to elongate the FL along the pillar axis, thereby supplementing interface-induced PMA with shape anisotropy. However, as Figure 7.2 shows for a 15 nm single FL under ± 2 V bias, the spatially non-uniform STT distribution initiates a domain wall that dissolves after approximately 3 ns without achieving switching, as the competing torques on either side of the wall cancel the net driving force.

This failure motivates the composite FL approach: dividing the elongated FL into shorter segments, separated by a thin spacer, suppresses domain-wall formation, enabling sequential switching.

7.3.3 Composite FL Solution

Splitting the FL into two 5 nm segments overcomes the domain wall barrier. Two separator variants were investigated: a 0.9 nm MgO TB and a 2 nm NMS.

With the MgO TB separator (Figure 7.3a), FL₁ reverses first, followed by FL₂, switching after approximately 2 ns. The NMS variant (Figure 7.3b) achieves switching approximately twice as fast (~ 1 ns), because the metallic spacer permits stronger spin current transmission between the FL segments compared to the insulating MgO barrier, resulting in a larger torque driving FL₂ reversal after FL₁ has switched. These sub-2 ns switching times at moderate bias are consistent with the experimentally demonstrated low-voltage switching in 2.3 nm diameter composite FL devices [190],




Figure 7.2: Magnetization dynamics in an ultra-scaled cell with a 15 nm elongated single FL (cf. Figure 7.1 (c)) at ± 2 V bias. (a), (b) Time evolution of magnetization components for AP to P and P to AP initial configurations showing failed switching. (c) Three-dimensional magnetization snapshot at 2 ns revealing the domain wall that prevents reversal. Figures based on [182].

and the predicted sequential reversal mechanism agrees qualitatively with the current-driven switching behavior reported for multi-interface ultra-scaled MTJs [86, 189]. The following subsections analyze the sequential switching mechanism in detail.

7.3.4 Switching Mechanisms in Composite FLs

The two FL segments switch sequentially rather than simultaneously, with profound implications for switching speed, intermediate-state exploitation, and failure modes. The following parametric study employs a lower bias of ± 1.5 V for systematic model comparison.

Figure 7.3 shows that both transitions exhibit a distinct plateau at $\langle m_x \rangle \approx 0$, corresponding to an intermediate configuration in which the two FL segments have opposite orientations.

Figure 7.4 visualizes the sequential switching process, confirming the intermediate state, where the two FL segments have opposite orientations. The torque hierarchy governing this sequential switching, why FL₁ switches first in AP $\rightarrow$ P and FL₂ first in P $\rightarrow$ AP, is analyzed quantitatively in Section 7.4.1.




Figure 7.3: Successful switching of composite FL structures at -2 V bias: (a) two 5 nm FL segments separated by a 0.9 nm MgO TB (cf. Figure 7.1 (d)), reaching switching after approximately 2 ns, (b) two 5 nm FL segments separated by a 2 nm non-magnetic metallic spacer (cf. Figure 7.1 (e)), achieving switching after approximately 1 ns. Figures based on [182].




Figure 7.4: Three-dimensional magnetization snapshots of the sequential switching process in ultra-scaled composite FL cells (cf. Figure 7.1 (d)), showing initial, intermediate (plateau), and final states for (a) AP $\rightarrow$ P and (b) P $\rightarrow$ AP transitions. The switching order is direction-dependent: FL₁ switches first in AP $\rightarrow$ P, while FL₂ switches first in P $\rightarrow$ AP, reflecting the asymmetric torque configurations at each interface [123].

7.3.5 FL Length Optimization

A systematic parameter study varied the FL segment length over 2, 3, 4, 5, and 7.5 nm at ± 1.5 V bias.

Figure 7.5 reveals several critical trends. The intermediate plateau duration at $\langle m_x \rangle \approx 0$ increases systematically with FL length: barely visible for 2 nm segments but extending for hundreds of picoseconds at 5 nm. Because the anisotropy barrier scales as $\mathcal{E}_B \propto K_{\text{eff}} \cdot l$ while the interfacial spin-transfer torque is approximately length-independent, the characteristic switching time scales linearly with segment length, $\tau_{\text{sw}} \propto l$.




Figure 7.5: Magnetization trajectories during (a) AP→P and (b) P→AP switching for FL segment lengths of 2, 3, 4, 5, and 7.5 nm with standard polarization ($P_1 = 0.62$, $P_2 = 0.57$, $P_3 = 0.2$) and Set A material parameters (Table 7.2). Critical design constraints emerge: segments shorter than approximately 5 nm exhibit back-hopping instability during P→AP switching, because the anisotropy barrier scales with segment length ($\mathcal{E}_B \propto K_{\text{eff}} \cdot l$) while the interfacial spin-transfer torque is approximately length-independent, so shorter segments satisfy the instability condition (7.1) more readily. Segments longer than approximately 6 nm fail to complete switching at moderate bias. The optimal range of 3 nm to 5 nm provides reliable switching in both directions. Based on Figure [193].

Minimum length constraint - back-hopping: The 3 nm to 4 nm FL segments exhibit back-hopping during P→AP switching, visible in Figure 7.5 (b) as a reversal after the initial transition. The back-hopping arises because the low anisotropy barrier of short segments allows the torque from the reversed FL₁ to drive FL₂ back after the initial switch. This deterministic back-hopping mechanism and its exploitation for multi-level cell operation are analyzed in detail in Section 7.4. Conversely, the 2 nm segments demonstrate a distinct failure mode: only the FL₂ layer undergoes switching, subsequently becoming locked in an intermediate state. This inability to complete the transition to the full AP configuration results in a permanent write error.

Maximum length constraint - incomplete switching: The 7.5 nm segments fail to complete P→AP switching at 1.5 V, remaining locked at the initial state because the stray fields generated by the composite free layer favor the parallel alignment, effectively stabilizing the configuration and making the P state more difficult to destabilize with the available STT. In contrast, the AP→P transition is still possible for the same 7.5 nm structure, although switching occurs only after approximately 8 ns, indicating significantly slower reversal dynamics. Since the switching time greatly exceeds the timescale relevant for this study, the trajectories shown in Figure 7.5 are truncated at 3 ns, and this case is not further investigated.

Optimal range: The FL segment length range of 4 nm to 5 nm provides the optimal balance for the ultra-scaled geometry:




Figure 7.6: Same configuration as Figure 7.5 but with Set B material parameters (Table 7.2), confirming robustness of the FL length constraints across material parameter sets. Based on Figure [194].




Figure 7.7: Magnetization trajectories for FL segment lengths of 2, 3, 4, and 5 nm with enhanced TB_2 polarization ($P_1 = 0.5$, $P_2 = 0.9$, $P_3 = 0.2$) and Set A material parameters. The enhanced TB_2 polarization modifies the back-hopping threshold compared to the standard case (Figure 7.5). Based on Figure [193].

- (i) Lengths below 2 nm to 3 nm risk back-hopping instability during $P \rightarrow AP$ switching
- (ii) Lengths above 5 nm may fail to complete switching at moderate bias voltages
- (iii) Intermediate lengths of 4 nm to 5 nm provide successful switching with adequate thermal stability. However, 4 nm segments still exhibit marginal back-hopping tendencies. This behavior can be circumvented by utilizing lower bias voltages or optimizing pulse duration to ensure the device settles into the desired state before instability triggers.




Figure 7.8: Same enhanced TB₂ polarization as Figure 7.7 but with Set B material parameters, demonstrating that the polarization-dependent switching behavior persists across material choices. Based on Figure [194].

These constraints are specific to the 2.3 nm diameter geometry studied. Still, the fundamental trade-offs back-hopping for too-short layers and incomplete switching for too-long layers remain applicable across device dimensions.

New physical insight. The three-dimensional drift-diffusion simulations reveal that composite FLs in ultra-scaled MRAM cells switch sequentially rather than coherently, with the switching order determined by the asymmetric torque configuration at each interface FL₁ first for AP→P, FL₂ first for P→AP. This sequential mechanism, inaccessible to macro-spin models, produces an observable plateau at $\langle m_x \rangle \approx 0$ whose duration scales with the FL segment length. The optimal segment length of 4 nm to 5 nm balances thermal stability ($\Delta > 60$) against the onset of back-hopping instability.

7.3.6 Material Parameter Robustness

To verify that back-hopping is a robust structural phenomenon rather than a material artifact, the simulations were repeated with an alternative parameter set (Set B) from the Boris Computational Spintronics database (Table 7.2): 32% lower M_S , doubled A , halved D_e , and 20% shorter λ_J , with identical K_i and α . Figure 7.7

Figure 7.6 and Figure 7.8 confirm that Set B reproduces the same directional asymmetry, FL-length dependence, and oscillation periods as Set A despite substantially different bulk transport parameters. This robustness follows from the instability condition (7.1): since K_i and the TB polarizations are identical in both sets, the onset threshold is preserved. This configuration is further validated by the dynamics shown in Figure 7.7 and Figure 7.8, which confirm that the middle barriers, referring to the tunnel barriers separating layers, dominate the switching outcome, establishing the TB polarization hierarchy as the primary design parameter for back-hopping control.

Table 7.2: Comparison of material parameter sets used to verify robustness of the back-hopping mechanism.

Parameter	Set A	Set B	Ratio
M_S [A m^{-1}]	1.2×10^6	0.81×10^6	0.68
A [J m^{-1}]	1×10^{-11}	2×10^{-11}	2.0
D_e [$\text{m}^2 \text{s}^{-1}$]	2×10^{-2}	1×10^{-2}	0.5
λ_J [nm]	1.0	0.8	0.8
K_i [J m^{-2}]	1.53×10^{-3}		1.0
α	0.015		1.0

7.3.7 Switching Characteristics in Asymmetric FLs

To further investigate deterministic back-hopping behavior, cells with asymmetric FL configurations are investigated. Figure 7.9 displays the magnetization trajectories for switching from AP→P and P→AP in which the segments of the composite FL have asymmetric lengths. In these asymmetric structures, either FL₁ is fixed at 5 nm while FL₂ is varied (2, 3, and 4 nm), or conversely, FL₂ is fixed at 5 nm and FL₁ is varied.

Figure 7.9 (a) demonstrates that successful switching from AP→P is always achieved regardless of the FL composition. However, Figure 7.9 (b) reveals that successful switching from P→AP is highly dependent on the layer geometry. Successful transitions are observed only for the symmetric configuration with a length of 5 nm. For asymmetrical configurations, the variations in FL₁ and FL₂ lengths result in distinct total magnetizations, leading to intermediate states where the cell fails to settle into the AP configuration. The differences in response times among these cells arise from their distinct uniaxial anisotropies, which scale with the layer's length. A shorter layer has a lower energy barrier separating the two magnetization states, thereby accelerating the response time but increasing susceptibility to instability.

To control this instability, Figure 7.10 reports switching simulations where the polarization of the middle TB (TB₂) is enhanced. This higher TB polarization yields a larger spin current, ensuring that the inter-segment torque acting between the two FL parts is stronger than the torques originating from the RL. In Figure 7.10 (a), all structures switch successfully during AP→P. Notably, the asymmetric structure with a 4 nm FL₁ and a 5 nm FL₂ exhibits a profoundly delayed switching initiation, remaining in a canted state ($\langle m_x \rangle \approx -0.75$) for over 3 ns. This delay arises from a torque deadlock: because FL₁ must initiate the AP→P transition, the forward torque from the RL is heavily countered by the opposing torque from the thicker 5 nm FL₂, creating a backward push amplified by the enhanced TB₂ polarization. Shorter segments (2 and 3 nm) avoid this deadlock because their reduced energy barriers are easily overcome by the STT. Furthermore, the symmetric 5 nm structure bypasses this trap entirely, due to its geometry, which prevents a larger layer from artificially pinning a smaller one, allowing it to transition smoothly and much faster (~ 2 ns).

Similarly, during the P→AP transition (Figure 7.10 (b)), the asymmetric structure with a 5 nm FL₁ and a 4 nm FL₂ demonstrates an analogous behavior. Because the torque hierarchy dictates that FL₂ must initiate this reversal, the 4 nm FL₂ presents




Figure 7.9: Magnetization trajectories with standard polarization ($P_1 = 0.62$, $P_2 = 0.57$, $P_3 = 0.2$) and asymmetric FL segment lengths for (a) AP→P and (b) P→AP switching across FL lengths of 2 nm to 5 nm. Based on Figure [188].




Figure 7.10: Magnetization trajectories with enhanced TB₂ polarization ($P_1 = 0.5$, $P_2 = 0.9$, $P_3 = 0.2$) and asymmetric FL segment lengths for (a) AP→P and (b) P→AP switching across FL lengths of 2 nm to 5 nm. The enhanced TB₂ polarization intensifies the cyclic transitions compared to the standard polarization case (Figure 7.9), confirming deterministic back-hopping. Based on Figure [188].

a substantial energy barrier compared to thinner segments. The STT struggles to destabilize this thick initiating layer against the influence of the larger 5 nm FL₁, significantly postponing the onset and preventing the reversal process. Consequently, the magnetization remains pinned near $\langle m_x \rangle \approx 0.75$ for the entire time. This deterministic process repeats indefinitely as long as the bias is maintained, confirming that the middle polarization can tune or deliberately trigger cyclic instability.

7.4 Back-Hopping and Multi-Level Cell Operation

The FL length optimization revealed that segments shorter than 3 nm exhibit back-hopping instability during $P \rightarrow AP$ switching. Especially as shown in Figure 7.7 with a higher polarization in the middle TB, initiating a cyclic instability rather than settling into the switched state. Rather than a failure mode, this instability follows deterministic dynamics governed by the inter-segment torque hierarchy and can be exploited for multi-level cell operation [188, 193]. As in the preceding section, all simulations use $J_{iec} = 0$.

Unlike thermally driven back-hopping in conventional STT-MRAM [195, 196], the back-hopping in composite FL structures is a deterministic consequence of the inter-segment torque hierarchy. When the middle barrier (TB₂) polarization exceeds that of the outer barriers, the torque from FL₁ on FL₂ through TB₂ destabilizes whichever state the composite FL achieves. After both segments reach the target AP state during $P \rightarrow AP$ switching, this torque can exceed the interface anisotropy restoring force:

$$|\mathbf{T}_{2,FL_1}| > \frac{2K_i}{\mu_0 M_S} \quad (7.1)$$

When this condition is satisfied, FL₂ reverses back toward the initial orientation, which in turn reverses the torque it exerts on FL₁, triggering a cascade that returns the entire composite FL to its initial configuration. As long as the bias is maintained, the cycle repeats indefinitely. This cyclic switching of the layers is also referred to as spin-torque nano-oscillators, spiking dynamics, or windmill structures [197, 198].

The mechanism produces a pronounced directional asymmetry: during $AP \rightarrow P$, the TB₁ and TB₂ torques act cooperatively, yielding clean transitions (Figure 7.5 (a), Figure 7.7 (a)). At the same time, during $P \rightarrow AP$ the reconfigured TB₂ spin-accumulation opposes switching, producing cyclic reversals (Figure 7.5 (b), Figure 7.7 (b)). The cycling frequency increases with decreasing segment length, and enhanced TB₂ polarization ($P_2 = 0.9$, Figure 7.7) intensifies the amplitude relative to standard polarization ($P_2 = 0.57$, Figure 7.5), confirming TB₂ polarization as the primary control parameter.

7.4.1 Torque Analysis

To quantify the sequential switching and back-hopping connection, the torques on each FL segment are examined with the RL fixed at $\mathbf{m}_{RL} = +\hat{x}$. The torque on FL₁ has two contributions:

$$\mathbf{T}_1 = \mathbf{T}_{1,RL} + \mathbf{T}_{1,FL_2} \quad (7.2)$$

where $\mathbf{T}_{1,RL}$ is the RL-polarized contribution through TB₁ and $\mathbf{T}_{1,FL_2}$ is the FL₂-polarized contribution through TB₂. Since the direct RL torque on FL₂ is attenuated by two barriers:

$$\mathbf{T}_2 \approx \mathbf{T}_{2,FL_1} \quad (7.3)$$

7.4.1.1 AP→P: Cooperative Torques

In the initial AP configuration ($\mathbf{m}_1 = \mathbf{m}_2 = -\hat{x}$), spin-polarized electrons tunnel through TB₁ and exert a damping-like torque rotating $\mathbf{m}_1$ toward $+\hat{x}$. Simultaneously, the secondary torque from FL₂ through TB₂ is:

$$\mathbf{T}_{1,FL_2} \propto \mathbf{m}_1 \times (\mathbf{m}_1 \times \mathbf{m}_2) = \mathbf{m}_1 \times (\mathbf{m}_1 \times (-\hat{x})) \quad (7.4)$$

Both torques act cooperatively, explaining why FL₁ switches first [182]. The torque on FL₂ from FL₁:

$$\mathbf{T}_{2,FL_1} \propto \mathbf{m}_2 \times (\mathbf{m}_2 \times \mathbf{m}_1) \quad (7.5)$$

is initially zero but after FL₁ reverses to $+\hat{x}$, changes sign:

$$\mathbf{T}_{2,FL_1} \propto \mathbf{m}_2 \times (\mathbf{m}_2 \times (+\hat{x})) \quad (7.6)$$

This drives FL₂ to switch, but with only the FL₁ torque available, the reversal is slower, producing the extended plateau in Figure 7.3 (a).




Figure 7.11: Spin-transfer torque heatmap showing $T_{S,x}$ (top) and $T_{S,z}$ (bottom) as a function of stack position and magnetization angle for standard polarization ($P_1 = 0.62$, $P_2 = 0.57$, $P_3 = 0.2$). Left column: FL₂ rotation, right column: FL₁ rotation. Contour lines indicate torque magnitude within ferromagnetic layers.

7.4.1.2 P→AP: Competing Torques and Back-Hopping Origin

In the initial P configuration ($\mathbf{m}_1 = \mathbf{m}_2 = +\hat{x}$), the reversed bias drives the RL torque on FL₁ toward $-\hat{x}$:

$$\mathbf{T}_{1,RL} \propto -\mathbf{m}_1 \times (\mathbf{m}_1 \times (+\hat{x})) \quad (7.7)$$

However, the torque from FL₂ through TB₂ opposes this reversal:

$$\mathbf{T}_{1,FL_2} \propto \mathbf{m}_1 \times (\mathbf{m}_1 \times (+\hat{x})) \quad (7.8)$$

These torques partially cancel on FL₁, while the torque on FL₂ from FL₁, given by:

$$\mathbf{T}_{2,FL_1} \propto -\mathbf{m}_2 \times (\mathbf{m}_2 \times (+\hat{x})) \quad (7.9)$$

drives FL₂ toward $-\hat{x}$ unopposed. After FL₂ reverses, both torques on FL₁ become cooperative:

$$\mathbf{T}_{1,RL} + \mathbf{T}_{1,FL_2} \propto -\mathbf{m}_1 \times (\mathbf{m}_1 \times (+\hat{x})) + \mathbf{m}_1 \times (\mathbf{m}_1 \times (-\hat{x})) \quad (7.10)$$

Both terms drive FL₁ toward $-\hat{x}$, producing a slow first phase (opposing torques) followed by a fast second phase (cooperative torques). After both segments reach the AP state ($\mathbf{m}_1 = \mathbf{m}_2 = -\hat{x}$), the torque landscape reconfigures. Substituting $\mathbf{m}_1 = -\hat{x}$ into (7.9):

$$\mathbf{T}_{2,FL_1} \propto -\mathbf{m}_2 \times (\mathbf{m}_2 \times (-\hat{x})) = \mathbf{m}_2 \times (\mathbf{m}_2 \times (+\hat{x})) \quad (7.11)$$

This torque drives $\mathbf{m}_2$ toward $+\hat{x}$ (the initial P orientation) rather than stabilizing the AP state, triggering back-hopping when (7.1) is satisfied.

After FL₂ backhops to $+\hat{x}$ (while FL₁ remains at $-\hat{x}$), the torque on FL₁ from FL₂ through TB₂ becomes:

$$\mathbf{T}_{1,FL_2} \propto \mathbf{m}_1 \times (\mathbf{m}_1 \times (+\hat{x})) \quad (7.12)$$

driving $\mathbf{m}_1$ toward $+\hat{x}$, while the RL torque through TB₁ opposes:

$$\mathbf{T}_{1,RL} \propto -\mathbf{m}_1 \times (\mathbf{m}_1 \times (+\hat{x})) \quad (7.13)$$

The net torque on FL₁, with magnitudes scaling as P_2 and P_1 :

$$\mathbf{T}_1 \propto (P_2 - P_1) \mathbf{m}_1 \times (\mathbf{m}_1 \times (+\hat{x})) \quad (7.14)$$

For $P_2 = 0.9 > P_1 = 0.5$, the FL₂ torque dominates, driving FL₁ back to $+\hat{x}$ and completing the cycle. The time-resolved data (Figure 7.15) confirms a 10× TB₂ torque surge at back-hopping onset.

The heatmaps are constructed by systematically sampling the torque landscape of the composite FL. For each configuration, one FL segment is held at a fixed orientation. At the same time, the other is rotated in the xz -plane (about $\hat{y}$) in 5° increments from the parallel (0°) to the anti-parallel (180°) alignment. Specifically, for the P→AP-relevant landscape: FL₁ is fixed in the P orientation ($+\hat{x}$) while FL₂ is rotated from P (0°) toward AP (180°). Subsequently, FL₂ is fixed in the P orientation,




Figure 7.12: Spin-transfer torque heatmap for enhanced TB₂ polarization ($P_1 = 0.5$, $P_2 = 0.9$, $P_3 = 0.2$). Same layout as Figure 7.11. The enhanced polarization introduces a pronounced torque imbalance: $T_{S,x}$ on FL₁ is significantly enhanced by the high $P_2 = 0.9$, while the torque on FL₂ is reduced by the lower $P_3 = 0.2$.

and FL₁ is rotated from AP (180°) toward P (0°). At each angular step, the coupled spin-transport equations are solved self-consistently to obtain the full spatial torque distribution. This procedure maps the torque experienced by each segment as a function of the other's orientation, capturing the non-additive inter-segment coupling.

The resulting heatmaps quantify the torque imbalance: in the standard case (Figure 7.11), $T_{S,x}$ peaks at comparable magnitudes ($\sim -20 \times 10^{15} \text{ A m}^{-1} \text{ s}^{-1}$) for both FL segments, while the enhanced configuration (Figure 7.12) doubles this to $\sim -40 \times 10^{15} \text{ A m}^{-1} \text{ s}^{-1}$ for FL₁, reflecting $P_2/P_3 = 0.9/0.2$. The perpendicular torque $T_{S,z}$ intensifies at TB₂ with a sign change near 90° , creating the destabilizing field responsible for back-hopping.

The three-dimensional surfaces (Figure 7.13, Figure 7.14) confirm the $2\times$ amplification across both torque components, with the $T_{S,x}$ angular profile exhibiting a clear $\sin(\theta)$ dependence characteristic of the drift-diffusion mechanism, ruling out numerical artifacts.

These spatial torque data provide an explanation for the directional asymmetry: during P $\rightarrow$ AP, the $T_{S,z}$ sign change near 90° at the TB₂ interface (Figure 7.12) constitutes the destabilizing field that, when the instability condition (7.1) is satisfied,




Figure 7.13: Three-dimensional surface representation of $T_{S,x}$ (a, b) and $T_{S,z}$ (c, d) for standard polarization ($P_1 = 0.62$, $P_2 = 0.57$, $P_3 = 0.2$). Panels (a, c): FL_2 rotation, (b, d): FL_1 rotation. The in-plane torque $T_{S,x}$ is negative throughout the angular range with a minimum near 90° . A Gaussian filter suppresses numerical noise from the FE mesh.

triggers the back-hopping cascade established in Section 7.4. During AP $\rightarrow$ P, both T_{B1} and T_{B2} torques cooperate to stabilize the final P state, and the torques vanish smoothly at $\theta \rightarrow 0^\circ$ without sign reversal. The slower initial P $\rightarrow$ AP phase arises because the flux-closure stray field from both FL segments at $+\hat{x}$ opposes the initial torque, whereas AP $\rightarrow$ P benefits from stray field assistance. This framework also explains the voltage-polarity asymmetry (Section 7.4.5): the cooperative torque configuration renders the P state immune to wrong-polarity switching, whereas the AP state becomes unstable above 3.0 V.

The T_{B1} torque (initial magnitude $\sim 5 \times 10^{15} \text{ A m}^{-1} \text{ s}^{-1}$) initiates the first segment reversal in each direction. During AP $\rightarrow$ P (Figure 7.15 (a)), the reconfigured T_{B2} torque on FL_2 reaches $T_{S,z} \sim +4 \times 10^{15} \text{ A m}^{-1} \text{ s}^{-1}$, cooperating with T_{B1} for a smooth transition. During P $\rightarrow$ AP (Figure 7.15 (b)), the T_{B2} torque surges to $T_{S,z} \sim +40 \times 10^{15} \text{ A m}^{-1} \text{ s}^{-1}$ after both segments reach the AP state, a $10\times$ asymmetry that quantitatively explains why only P $\rightarrow$ AP satisfies the instability condition (7.1).




Figure 7.14: Three-dimensional torque surfaces for enhanced TB₂ polarization ($P_1 = 0.5$, $P_2 = 0.9$, $P_3 = 0.2$). Compared to the standard case (Figure 7.13), $T_{S,x}$ on FL₁ (b) reaches approximately twice the magnitude of FL₂ (a), reflecting the $P_2/P_3 = 0.9/0.2$ ratio. This imbalance drives sequential switching and back-hopping.

The field-like torque $T_{S,x}$ remains comparatively small, modifying the precession trajectory without altering the switching sequence [200].

The combined trajectory and torque data confirm three key properties: (1) the oscillation period scales with FL segment length, establishing back-hopping as deterministic, (2) shorter segments (2 nm to 3 nm) exhibit more pronounced cycling due to lower anisotropy barriers, and (3) the $2\times$ torque amplification from $P_2 = 0.57$ to 0.9 (Figure 7.11–Figure 7.14) provides a quantitative design knob for back-hopping control.

7.4.2 Cyclic Switching Visualization

Figure 7.16 visualizes the complete P $\rightarrow$ AP back-hopping cycle through three-dimensional magnetization snapshots. The five stages trace the trajectory through all four composite FL states: (i) initial P ($\rightarrow \rightarrow \rightarrow$, State 1), (ii) FL₂ reversal ($\rightarrow \rightarrow \leftarrow$, State 2), (iii) cooperative FL₁ switching, full AP ($\rightarrow \leftarrow \leftarrow$, State 3), (iv) FL₂ back-hopping




Figure 7.15: Decomposition of damping-like and field-like spin-transfer torques on FL₁ and FL₂ during (a) AP→P and (b) P→AP switching. The torque hierarchy between TB₁ and TB₂ governs the sequential switching order and determines back-hopping onset [199].

(7.11) ($\rightarrow \leftarrow \rightarrow$, State 4), and (v) FL₁ return (7.14) to State 1. The critical transition (iv)→(v) occurs when $T_{S,z}$ at the TB₂ interface increases by approximately one order of magnitude (Figure 7.12), crossing the instability threshold (7.1).

7.4.3 Multi-Level Cell Realization

The deterministic cyclic switching enables multi-level cell operation: a composite FL with controlled back-hopping supports four distinct magnetization states [193], summarized in Table 7.3.

Table 7.3: Multi-level cell state definitions and characteristics.

State	FL ₁ vs. RL	FL ₂ vs. RL	Overall Resistance	Binary Code
1	Parallel	Parallel	Lowest	00
2	Parallel	Anti-parallel	Low-Medium	01
3	Anti-parallel	Anti-parallel	Highest	11
4	Anti-parallel	Parallel	Medium-High	10



Figure 7.16: Cyclic switching visualization during P→AP transition: three-dimensional magnetization snapshots at five characteristic stages (i)–(vi) of the FL reversal cycle, colored by the m_x magnetization component. Based on [199].

The cyclic P→AP trajectory traverses all four states ($1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 1 \rightarrow \dots$), visiting four unique resistance levels per cycle and validating four-level MLC operation. By controlling the write pulse duration, any state can be accessed and stabilized when the bias is removed.

Figure 7.17 confirms $>20\%$ resistance separation between adjacent states, stable retention after pulse removal, and reproducibility over multiple cycles. Pulse duration precision of approximately 100 ps is required.

7.4.4 FL Thickness Variation

Since FL thickness controls the anisotropy barrier and back-hopping onset, the pulse duration dependence is examined for the 3 nm symmetric structure to establish the MLC operational window.

Figure 7.18 maps the pulse-duration operational window for the 3 nm symmetric structure. Short pulses ($t_{\text{pulse}} \lesssim 0.8$ ns, red curves) return the magnetization to the initial state. Intermediate durations (1 ns to 2 ns) access distinct partial-switching configurations, and the longest pulses ($t_{\text{pulse}} > 2.5$ ns, dark blue) drive a full back-




Figure 7.17: Resistance (purple, left axis) and average magnetization m_x (orange dashed, right axis) as a function of time for four pulse durations: (a) $t_{\text{pulse}} = 1.0 \text{ ns}$, (b) 1.3 ns , (c) 2.0 ns , and (d) 3.0 ns . The vertical dotted line marks pulse termination t_{off} . The four resistance levels are clearly distinguishable with $>20\%$ separation between adjacent states. Based on Figure [194].

hopping cycle, confirming the cyclic trajectory of Section 7.4.2. The FL length dependence of the stable-state window was established in Section 7.3.5.

7.4.5 Voltage Polarity Robustness

Figure 7.19 shows the magnetization response under reversed voltage polarity (opposite to the polarity required for switching) for both initial states.

The P state is immune to reversed-polarity torques across the entire voltage range: the cooperative torque configuration (both TB_1 and TB_2 torques stabilize the parallel alignment, cf. Section 7.4.1) prevents destabilization regardless of voltage magnitude. In contrast, the AP state, characterized by competing torques between TB_1 and TB_2 , becomes unstable above a voltage threshold where the net torque overcomes the anisotropy barrier and initiates cyclic back-hopping. This asymmetry follows directly from the torque hierarchy established in Section 7.4.1 and confirms back-hopping as a deterministic torque-driven phenomenon [200].

New physical insight. Back-hopping in composite FL structures is a deterministic, torque-hierarchy-driven phenomenon robust across material parameter sets, not the




Figure 7.18: Average magnetization $\langle m_x \rangle$ versus time for P→AP switching at 1.5 V bias with write pulse durations ranging from 0 to 3.5 ns (color-coded from red for short pulses to blue for long pulses). By terminating the write pulse at different times, distinct final magnetization configurations are accessed: short pulses ($t_{\text{pulse}} \lesssim 0.8$ ns) return the composite FL to its initial state, intermediate pulses stabilize partial-switching configurations, and the longest pulses ($t_{\text{pulse}} > 2.5$ ns) drive a complete back-hopping cycle. Based on [199].




Figure 7.19: Magnetization response under reversed voltage polarity (opposite to the polarity required for switching): (a) AP state with positive voltage, stable at $m_x \approx -1$ for $V \leq 2.5$ V, but cyclic oscillations between $m_x = -1$ and $+1$ emerge at $V \geq 3.0$ V with increasing frequency, (b) P state with negative voltage, stable at $m_x = +1$ across the entire range (2.0 V to 4.5 V). The asymmetry arises from the cooperative torque configuration of the P state: both TB_1 and TB_2 torques act to stabilize the parallel alignment (see Section 7.4.1), whereas the AP state suffers from competing torques that enable destabilization at sufficient voltage.

stochastic thermal effect assumed in prior literature [188,195,201]. The cyclic dynamics are controllable through TB₂ polarization and FL segment length, enabling four-state MLC operation with distinct resistance levels addressable via pulse duration modulation at 100 ps resolution [193,194], transforming a write error into a 4× storage density increase.

7.5 Interlayer Exchange Coupling Effects

The preceding analysis established that inter-segment spin-transfer torques govern the sequential and cyclic switching dynamics in composite FL structures and can be harnessed for MLC operation, demonstrating that coupling between magnetic layers is a resource rather than a liability. This section systematically introduces IEC as a design parameter by sweeping the coupling strength across different stack structures, establishing the distinct roles of IEC in each configuration and motivating the advanced structures investigated in Section 7.6.

7.5.1 Interlayer Exchange Coupling as Control Parameter

The back-hopping analysis in Section 7.4 revealed that the inter-segment spin-transfer torque hierarchy governs the cyclic dynamics of composite FL structures. A natural question is whether IEC, which adds an effective field coupling between magnetic layers, can modify or control these dynamics. This motivates a systematic investigation of IEC as a design parameter. The IEC adds an effective field contribution to the LLG equation:

$$\mathbf{H}_{\text{IEC}} = \frac{J_{\text{iec}}}{M_S t_{\text{FL}}} \mathbf{m}_{\text{adjacent}} \quad (7.15)$$

where J_{iec} is the coupling strength (positive for FM, negative for AFM alignment), t_{FL} is the FL thickness, and $\mathbf{m}_{\text{adjacent}}$ is the magnetization of the adjacent layer, as derived in Section 6.4 [202], this effective field can either assist or oppose the switching torques, depending on the sign and magnitude of the coupling.

The first sweep of $J_{\text{iec},1}$ and $J_{\text{iec},2}$ was conducted independently through negative (AFM), zero, and positive (FM) values for the composite FL with MgO TB separator (Figure 7.1 (d)). The $\pm 2 \text{ mJ m}^{-2}$ sweep covers experimentally reported coupling magnitudes for CoFeB|MgO interfaces [204,205] and metallic spacers [206,207]. Both FM and AFM coupling have been observed, contingent upon spacer thickness and crystalline quality [195]. Figure 7.20 displays the AP→P switching dynamics during this sweep. Specifically, AFM coupling at the FL₁–FL₂ interface increases the energy barrier by opposing ferromagnetic alignment, thereby delaying switching; this occurs because the torque must first overcome this added barrier. In contrast, zero coupling leads to the baseline sequential switching of Section 7.3, as there is neither assistance nor opposition. FM coupling facilitates switching by inducing cooperative alignment between FL₁ and FL₂; however, this assistance can suppress the intermediate plateau, consequently eliminating the multi-level cell capability described in Section 7.4.




Figure 7.20: Influence of IEC strengths $J_{\text{iec},1}$ and $J_{\text{iec},2}$ on the AP→P switching dynamics in the ultra-scaled composite FL with MgO TB separator (Figure 7.1 (d)). The IEC is swept at the FL₁–FL₂ interface through the middle MgO barrier: (a) negative $J_{\text{iec},1}$ (AFM), (b) $J_{\text{iec},1} = 0$, and (c) positive $J_{\text{iec},1}$ (FM), each producing distinct switching dynamics. Based on Figures [203].

Next, the metallic NMS separator was analyzed (Figure 7.1 (e)). Both AP→P and P→AP transitions were included (Figure 7.21). The comparison shows a clear asymmetry between the two separator structures. The NMS variant transmits spin current more efficiently between the FL segments, making switching dynamics more sensitive to J_{iec} variations.

Under AFM coupling, the NMS-separated structure exhibits stronger suppression of the second-segment reversal, while under FM coupling, it shows faster synchronized switching than the TB variant. The P→AP direction is consistently more sensitive to IEC than AP→P, confirming the directional asymmetry established by the torque hierarchy analysis in Section 7.4.1. When the bias is turned off at 2.5 ns, the system relaxes to the desired final state in the AP→P direction. In contrast, certain P→AP configurations settle into an intermediate anti-parallel FL₁–FL₂ arrangement rather than a complete AP state, confirming that properly tuned AFM IEC can stabilize the final magnetic state and suppress back-transitions. An alternative mitigation strategy is to engineer structural asymmetry in the MTJ stack: Manchon *et al.* [209] showed that breaking inversion symmetry modifies the bias dependence of the spin torque, and Oh *et al.* [210] demonstrated experimentally that asymmetries in material composition expand the stability region of the parallel configuration and suppress back-hopping.




Figure 7.21: Systematic IEC sweep for the two ultra-scaled composite FL structures from [208]: (a–c) AP→P and (d–f) P→AP transitions for the composite FL with MgO TB separator (Figure 7.1 (d)), and (g–i) the same analysis for the composite FL with NMS separator (Figure 7.1 (e)). In each row, the left panel applies negative $J_{\text{iec},1}$ (AFM), the middle panel sets $J_{\text{iec},1} = 0$, and the right panel applies positive $J_{\text{iec},1}$ (FM). The coupling strength range reflects experimentally reported values for CoFeB|MgO interfaces [204, 205] and metallic spacers [206, 207]. Based on Figures [208].

These initial sweeps on ultra-scaled structures established IEC as a tunable design parameter. The follow-up study [116] extended this investigation to three distinct stack families spanning different IEC roles, from back-hopping suppression to reliability threshold control, as detailed in the following subsection.

7.5.2 Systematic IEC Sweep Across Three Stack Configurations

Stack A (ultra-scaled composite FL, diameter 2.3 nm Figure 7.22 (a)) features IEC between FL₁ and FL₂ through an MgO TB. A weak FM coupling of $J_{\text{iec}} \approx 0.01 \text{ mJ m}^{-2}$ [116, 205], consistent with the weak FM coupling observed experimentally in

CoFeB|MgO|CoFeB structures [211], is sufficient to suppress back-hopping while preserving fast switching: the FM coupling prevents the field-like torque from reversing the FL₂ magnetization, while the coupled FL₁–FL₂ system initiates reversal more uniformly. Figure 7.23(a) illustrates this contrast: with $J_{\text{iec}} = 0.01 \text{ mJ m}^{-2}$, the averaged FL magnetization $\langle m_x \rangle$ switches cleanly from +1 to –1 within $\sim 1 \text{ ns}$, whereas the uncoupled case ($J_{\text{iec}} = 0$) exhibits back-hopping with the magnetization reverting toward




Figure 7.22: Layer stacks of the three configurations studied in [116]: (a) Stack A (ultra-scaled composite FL, same architecture as Figure 7.1 (d)), (b) Stack B (SAF without polarization layer (PL)), and (c) Stack C (SAF with PL). IEC acts at different interfaces in each stack: FL_1 – FL_2 coupling through TB_2 in Stack A, HL–RL coupling through Ru in Stack B, and both HL–RL (AFM) and RL–PL (FM) coupling in Stack C.

$\langle m_x \rangle \approx +0.2$. The torque distribution at the onset of back-hopping (Figure 7.24(a)) reveals that the damping-like component $T_{S,z}$ peaks at each MgO interface, with the torque contributions from FL_1 driving the reversal of FL_2 by overcoming the interface anisotropy. Note that J_{iec} is highly sensitive to the MgO barrier thickness and the quality of the CoFeB|MgO interfaces. Consequently, maintaining this specific coupling is experimentally demanding, and subtle deviations in the fabrication process, particularly regarding interface morphology and barrier thickness, can result in an insufficient coupling strength, thereby reintroducing back-hopping failures during operation.

Stack B (SAF without PL, diameter 70 nm, Figure 7.22 (b)), a diameter widely used in experimental characterization and demonstration of perpendicular MTJs [185], employs AFM coupling between the CoPt hard layer (HL) and the CoFeB RL through a Ru NMS [212]. At the nominal coupling strength of $J_{iec} = -1.32 \text{ mJ m}^{-2}$ [207], the strong AFM coupling prevents RL reversal, maintaining stable AP alignment between HL and RL. At a reduced coupling of -0.5 mJ m^{-2} , however, the torques from the FL overcome the IEC, allowing domain wall formation in both FL and RL and triggering back-and-forth switching. Figure 7.23(b) captures this behavior: the strong-coupling curve reaches $\langle m_x \rangle = 0$, confirming stable AP alignment between RL and FL, while the weak-coupling curve oscillates as domain walls form in both the FL and the RL. The torque distribution (Figure 7.24(b)) shows that the HL generates a large damping-like torque $T_{S,z}$ at the Ru interface that opposes RL reversal. In contrast, the FL exerts a competing torque through the MgO TB that drives the RL toward P alignment. When the AFM coupling is sufficiently strong, the IEC effective field reinforces the HL torque and stabilizes the RL. The three-dimensional




Figure 7.23: Switching trajectories during P→AP transition for the three stack configurations: (a) Stack A (ultra-scaled composite FL), (b) Stack B (SAF without PL), and (c) Stack C (SAF with PL). Weak FM coupling suppresses back-hopping in Stack A, strong AFM coupling maintains RL stability in Stack B, and PL reversal in Stack C depends on the RL–PL coupling strength [116].




Figure 7.24: Spin-torque distribution at the moment of back-hopping onset for (a) Stack A, (b) Stack B, and (c) Stack C. The field-like component $T_{S,x}$ and damping-like component $T_{S,z}$ reveal distinct torque patterns at each ferromagnetic interface, with IEC modifying both the magnitude and spatial distribution of the torques [116].

magnetization snapshots in Figure 7.25 visualize this contrast: under weak coupling ((a)–(f)), domain walls nucleate in the RL by ~ 0.4 ns and its magnetization reverses almost completely by ~ 1 ns, whereas under strong coupling ((g)–(l)) the RL remains uniformly magnetized throughout the switching process.

Stack C (SAF with PL, diameter 70 nm, Figure 7.22 (c)) adds a CoFeB PL coupled to the RL through a Ta NMS, with AFM coupling between HL and RL through Ru ($J_{\text{iec}} = -1.5 \text{ mJ m}^{-2}$ [207]). The FM coupling between RL and PL controls PL stability: at the originally proposed coupling of 0.8 mJ m^{-2} [213], back-hopping is absent, whereas a revised weaker estimate of 0.21 mJ m^{-2} [195, 214] allows the torques from the FL to overcome the IEC and reverse the PL magnetization. As shown in Figure 7.23(c), the strong-coupling case produces a clean transition to $\langle m_x \rangle \approx 0$, confirming that only the FL switches while the PL remains fixed. Under weak coupling, the trajectory overshoots to $\langle m_x \rangle \approx -0.5$ as the PL reverses before partially recovering. The torque analysis (Figure 7.24(c)) reveals that the FL, near complete reversal, exerts a parasitic torque on the PL through the Ta NMS. The coupling strength determines




Figure 7.25: Three-dimensional magnetization states during P→AP switching in Stack B (SAF without PL), contrasting weak ($J_{\text{iee}} = -0.5 \text{ mJ m}^{-2}$) and strong ($J_{\text{iee}} = -1.32 \text{ mJ m}^{-2}$) AFM HL–RL coupling. (a)–(c) weak coupling, full view and (d)–(f) weak coupling, FL clipped to reveal RL, (g)–(i) strong coupling, full view and (j)–(l) strong coupling, FL clipped to reveal RL. Columns show the initial state, mid-switching ($\sim 0.4 \text{ ns}$), and final state ($\sim 1 \text{ ns}$). With strong coupling, the RL remains uniformly oriented throughout switching, with weak coupling, the RL develops domain walls and its magnetization reverses almost completely [215].

whether the IEC effective field can counteract this torque. Figure 7.26 provides three-dimensional confirmation: under weak coupling ((a)–(f)), the PL magnetization gradually reverses from its initial orientation to nearly complete inversion by $\sim 1 \text{ ns}$, while under strong coupling ((g)–(l)) the PL retains its original orientation throughout.

The three-stack comparison demonstrates that IEC plays distinct roles depending on architecture: modifier of sequential switching dynamics (Stack A), RL stability guarantee (Stack B), and PL alignment control (Stack C). These roles motivate the dedicated investigations of hybrid FL structures and ds-MTJ speed enhancement in the following section.




Figure 7.26: Three-dimensional magnetization states during P→AP switching in Stack C (SAF with PL), contrasting weak ($J_{\text{iec}} = 0.2 \text{ mJ m}^{-2}$) and strong ($J_{\text{iec}} = 0.8 \text{ mJ m}^{-2}$) FM RL–PL coupling. (a)–(c) weak coupling, full view and (d)–(f) weak coupling, FL clipped to reveal PL, (g)–(i) strong coupling, full view and (j)–(l) strong coupling, FL clipped to reveal PL. Columns show the initial state, mid-switching ($\sim 0.4 \text{ ns}$), and final state ($\sim 1 \text{ ns}$). With strong coupling, the PL remains stable, with weak coupling, the PL reverses almost completely due to parasitic spin-torques overcoming the IEC [215].

Table 7.4: Material and geometry parameters for the three MDPI 2024 stack configurations [116] and the MME 2025 SAF-enhanced stacks [216].

Common parameters are listed in Table 7.1.

Parameter	Stack A (MDPI)	Stack B (MDPI)	Stack C (MDPI)	MME Stacks D/E (MME 2025)
Cell diameter	2.3 nm	70 nm	70 nm	70 nm
FL configuration	Composite (2 seg.)	Single	Single	Composite (2 seg.)
FL thickness	$2 \times 5 \text{ nm}$	1.4 nm	1.4 nm	$\text{FL}_1=0.8 \text{ nm}, \text{FL}_2=3.5 \text{ nm}$
NMS material	MgO (TB ₂)	Ru	Ru / Ta	W / Ta
NMS thickness	0.9 nm	0.85 nm	0.85 nm / 0.4 nm	0.3 nm
J_{iec} (FL–FL / HL–RL)	0.01 mJ m^{-2} (FM)	-1.32 mJ m^{-2} (AFM)	-1.5 mJ m^{-2} (AFM)	-1.32 mJ m^{-2} (AFM)
J_{iec} (RL–PL)	—	—	0.8 mJ m^{-2} (FM)	0.37 mJ m^{-2} (W) / 0.2 mJ m^{-2} (Ta)
M_S [MA m^{-1}]	1.2	1.0	1.0	1.0
K_i [mJ m^{-2}]	1.53	1.53	1.53	1.53
α	0.015	0.02	0.02	0.02

7.6 Advanced Structures

The preceding analysis established IEC as a versatile design parameter whose role depends on the device architecture: a modifier of sequential switching, a guarantor of RL stability, and a controller of PL alignment. This section advances to structures that combine these roles: SAF-enhanced hybrid FLs with metallic spacers, systematic torque decomposition, and ds-MTJ structures with dual reference layers [216, 217, 218].

7.6.1 SAF Architecture and Operating Principle

A SAF couples two ferromagnetic layers anti-parallel through RKKY exchange via a metallic spacer (Section 6.4), providing stray-field cancellation, enhanced stability through the AFM coupling energy barrier, and external field immunity.

The material and geometry parameters for all three configurations are given in Table 7.4. Figure 7.28 compares the switching dynamics across the three configurations. All three share the SAF reference structure of Stack C (Figure 7.22 (c)): a CoPt HL AFM-coupled to a CoFeB RL through Ru at $J_{\text{iec}} \approx -1.32 \text{ mJ m}^{-2}$ [207], with a PL FM-coupled to the RL through Ta at $J_{\text{iec}} \approx 0.8 \text{ mJ m}^{-2}$ [213]. **Stack C** employs a single FL, serving as the reference.




Figure 7.27: Layer stack of the SAF-enhanced composite FL configuration from [216]. Stack D and Stack E share the same architecture, differing only in the NMS material between the FL segments: tungsten (W, $\lambda_{sf} = 2.4 \text{ nm}$) for Stack D, tantalum (Ta, $\lambda_{sf} = 1.9 \text{ nm}$) for Stack E. IEC acts at the HL–RL interface (through Ru, AFM coupling) and at the FL₁–FL₂ interface (through the metallic spacer, FM coupling). Stack C employs a single FL (same architecture as Figure 7.22 (c)).




Figure 7.28: Switching duration with different FM coupling for three stack configurations: (a) Stack C with single FL at strong (0.8 mJ m^{-2}) and weak (0.2 mJ m^{-2}) FM coupling, (b) Stack D with W spacer composite FL at $J_{\text{iec}} = 0.37 \text{ mJ m}^{-2}$ and 0.8 mJ m^{-2} , and (c) Stack E with Ta spacer composite FL at $J_{\text{iec}} = 0.2 \text{ mJ m}^{-2}$ and 0.8 mJ m^{-2} . Weaker coupling delays switching in all three configurations, with the Ta spacer (c) exhibiting the longest switching times [216].

Stack D and **Stack E** replace the single FL with a composite FL (Figure 7.27), differing in the NMS material between the FL segments: W for Stack D with FM coupling of 0.37 mJ m^{-2} [192], and Ta for Stack E with 0.2 mJ m^{-2} [195, 214].

7.6.2 Reliability Requirements: The Critical Role of Coupling Strength

The reliability of SAF-enhanced structures depends critically on the IEC strength at each interface (Table 7.4). If the coupling is too weak, the reference structure can become unstable during FL switching, leading to device failure.

The simulations reveal two distinct reliability regimes [215] (Figures 7.25 and 7.26). With strong coupling ($|J_{\text{iec}}| > 1 \text{ mJ m}^{-2}$), the RL remains rigidly coupled to the hard layer throughout FL switching. The IEC effective field dominates over parasitic torques and stray fields, producing symmetric $P \rightarrow AP$ and $AP \rightarrow P$ behavior with no instability. With weak coupling ($|J_{\text{iec}}| < 0.5 \text{ mJ m}^{-2}$), the parasitic spin-torque on the RL overcomes the IEC field and initiates RL reversal, leading to (1) TMR collapse when the RL aligns parallel to the FL in both states, (2) asymmetric switching with one direction showing instability, and (3) permanent reference corruption in severe cases.

The three-dimensional visualizations confirm this contrast: in Stack B (Figure 7.25), strong AFM coupling maintains uniform RL magnetization throughout FL switching, while weak coupling produces RL domain formation and a corrupted final state. In Stack C (Figure 7.26), the analogous failure mode involves PL reversal when the RL-PL FM coupling is insufficient, producing a write failure indistinguishable from the initial P state by resistance measurement.

7.6.3 Spin-Torque Analysis and IEC Design Criterion

The spin-torque distribution across the full multilayer stack [216] reveals parasitic torques on the reference structure during FL switching: the spin current traverses the entire stack with absorption at every ferromagnetic interface, generating significant torque on the RL and PL. When the RL–PL coupling is insufficient, the PL is the most vulnerable element due to lower anisotropy and weaker FM coupling through Ta. In this case, PL reversal disrupts spin injection and causes write errors. The comparison between W and Ta NMS materials shows that W provides larger torques due to its longer spin-flip length (2.4 nm vs. 1.9 nm) [219]. A detailed spacer material comparison is provided in the hybrid FL analysis (Section 7.6.4).

The IEC must exceed the maximum parasitic torque on the reference structure under worst-case switching conditions:

$$|J_{\text{iec}}| > \frac{|\mathbf{T}_{\text{RL/PL}}|_{\text{max}} \cdot t_{\text{RL/PL}}}{M_S}, \quad (7.16)$$

where $|\mathbf{T}_{\text{RL/PL}}|_{\text{max}}$ is the maximum torque on the RL or PL during switching, and $t_{\text{RL/PL}}$ is the layer thickness. For typical parameters, this requirement yields $|J_{\text{iec}}| > 1 \text{ mJ m}^{-2}$, consistent with the empirical observations from the switching simulations.

7.6.4 Hybrid FL with Metallic Spacers

The hybrid FL replaces the MgO TB between FL segments with a metallic NMS (W or Ta, thickness 0.3 nm), preserving the composite architecture ($\text{FL}_1 = 0.8 \text{ nm}$, NMS, $\text{FL}_2 = 3.5 \text{ nm}$) while adding a direct metallic path for spin transport between the segments. The three stack configurations from Section 7.6.1 single FL (Stack C), W spacer composite (Stack D), and Ta spacer composite (Stack E) are now investigated with a 70 nm device diameter at a 2 V bias.

W ($\lambda_{sf} = 2.4 \text{ nm}$) transmits more spin current than Ta ($\lambda_{sf} = 1.9 \text{ nm}$) [219], producing stronger IEC-induced effective fields and faster switching. At weak FM coupling (0.37 mJ m^{-2} for W [192], 0.2 mJ m^{-2} for Ta [195, 214]), both spacers enable sequential switching of FL_1 and FL_2 but with distinct dynamics visible in the three-dimensional magnetization states. At the stronger coupling of 0.8 mJ m^{-2} [213, 220], the rigid alignment between FL_1 and FL_2 forces a more synchronized reversal but at reduced switching speed. Figure 7.28 confirms that weaker coupling systematically delays switching across all three configurations, with the Ta spacer (Stack E) exhibiting the longest switching times due to its shorter spin-flip length attenuating the inter-segment coupling.

The three-dimensional magnetization states visualize the effect of coupling strength on switching dynamics across all three configurations. In the single-FL baseline (Figure 7.28 (a)), weak RL–PL coupling allows PL reversal during FL switching, while strong coupling maintains PL stability. In the composite FL stacks, both the W spacer (Figure 7.29 (a)–(f)) and Ta spacer (Figure 7.30 (a)–(f)) exhibit sequential FL_1/FL_2 reversal under weak coupling, while strong coupling (Figure 7.29 (g)–(l), Figure 7.30 (g)–(l)) forces synchronized reversal. The sequential switching arises because the dominant




Figure 7.29: Three-dimensional magnetization states during P→AP switching in Stack D (W spacer composite FL), contrasting weak ($J_{\text{iec}} = 0.37 \text{ mJ m}^{-2}$) and strong ($J_{\text{iec}} = 0.8 \text{ mJ m}^{-2}$) FM FL₁–FL₂ coupling. (a)–(c) weak coupling, full view and (d)–(f) weak coupling, FL₂ clipped to reveal FL₁, (g)–(i) strong coupling, full view and (j)–(l) strong coupling, FL₂ clipped to reveal FL₁. Rows show the initial state, mid-switching ($\sim 0.5 \text{ ns}$), and final state ($\sim 1.7 \text{ ns}$). With weak coupling, FL₁ reverses first producing a transient plateau at $\langle m_x \rangle \approx 0.5$, with strong coupling, synchronized reversal eliminates the plateau [216].

spin-transfer torque is concentrated at the PL|TB and TB|FL₁ interfaces, while the torque at the metallic NMS is negligible: FL₁ therefore reaches its critical switching threshold first. Once FL₁ reverses, the resulting FL₁–FL₂ misalignment increases the IEC effective field, which, combined with spin-accumulation from the reversed FL₁, drives FL₂ to complete the transition. With the W spacer, this two-step process produces a transient plateau at $\langle m_x \rangle \approx 0.5$ visible in the switching trajectory (Figure 7.28 (b)) and the mid-switching snapshots (Figure 7.29 (b),(e)). The Ta spacer's shorter spin-flip length suppresses this plateau (Figure 7.30 (b),(e)). Strong coupling forces synchronized reversal at reduced speed.

7.6.5 Four-Stage Torque Decomposition

To understand how the metallic spacer affects switching dynamics, the P→AP reversal is decomposed into four characteristic stages defined by the FL magnetization angle θ relative to the initial P state. Figure 7.31 presents the spin-transfer torque distribution across the full stack for all three configurations at each stage.




Figure 7.30: Three-dimensional magnetization states during P→AP switching in Stack E (Ta spacer composite FL), contrasting weak ($J_{\text{iec}} = 0.2 \text{ mJ m}^{-2}$) and strong ($J_{\text{iec}} = 0.8 \text{ mJ m}^{-2}$) FM FL₁–FL₂ coupling. (a)–(c) weak coupling, full view and (d)–(f) weak coupling, FL₂ clipped to reveal FL₁, (g)–(i) strong coupling, full view and (j)–(l) strong coupling, FL₂ clipped to reveal FL₁. Rows show the initial state, mid-switching ($\sim 0.75 \text{ ns}$), and final state ($\sim 1.7 \text{ ns}$). The shorter spin-flip length of Ta (1.9 nm) compared to W (2.4 nm) produces weaker inter-segment torques, eliminating the transient plateau and slowing overall switching [216].

At the initial P state ($\theta \approx 0^\circ$, Figure 7.31 (a)–(c)), torques concentrate at the PL|TB and TB|FL interfaces, with all three stacks exhibiting similar distributions because the FL magnetization remains aligned. At $\theta = 45^\circ$ ((d)–(f)), the field-like torque $T_{S,x}$ increases significantly, driving the magnetization toward the transverse plane, critically, the composite FL stacks ((e),(f)) develop additional torque peaks at the NMS interfaces that are absent in the single-FL configuration ((d)). At $\theta = 135^\circ$ ((g)–(i)), torques increase by an order of magnitude compared to Stage 1: the damping-like torque $T_{S,z}$ dominates and completes the FL reversal. At the same time, the PL becomes destabilized by parasitic torques. The W spacer ((h)) generates systematically larger torques than the Ta spacer ((i)) due to its longer spin-flip length (2.4 nm vs. 1.9 nm). At the final AP state ($\theta \approx 180^\circ$, (j)–(l)), residual torques persist as the system equilibrates. Notably, the PL experiences significant residual torque that must be counteracted by the SAF coupling to prevent reversal.

The composite FL thus introduces additional torque injection points at NMS interfaces, with W generating systematically larger torques than Ta across all four stages, quantifying the mechanism behind the faster switching in Figure 7.28 (b).




Figure 7.31: Four-stage torque decomposition during P→AP switching for (a)–(d) Stack C with single FL, (e)–(h) Stack D with W spacer composite FL, and (i)–(l) Stack E with Ta spacer composite FL. Rows show the spin-transfer torque distribution at the initial P state ($\theta \approx 0^\circ$), $\theta = 45^\circ$, $\theta = 135^\circ$, and the final AP state ($\theta \approx 180^\circ$). The composite FL stacks show additional torque peaks at the NMS interfaces absent in the single-FL configuration [216].

7.6.6 Write Efficiency Analysis

The practical relevance of IEC engineering depends on whether optimizing switching speed compromises data retention. Write efficiency, defined as the ratio of thermal stability to critical switching current:

$$\eta_{\text{write}} = \frac{\Delta}{I_{c0}} \quad (7.17)$$

captures this fundamental tradeoff. To anchor this metric, the macrospin critical current [31] for the 70 nm single-FL Stack C yields $I_{c0} \approx 585 \mu\text{A}$ (critical current density $J_{c0} \approx 15 \text{ MA cm}^{-2}$, within the 1 MA cm^{-2} to 20 MA cm^{-2} range reported for CoFeB|MgO pMTJs [185, 187]), with a thermal stability factor $\Delta \approx 600$ that far exceeds the retention requirement. The 2 V bias used in the simulations corresponds to approximately $2 I_{c0}$, consistent with standard overdrive conditions. For symmetric composite FL layers ($\text{FL}_1 = \text{FL}_2$ thickness), the IEC contributions to both the energy barrier $\mathcal{E}_B$ and the critical current I_{c0} cancel, rendering write efficiency weakly sensitive to coupling strength IEC adds the same energy to both the parallel and anti-parallel states. For the asymmetric configuration studied here ($\text{FL}_2 = 3.5 \text{ nm}$, $\text{FL}_1 = 0.8 \text{ nm}$), incomplete cancellation means IEC modestly affects η_{write} : weaker coupling reduces I_{c0} more than $\mathcal{E}_B$, yielding a slight efficiency gain. The practical implication is that the designer can optimize IEC for switching speed via the spin-torque redistribution mechanism identified in the preceding subsection without significantly compromising write efficiency, effectively decoupling two otherwise conflicting design objectives [216].

While the hybrid FL analysis established IEC and spacer material as parameters governing switching speed and write efficiency, a complementary approach to improving switching performance is to engineer additional spin-torque sources. The ds-MTJ architecture, investigated in parallel [218], positions a second reference layer opposite the conventional RL to generate cooperative torques that accelerate switching.

7.6.7 Device Concept and Dual Torque Mechanism

A ds-MTJ positions the FL between two reference layers with anti-parallel magnetizations (Figure 7.32): RL_1 is separated from the FL by a conventional MgO TB providing the Slonczewski STT [19], while RL_2 is separated by a thin NMS providing an additional GMR-type torque [119]. Because RL_1 and RL_2 point in opposite directions ($-\hat{x}$ and $+\hat{x}$, respectively), both torques favor the same switching direction, producing cooperative torque action that accelerates switching compared to single-RL structures [218]. The complementary contributions of a predominantly damping-like Slonczewski torque from RL_1 and a combined damping-like and field-like GMR torque from RL_2 add constructively at the two FL interfaces. The GMR torque magnitude depends on the spin-flip length λ_{sf} of the NMS material, with longer λ_{sf} allowing more spin current from RL_2 to reach the FL.

7.6.8 Influence of Spacer Material

The NMS material determines the GMR torque magnitude [219]. The layer parameters are summarized in Table 7.5. Two commonly used spacer materials were compared: Ru ($\lambda_{sf} \approx 4 \text{ nm}$), which enables efficient spin transport and strong RKKY-type IEC [108, 109, 110] tunable between FM and AFM character via spacer thickness, and Ta ($\lambda_{sf} \approx 1.9 \text{ nm}$), which produces weaker GMR torques but is preferred for process integration as a standard seed/capping material in MTJ stacks.

Figure 7.33 quantifies the enhancement in switching speed. (1) The Ru-NMS ds-MTJ achieves a $2\text{--}3\times$ speedup over conventional single-RL pMTJ structures across the




Figure 7.32: Layer stack of the double spin-torque MTJ (ds-MTJ) architecture, in which the FL is sandwiched between two reference layers with anti-parallel magnetizations separated by a TB and a non-magnetic metallic spacer. The anti-parallel RL configuration produces cooperative torques from both interfaces that constructively add, and the NMS material choice (Ru vs. Ta) controls the secondary torque magnitude via its spin-flip length [218].

full voltage range, reaching sub-nanosecond switching times above 2 V with standard CoFeB|MgO materials [218]. (2) The Ta-NMS variant provides a $1.2\text{--}1.5\times$ enhancement, consistent with the $\sim 2\times$ ratio of the Ru and Ta spin-flip lengths, confirming that the secondary torque scales with spin current transmission efficiency through the spacer. (3) The speedup factor is larger at lower voltages, where the primary RL_1 torque is marginal, and the secondary RL_2 contribution becomes proportionally more significant in determining switching success.

7.6.9 Role of Interlayer Exchange Coupling

The IEC effective field (7.15) adds a contribution to the switching dynamics beyond the spin-transfer torques. A systematic parameter study varied $J_{\text{iec},1}$ (RL_1 –FL coupling through the TB) and $J_{\text{iec},2}$ (FL– RL_2 coupling through the NMS) from -2 to $+2\text{ mJ m}^{-2}$, covering the experimentally observed range for CoFeB|MgO and CoFeB/Ru/CoFeB interfaces [204].

Figure 7.34 reveals four key trends: (1) zero coupling ($J_{\text{iec},1} = J_{\text{iec},2} = 0$) yields the fastest switching by eliminating IEC-induced barriers [116], (2) FM coupling assists one transition direction but opposes the other, creating asymmetry: for RL_1 in $-\hat{x}$, positive $J_{\text{iec},1}$ assists $P \rightarrow AP$ but hinders $AP \rightarrow P$, (3) AFM coupling produces the opposite asymmetry, favoring $AP \rightarrow P$ over $P \rightarrow AP$, and (4) strong coupling can introduce writing errors by trapping the magnetization in unintended states when the IEC-induced effective field competes with the spin-transfer torque and distorts the energy landscape.




Figure 7.33: Switching performance of ds-MTJ vs. conventional pMTJ structures, showing (a) the switching time ratio and (b) absolute switching times for both transition directions. The ds-MTJ with Ru-NMS achieves a 2–3 $\times$ speedup and sub-nanosecond switching at moderate voltages, while Ta-NMS provides a 1.2–1.5 $\times$ enhancement, consistent with the ratio of their spin-flip lengths (4 nm vs. 1.9 nm) [218].

Table 7.5: Layer parameters of the ds-MTJ architecture [218]. The FL is sandwiched between RL₁ (MgO tunnel barrier) and RL₂ (metallic spacer). Two NMS materials are compared.

Parameter	Ru-NMS	Ta-NMS
Cell diameter	70 nm	
RL ₁ (CoFeB) thickness	2.5 nm	
FL (CoFeB) thickness	1.3 nm	
RL ₂ (CoFeB) thickness	2.5 nm	
MgO TB thickness	0.9 nm	
NMS material	Ru	Ta
NMS thickness	2 nm	2 nm
λ_{sf} (NMS)	4 nm	1.9 nm
J_{iec} (RL ₁ –FL)	0 (baseline)	
J_{iec} (FL–RL ₂)	0 (baseline)	

The torque distribution confirms the dual-torque mechanism: both RL₁ (through the TB) and RL₂ (through the NMS) contribute torque peaks at opposite FL interfaces that constructively add for the anti-parallel RL configuration, with the RL₂ peak larger for Ru than for Ta, consistent with the longer spin-flip length.

7.6.10 SAF-Enhanced Double Spin-Torque MTJ

The preceding subsections examined how metallic spacers and IEC affect the effective field and torque distributions within a single reference layer configuration. The natural extension combines the ds-MTJ concept from Section 7.6.7, where dual reference layers




Figure 7.34: Influence of IEC on ds-MTJ switching dynamics with Ru-NMS, exploring a 3×5 parameter matrix of AFM, zero, and FM coupling through TB_1 (columns) with varying NMS coupling (curves within each panel) for (a) AP→P and (b) P→AP transitions. Zero coupling ($J_{\text{iec},1} = J_{\text{iec},2} = 0$) generally yields the fastest switching, while strong coupling can introduce energy landscape distortions that trap the magnetization in unintended states [218].

generate cooperative torques, with SAF stabilization from Section 7.6.1, replacing one or both reference layers with SAF stacks. In this configuration, RL_1 is replaced with a CoPt hard layer antiferromagnetically coupled through Ru, while RL_2 uses a CoFeB layer ferromagnetically coupled through Ta. The device geometry is 70 nm diameter at 0.6 V bias.

The NMS thickness determines the coupling strength and hence the switching behavior (Table 7.6). For the Ru spacer, 0.4 nm produces excessively strong AFM coupling ($J_{\text{iec}} = -2.1 \text{ mJ m}^{-2}$), causing oscillatory behavior, where the FL cannot settle into the AP state. At 0.7 nm, coupling moderates and switching completes, though slowly. The optimal Ru thickness is 1.0 nm ($J_{\text{iec}} = -0.65 \text{ mJ m}^{-2}$), yielding clean sub-nanosecond switching for both AP→P and P→AP transitions. For the Ta spacer, 0.25 nm is optimal ($J_{\text{iec}} = 0.4 \text{ mJ m}^{-2}$, FM), achieving the fastest switching among Ta thicknesses, while 0.5 nm and 0.75 nm produce progressively weaker coupling and slower switching.

The torque analysis (Figure 7.37) confirms that the double SAF configuration provides enhanced reference layer stability through additional pinning fields while maintaining the cooperative dual torque mechanism essential for sub-nanosecond




Figure 7.35: Influence of IEC on ds-MTJ switching dynamics with Ta-NMS, exploring a 3×5 parameter matrix of AFM, zero, and FM coupling through TB_1 (columns) with varying NMS coupling (curves within each panel) for (a) AP $\rightarrow$ P and (b) P $\rightarrow$ AP transitions. Zero coupling ($J_{ie,1} = J_{ie,2} = 0$) generally yields the fastest switching, while strong coupling can introduce energy landscape distortions that trap the magnetization in unintended states [218].

switching. The Ru spacer generates stronger RL_2 torques due to its longer spin-flip length, consistent with the W vs Ta comparison in the hybrid FL analysis.

New physical insight. The systematic investigation across four stack families reveals IEC as a versatile design parameter whose role evolves with device complexity. In ultra-scaled composite structures, weak FM coupling (0.01 mJ m^{-2}) suppresses back-hopping while preserving fast switching [116]. In SAF-enhanced structures, a quantitative threshold $|J_{ie}| > 1 \text{ mJ m}^{-2}$ is required to prevent parasitic spin-torque-induced RL reversal during FL switching [215]. The hybrid FL with metallic spacers reveals that the spin-flip length of the NMS material is the fundamental parameter controlling torque transmission between FL segments, with W (2.4 nm) enabling systematically faster switching than Ta (1.9 nm). The four-stage torque decomposition (Figure 7.31) shows that composite FL structures introduce additional torque injection points at NMS interfaces, producing parasitic torques on the reference structure that scale with the spacer's spin-flip length. In ds-MTJ structures, anti-parallel dual reference layers with zero IEC generate cooperative torques that achieve a 2–3 $\times$ switching speed enhancement and sub-nanosecond operation with standard CoFeB|MgO materials [218]. The integration of SAF layers into ds-MTJ structures

Table 7.6: SAF-enhanced ds-MTJ spacer parameters and coupling values [217]. The optimal thickness for each spacer material is highlighted.

Spacer	Thickness	J_{iec} [mJ m^{-2}]	Behavior
Ru	0.4 nm	-2.1 (AFM)	Oscillatory (too strong)
Ru	0.7 nm	-1.0 (AFM)	Slow switching
Ru	1.0 nm	-0.65 (AFM)	Optimal
Ta	0.25 nm	$+0.4$ (FM)	Optimal
Ta	0.5 nm	$+0.2$ (FM)	Slower switching
Ta	0.75 nm	$+0.1$ (FM)	Slowest switching




Figure 7.36: Switching trajectories for SAF-enhanced ds-MTJ structures comparing Ru and Ta NMS at varying thicknesses: (a) $P \rightarrow AP$ and (b) $AP \rightarrow P$ transitions. Optimal performance at Ru 1.0 nm ($J_{\text{iec}} = -0.65 \text{ mJ m}^{-2}$, AFM) and Ta 0.25 nm ($J_{\text{iec}} = 0.4 \text{ mJ m}^{-2}$, FM). Ru at 0.4 nm (-2.1 mJ m^{-2}) exhibits oscillatory behavior from excessively strong AFM coupling [217].

(Figure 7.36) demonstrates that optimal NMS thickness balances coupling strength against spin transport efficiency: Ru at 1.0 nm (-0.65 mJ m^{-2} AFM) and Ta at 0.25 nm (0.4 mJ m^{-2} FM). This completes the IEC role progression established across the chapter: from modifier (back-hopping suppression, Section 7.4) to controller (cyclic switching) to design parameter (ds-MTJ speed, Section 7.6.7) to reliability threshold (Section 7.6.2) to engineering parameter (hybrid FL spacer selection and NMS thickness optimization).




Figure 7.37: Spin-transfer torque distribution across the full multilayer stack for (a) single SAF ds-MTJ and (b) double SAF ds-MTJ, comparing Ru (AFM) and Ta (FM) NMS. The Ru spacer generates stronger torques on RL₂ due to its longer spin-flip length, while the double SAF configuration provides enhanced reference layer stability with minimal switching speed penalty [217].

7.7 Synthesis

The comprehensive simulation studies presented in this chapter yield a consistent picture: IEC emerges as the central design parameter across all device families studied, with its role evolving systematically from a modifier to an engineering parameter. This section synthesizes the key findings into practical design rules and summarizes the main conclusions.

7.7.1 Design Guidelines for Ultra-Scaled Composite FL Structures

For 2.3 nm diameter devices with CoFeB|MgO materials, the optimal FL segment length is 3 nm to 5 nm: shorter segments (2 nm to 3 nm) exhibit back-hopping instability during P→AP switching due to insufficient anisotropy barriers. In comparison, longer segments (above 6 nm to 7 nm) fail to complete switching at moderate bias (≤ 2 V) from stray field stabilization of intermediate states. Interface-localized anisotropy models should be used for elongated structures, where they predict 10–15% faster P→AP switching compared to bulk averaging (transition region $\tau = 0.5$ nm to 1.0 nm). The middle TB polarization determines the switching regime: keeping TB₂ below TB₁ ensures binary operation, whereas enhancing TB₂ above TB₁ enables controlled back-hopping with four accessible MLC states.

7.7.2 Design Guidelines for Multi-Level Cell Implementation

Symmetric FL segments (both 3 nm to 4 nm) maximize resistance separation between intermediate states. Reliable state addressing requires write pulse duration control with approximately 100 ps resolution, a minimum 20% resistance margin between adjacent

states through TB thickness optimization, and verified thermal stability ($\Delta > 40$) for all four configurations, particularly the intermediate anti-parallel FL_1/FL_2 states.

7.7.3 Design Guidelines for ds-MTJ and SAF-Enhanced Structures

For NMS material selection, Ru provides $2\text{--}3\times$ switching speed enhancement via its longer spin-flip length (4 nm), while Ta offers $1.2\text{--}1.5\times$ enhancement with better process compatibility. SAF coupling strength must maintain $|J_{\text{iec}}| > 1.0 \text{ mJ m}^{-2}$ (AFM) at the HL–RL interface and $|J_{\text{iec}}| > 0.5 \text{ mJ m}^{-2}$ (FM) at the RL–PL interface, with margins for $\pm 30\%$ process variation and temperature dependence. NMS thickness should target the first AFM coupling peak for Ru (approximately 0.8 nm to 0.9 nm). For Ta, select a thickness that provides the desired FM coupling while maintaining adequate spin transport, and verify the coupling sign and magnitude experimentally.

7.7.4 Cross-Architecture Performance Comparison

Table 7.7 consolidates the quantitative performance of all seven device configurations. As device complexity increases from single FL composites to SAF-enhanced and hybrid structures, the role of IEC escalates from an incidental modifier to an engineering parameter jointly optimized across multiple interfaces. The final two columns, derived from Section 7.6, represent the culmination of this progression: spacer material, thickness, and coupling strength are jointly tuned for switching speed, write efficiency, and reference layer stability.

IEC as unifying design parameter. Across all structures studied in this chapter, inter-layer coupling emerges as the single physical theme whose role evolves most dramatically with device complexity. In ultra-scaled composite FLs, inter-segment spin-transfer torques through TB_2 drive the sequential switching of individual ferromagnetic segments without any IEC ($J_{\text{iec}} = 0$, Section 7.3). In composite structures operating in the back-hopping regime, these same spin-transport torques govern the cyclic dynamics that generate windmill states, enabling multilevel storage (Section 7.4). In ds-MTJ structures with dual reference layers, IEC becomes a tunable design parameter that governs the switching speed enhancement achievable through constructive torque cooperation (Section 7.5). Finally, in SAF-enhanced structures, IEC strength exceeding 1 mJ m^{-2} is required as a quantitative threshold for reference layer stability, without which the device fails catastrophically during write operations (Section 7.6.1). In advanced structures combining hybrid FLs with SAF-stabilized double spin-torque reference layers (Section 7.6), IEC transitions from a constraint to be satisfied to an engineering parameter to be optimized, with the NMS material's spin-flip length and thickness jointly determining switching speed, write efficiency, and reference layer stability. This progression from modifier to controller to design parameter to reliability threshold to engineering parameter constitutes the chapter's central physical narrative. It also provides a systematic framework for future STT-MRAM development. As

device structures grow in complexity, IEC management transitions from a secondary consideration to a primary design axis alongside thermal stability and switching speed.

Table 7.7: Cross-architecture comparison of the seven device configurations studied in this chapter. Each column represents a distinct architecture family, and rows compare key design metrics and the role of IEC. Switching speedup is relative to a conventional single-RL pMTJ. Dashes indicate metrics not applicable or not directly quantified for that architecture. The ultra-scaled and MLC studies (Sections 7.3 and 7.4) use $J_{\text{iec}} = 0$, the IEC role listed refers to the effect observed when coupling is subsequently introduced (Section 7.5).

Metric	Ultra-Scaled	MLC	ds-MTJ (Ru)	ds-MTJ (Ta)	SAF-A	Hybrid FL (W)	SAF+ds-MTJ
Cell diameter	2.3 nm	2.3 nm	70 nm	70 nm	70 nm	70 nm	70 nm
FL type	Composite (2 seg.)	Composite (2 seg.)	Single	Single	Single	Composite (W)	Single
Optimal FL length	3 nm to 5 nm	3 nm to 4 nm	1.3 nm	1.3 nm	1.4 nm	0.8 nm+3.5 nm	1.3 nm
Bits per cell	1	2	1	1	1	1	1
Switching speedup	—	—	2–3×	1.2–1.5×	~1×	1.5–2×	2–3×
IEC role	Modifier ¹	Controller ¹	Design param.	Design param.	Threshold	Eng. param.	Eng. param.
Key design param.	FL length	TB ₂ polar.	NMS material	NMS material	$ J_{\text{iec}} $	Spacer λ_{sf}	NMS thickness
Critical IEC value	—	TB ₂ > TB ₁	Zero coupling	Zero coupling	$>1.0 \text{ mJ m}^{-2}$	0.37 mJ m^{-2}	-0.65 mJ m^{-2}

¹Role observed when IEC is introduced in Section 7.5. Base studies use $J_{\text{iec}} = 0$.

Summary and Outlook

8.1 Summary

This thesis presents a comprehensive computational framework for modeling magnetization dynamics and spin transport in ultra-scaled, multilayered spin-transfer torque magnetoresistive random access memory (STT-MRAM) devices. The framework addresses three tightly coupled aspects of next-generation magnetoresistive random access memory (MRAM) design: the micromagnetic solver and its time integration, the self-consistent treatment of spin-transfer torques via drift-diffusion transport, and the application of the unified solver to advanced device architectures. The principal contributions are summarized below.

8.1.1 Micromagnetic Solver and Time Integration

This work implements a three-dimensional micromagnetic solver was implemented using the finite element method (FEM) for spatial discretization. The demagnetization field, which poses a primary computational bottleneck due to its long-range nature, is computed using a hybrid finite element-boundary element method (FE-BEM) that avoids discretizing the surrounding free space. Furthermore, three time integration strategies are developed, implemented, and systematically compared for the spatially discretized Landau–Lifshitz–Gilbert (LLG) equation.

The first approach employs adaptive higher-order backward differentiation formula (BDF) methods provided by the *SUNDIALS* library, which achieve unconditional A-stability and handle the stiffness of the exchange interaction through a fully implicit treatment. The second approach uses implicit-explicit (IMEX) schemes, in which the stiff exchange contribution is integrated implicitly, whereas the remaining effective-field terms are treated explicitly. The third approach employs a tangent-plane time integration scheme that reformulates the nonlinear, constrained LLG equation as a sequence of linear saddle-point problems, thereby preserving the unit-length constraint on the magnetization at the algebraic level.

A systematic comparison using the μ MAG Standard Problem 4 demonstrated that the tangent-plane scheme offers superior stability and robust constraint preservation for the stiff dynamics characteristic of ultra-scaled magnetic elements. At the same time, the fully implicit BDF method achieves the largest time steps during current-

driven switching events. These findings provide concrete guidance for selecting the appropriate time integrator based on the problem's dominant physics.

8.1.2 Coupled Spin and Charge Drift-Diffusion Transport

The micromagnetic solver is extended with a coupled spin and charge drift-diffusion formalism to compute spin-transfer torques from the first principles of semiclassical transport, rather than relying on the simplified Slonczewski macro-spin approximation. This extension provides a unified framework for capturing both interface-induced Slonczewski torques at tunnel barriers and bulk Zhang–Li torques arising from magnetization textures within ferromagnetic layers.

A key contribution of this work is the derivation and implementation of specialized boundary conditions at material interfaces. At the magnesium oxide (MgO) tunnel barrier, spin-filtering boundary conditions relate the tunneling spin current to the angular and voltage-dependent conductance of the junction, reproducing the expected angular and bias dependencies of the torque. At metallic interfaces between ferromagnetic layers and nonmagnetic spacers, spin-dephasing boundary conditions account for the rapid loss of transverse spin coherence. This rigorous treatment of interfaces proved essential for accurately resolving the spatially varying torque landscape in multi-layered structures, where the spin-accumulation and its gradient exhibit sharp variations near each interface.

8.1.3 Interlayer Exchange Coupling and Interface-Mapping

Additionally, a novel numerical treatment for interlayer exchange coupling (IEC) is introduced via an interface-mapping algorithm. In conventional approaches, simulating synthetic antiferromagnet (SAF) structures or composite free layers separated by ultrathin nonmagnetic spacers (on the order of angstroms) would require volumetric meshing of the spacer layer, thereby imposing severe constraints on mesh resolution and computational cost. The interface-mapping algorithm circumvents this limitation by projecting the magnetization degrees of freedom from one interface onto the opposing interface, effectively treating the IEC as a surface coupling without discretizing the spacer volume.

This approach enables the simulation of realistic multi-layered device stacks, including SAF-enhanced reference layers and composite free layers with multiple ferromagnetic segments, at a computational cost comparable to that of simulating a single ferromagnetic layer.

8.1.4 Application to Advanced MRAM Architectures

The unified solver is applied to investigate three interconnected phenomena critical to the scaling and performance of next-generation STT-MRAM.

First, the switching dynamics in ultra-scaled cells with composite free layers were analyzed. The three-dimensional simulations reveal that composite structures exhibit sequential switching of individual ferromagnetic segments, with the switching order

and speed governed by the interplay of spin torques and stray fields. At the same time, IEC provides an additional control parameter when introduced in multi-layered configurations. This sequential behavior, which is fundamentally inaccessible to macro-spin models, has direct implications for the write error rate and the distribution of switching currents.

Second, the back-hopping phenomenon, commonly regarded as a failure mode in composite magnetic tunnel junction (MTJ) structures, is shown to be a deterministic consequence of the torque dynamics rather than a stochastic effect. The simulations demonstrate that, by tuning the relative tunnel-barrier polarizations and the free layer geometry, back-hopping can be exploited to achieve stable intermediate magnetization configurations. These distinct states enable multi-level cell (MLC) operation, in which multiple bits are stored per physical cell through distinct resistance levels accessible via pulse-duration modulation. This represents a paradigm shift from treating back-hopping as a defect to recognizing it as a design opportunity.

Third, advanced device architectures are systematically investigated. The role of IEC is first quantified across three representative stack configurations (SAF without and with a polarization layer, and ultra-scaled composite free layer (FL)), identifying specific coupling windows that ensure reference layer stability without compromising write efficiency. Double-spin torque magnetic tunnel junction (ds-MTJ) structures, which employ dual reference layers to generate both giant magnetoresistance (GMR) and tunneling magnetoresistance (TMR) torques simultaneously, are shown to achieve sub-nanosecond switching through the constructive cooperation of the two torque contributions. Hybrid FL architectures with metallic spacers (W and Ta) demonstrate that the spacer material's spin-flip length is the key parameter governing switching speed enhancement, and a systematic four-stage torque decomposition revealed the distinct contributions of each ferromagnetic interface during the switching process. A cross-architecture comparison consolidating all seven device configurations studied throughout the chapter provides quantitative design guidelines for high-density MRAM arrays, in which IEC management, stray-field control, and reference-layer integrity are jointly optimized.

8.2 Outlook: Advanced Numerical Methods for Micromagnetics

While the physical insights presented in this thesis have advanced the understanding of STT-MRAM, they have also highlighted the limitations of current numerical frameworks. The resolution of complex, non-uniform textures and domain wall curvatures, which are essential for MLC operation, demands a higher degree of mathematical rigor. Future research must therefore complement ongoing device analysis with advances in computational methods.

8.2.1 Implicit Treatment of Spin-Transfer Torque in IMEX Schemes

The IMEX time integration scheme implemented in this work treats the exchange interaction implicitly to circumvent the mesh-dependent Courant–Friedrichs–Lewy (CFL) restriction, $\Delta t \sim h^2$. However, as demonstrated in Section 5.5.2, the spin-transfer torque (STT) terms were treated explicitly, introducing a separate stability constraint $\Delta t \sim 1/|\tau_{\text{STT}}|$ that dominates during current-driven switching events.

Future implementations should incorporate the STT terms into the implicit partition:

$$f_I = -\frac{\alpha\gamma}{1+\alpha^2} \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{exch}}) + \tau_{\text{STT}}(\mathbf{m}, \mathbf{S}) \quad (8.1)$$

This requires extending the preconditioner to include the STT Jacobian, which involves both the field-like (skew-symmetric) and Slonczewski (rank-one update) contributions:

$$\mathbf{J}_{\text{STT}} = c_{\text{FL}}[\mathbf{S}]_{\times} - c_{\text{SL}}(\mathbf{m} \otimes \mathbf{S} + (\mathbf{m} \cdot \mathbf{S})\mathbf{I}) \quad (8.2)$$

The fully-implicit BDF method already incorporates this Jacobian and achieves a $50\times$ reduction in the number of time steps compared to explicit methods. Extending the IMEX framework similarly would enable adaptive time-stepping during STT-driven switching while retaining the computational advantages of explicitly treating non-stiff contributions, such as demagnetization.

8.2.2 Higher-Order Finite Elements for Exchange Resolution

The accurate discretization of the exchange interaction is paramount in micromagnetics. The exchange energy density is proportional to $(\nabla \mathbf{m})^2$, and the resulting effective field involves the Laplacian $\nabla^2 \mathbf{m}$.

Current implementations often rely on linear (first-order) Lagrange basis functions. While computationally inexpensive, linear elements yield constant gradients within elements, making evaluation of the second derivative (the exchange field) discontinuous or requiring recovery techniques that introduce numerical noise. Future work should integrate higher-order finite elements (e.g., quadratic or cubic basis functions). By increasing the polynomial order (p -refinement) rather than only the mesh density (h -refinement), one can achieve spectral convergence, thereby drastically reducing the error in the exchange-field calculation, while accurately capturing the smooth rotation of magnetization in magnetic vortices without the numerical artifacts associated with linear meshes.

8.2.3 Parallel-in-Time Integration

The timescale disparity in spintronics is a major computational bottleneck. The LLG equation is stiff, requiring femtosecond time steps ($\Delta t \approx 1$ fs) to resolve precession, while write error rate studies require simulations spanning μ s or ms. The clock speed of a single processor core inherently limits the time step in sequential time-stepping.

To break this barrier, future solvers should adopt parallel-in-time integration strategies, such as the Parallel Full Approximation Scheme in Space and Time (PFASST). Unlike standard domain decomposition, which parallelizes only spatial degrees of freedom, PFASST iterates over multiple time slices simultaneously. By employing a hierarchy of discretizations (coarse propagators on CPU, fine propagators on GPU), this method could accelerate long-duration switching statistical studies by orders of magnitude, making the direct simulation of low-probability failure events computationally feasible.

8.2.4 Unified Space-Time Finite Element Methods

The traditional Method of Lines approach, discretizing space first to obtain a system of ordinary differential equations (ODEs) and then discretizing time, introduces operator-splitting errors and decouples spatial and temporal resolution. However, in fast switching events such as back-hopping, spatial gradients and temporal transients are intimately coupled.

A promising direction is the development of a fully coupled space-time finite element method (ST-FEM). By treating time as a fourth dimension, the entire space-time cylinder (x, y, z, t) can be discretized using 4D finite elements. This approach offers the potential for unconditional stability, bypassing the CFL conditions that limit explicit solvers, as well as space-time adaptivity, in which the mesh is refined dynamically in localized space-time regions where high activity occurs (e.g., the trajectory of a domain wall), while leaving the rest of the 4D domain coarse. Implementing these advanced numerical schemes would not only improve the accuracy of IEC and SAF simulations but also unveil subtle, high-frequency switching modes that are currently smoothed out by lower-order temporal approximations.

8.3 Conclusion

This thesis has demonstrated that the complex physics of ultra-scaled STT-MRAM devices, spanning spin transport across multiple interfaces, interlayer exchange coupling in composite stacks, and deterministic switching phenomena in sub-20 nm geometries, can be captured within a single, unified computational framework. The three-dimensional drift-diffusion approach developed here provides spatial resolution of torque densities, spin-accumulation profiles, and magnetization textures that are fundamentally inaccessible to macro-spin or one-dimensional models. This resolution reveals phenomena, including sequential switching, deterministic back-hopping, and the inherently non-additive coupling of spin-transfer torques in multi-segment stacks, that collectively reshape the design paradigm for next-generation devices.

From a practical standpoint, the results establish quantitative design guidelines across seven device configurations spanning four architecture families (ultra-scaled composite FL, SAF-enhanced, hybrid FL, and ds-MTJ). The identification of specific IEC windows for SAF stability, the demonstration of controlled MLC operation through intermediate magnetization states, and the optimization of double spin-torque

and hybrid structures for sub-nanosecond switching provide an actionable roadmap for the development of higher-density, faster, and more reliable MRAM.

Looking ahead, the next leap in device performance will likely stem not only from new materials and advances in fabrication, but also from the mathematical capability to simulate them. The transition to higher-order finite elements, parallel-in-time integration, and fully coupled space-time solvers promises to reveal hidden dynamics at the quantum-classical boundary, enabling predictive simulation of phenomena that remain beyond the reach of current methods. This thesis provides both the physical foundation and the computational starting point for that pursuit.

Bibliography

- [1] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient Processing of Deep Neural Networks: A Tutorial and Survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017. doi: [10.1109/JPROC.2017.2761740](https://doi.org/10.1109/JPROC.2017.2761740).
- [2] J. E. Lilienfeld, "Method and Apparatus for Controlling Electric Currents," US Patent US1,745,175A, January 28, 1930.
- [3] J. Bardeen and W. H. Brattain, "The Transistor, a Semiconductor Triode," *Physical Review*, vol. 74, pp. 230–231, 1948. doi: [10.1103/PhysRev.74.230](https://doi.org/10.1103/PhysRev.74.230).
- [4] F. M. Wanlass, "Low Stand-By Power Complementary Field Effect Circuitry," US Patent US3 356 858A, December 5, 1967.
- [5] J. W. Forrester, "Digital Information Storage in Three Dimensions Using Magnetic Cores," *Journal of Applied Physics*, vol. 22, no. 1, pp. 44–48, 1951. doi: [10.1063/1.1699817](https://doi.org/10.1063/1.1699817).
- [6] B. Dieny, I. L. Prejbeanu, K. Garello, P. Gambardella, P. Freitas *et al.*, "Opportunities and Challenges for Spintronics in the Microelectronics Industry," *Nature Electronics*, vol. 3, no. 8, pp. 446–459, 2020. doi: [10.1038/s41928-020-0461-5](https://doi.org/10.1038/s41928-020-0461-5).
- [7] J. Kim, A. Paul, P. A. Crowell, S. J. Koester, S. S. Sapatnekar *et al.*, "Spin-Based Computing: Device Concepts, Current Status, and a Case Study on a High-Performance Microprocessor," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 106–130, 2015. doi: [10.1109/JPROC.2014.2361767](https://doi.org/10.1109/JPROC.2014.2361767).
- [8] X. Wang, "RRAM-Based In-Memory Computing Architecture Designs," Ph.D. dissertation, University of Michigan, Ann Arbor, MI, 2023. [Online]. Available: <https://deepblue.lib.umich.edu/handle/2027.42/178016>
- [9] J. Liao, S. Dai, R.-C. Peng, J. Yang, B. Zeng *et al.*, "HfO₂-Based Ferroelectric Thin Film and Memory Device Applications in the Post-Moore Era: A Review," *Fundamental Research*, vol. 3, no. 3, pp. 332–345, 2023. doi: [10.1016/j.fmre.2023.02.010](https://doi.org/10.1016/j.fmre.2023.02.010).
- [10] M. Le Gallo, R. Khaddam-Aljameh, M. Stanisavljevic, A. Vasilopoulos, B. Kersting *et al.*, "A 64-Core Mixed-Signal In-Memory Compute Chip Based on Phase-Change Memory for Deep Neural Network Inference," *Nature Electronics*, vol. 6, no. 9, pp. 680–693, 2023. doi: [10.1038/s41928-023-01010-1](https://doi.org/10.1038/s41928-023-01010-1).
- [11] A. D. Kent and D. C. Worledge, "A New Spin on Magnetic Memories," *Nature Nanotechnology*, vol. 10, pp. 187–191, 2015. doi: [10.1038/nnano.2015.24](https://doi.org/10.1038/nnano.2015.24).

- [12] D. Apalkov, B. Dieny, and J. M. Slaughter, “Magnetoresistive Random Access Memory,” *Proceedings of the IEEE*, vol. 104, no. 10, pp. 1796–1830, 2016. doi: [10.1109/JPROC.2016.2590142](https://doi.org/10.1109/JPROC.2016.2590142).
- [13] T. Hanyu, T. Endoh, D. Suzuki, H. Koike, Y. Ma *et al.*, “Standby-Power-Free Integrated Circuits Using MTJ-Based VLSI Computing,” *Proceedings of the IEEE*, vol. 104, no. 10, pp. 1844–1863, 2016. doi: [10.1109/JPROC.2016.2574939](https://doi.org/10.1109/JPROC.2016.2574939).
- [14] M. Jullière, “Tunneling Between Ferromagnetic Films,” *Physics Letters A*, vol. 54, no. 3, pp. 225–226, 1975. doi: [10.1016/0375-9601\(75\)90174-7](https://doi.org/10.1016/0375-9601(75)90174-7).
- [15] C. Chappert, A. Fert, and F. N. Van Dau, “The Emergence of Spin Electronics in Data Storage,” *Nature Materials*, vol. 6, no. 11, pp. 813–823, 2007. doi: [10.1038/nmat2024](https://doi.org/10.1038/nmat2024).
- [16] S. Maekawa and U. Gafvert, “Electron Tunneling Between Ferromagnetic Films,” *IEEE Transactions on Magnetics*, vol. 18, no. 2, pp. 707–708, 1982. doi: [10.1109/TMAG.1982.1061834](https://doi.org/10.1109/TMAG.1982.1061834).
- [17] J. S. Moodera and G. Mathon, “Spin Polarized Tunneling in Ferromagnetic Junctions,” *Journal of Magnetism and Magnetic Materials*, vol. 200, no. 1, pp. 248–273, 1999. doi: [10.1016/S0304-8853\(99\)00515-6](https://doi.org/10.1016/S0304-8853(99)00515-6).
- [18] L. Berger, “Emission of Spin Waves by a Magnetic Multilayer Traversed by a Current,” *Physical Review B*, vol. 54, pp. 9353–9358, 1996. doi: [10.1103/PhysRevB.54.9353](https://doi.org/10.1103/PhysRevB.54.9353).
- [19] J. Slonczewski, “Current-Driven Excitation of Magnetic Multilayers,” *Journal of Magnetism and Magnetic Materials*, vol. 159, no. 1, pp. L1–L7, 1996. doi: [10.1016/0304-8853\(96\)00062-5](https://doi.org/10.1016/0304-8853(96)00062-5).
- [20] D. C. Ralph and M. D. Stiles, “Spin Transfer Torques,” *Journal of Magnetism and Magnetic Materials*, vol. 320, pp. 1190–1216, 2008. doi: [10.1016/j.jmmm.2007.12.019](https://doi.org/10.1016/j.jmmm.2007.12.019).
- [21] K. Ando, S. Fujita, J. Ito, S. Yuasa, Y. Suzuki *et al.*, “Spin-Transfer Torque Magnetoresistive Random-Access Memory Technologies for Normally Off Computing,” *Journal of Applied Physics*, vol. 115, no. 17, p. 172607, 2014. doi: [10.1063/1.4869828](https://doi.org/10.1063/1.4869828).
- [22] C. Safranski, G. Hu, J. Z. Sun, P. Hashemi, S. L. Brown *et al.*, “Reliable Sub-Nanosecond MRAM with Double Spin-Torque Magnetic Tunnel Junctions,” in *2022 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*, 2022, pp. 288–289. doi: [10.1109/VLSITechnologyandCir46769.2022.9830306](https://doi.org/10.1109/VLSITechnologyandCir46769.2022.9830306).
- [23] J. Grollier, D. Querlioz, K. Y. Camsari, K. Everschor-Sitte, S. Fukami *et al.*, “Neuromorphic Spintronics,” *Nature Electronics*, vol. 3, no. 7, pp. 360–370, 2020. doi: [10.1038/s41928-019-0360-9](https://doi.org/10.1038/s41928-019-0360-9).
- [24] D. Marković, A. Mizrahi, D. Querlioz, and J. Grollier, “Physics for Neuromorphic Computing,” *Nature Reviews Physics*, vol. 2, no. 9, pp. 499–510, 2020. doi: [10.1038/s42254-020-0208-2](https://doi.org/10.1038/s42254-020-0208-2).
- [25] S. S. P. Parkin, C. Kaiser, A. Panchula, P. M. Rice, B. Hughes *et al.*, “Giant Tunneling Magnetoresistance at Room Temperature with MgO (100) Tunnel Barriers,” *Nature Materials*, vol. 3, no. 12, pp. 862–867, 2004. doi: [10.1038/nmat1256](https://doi.org/10.1038/nmat1256).
- [26] S. Yuasa, T. Nagahama, A. Fukushima, Y. Suzuki, and K. Ando, “Giant Room-Temperature Magnetoresistance in Single-Crystal Fe/MgO/Fe Magnetic Tunnel Junctions,” *Nature Materials*, vol. 3, no. 12, pp. 868–871, 2004. doi: [10.1038/nmat1257](https://doi.org/10.1038/nmat1257).

- [27] A. Brataas, A. D. Kent, and H. Ohno, “Current-Induced Torques in Magnetic Materials,” *Nature Materials*, vol. 11, no. 5, pp. 372–381, 2012. doi: [10.1038/nmat3311](https://doi.org/10.1038/nmat3311).
- [28] S. Ikeda, K. Miura, H. Yamamoto, K. Mizunuma, H. D. Gan *et al.*, “A Perpendicular-Anisotropy CoFeB–MgO Magnetic Tunnel Junction,” *Nature Materials*, vol. 9, no. 9, pp. 721–724, 2010. doi: [10.1038/nmat2804](https://doi.org/10.1038/nmat2804).
- [29] M. Beg, M. Lang, and H. Fangohr, “Ubermag: Toward More Effective Micromagnetic Workflows,” *IEEE Transactions on Magnetics*, vol. 58, no. 2, pp. 1–5, 2022. doi: [10.1109/TMAG.2021.3078896](https://doi.org/10.1109/TMAG.2021.3078896).
- [30] J. J. Joos, P. Bassirian, P. Gypens, J. Mulkers, K. Litzius *et al.*, “Tutorial: Simulating Modern Magnetic Material Systems in Mumax3,” *Journal of Applied Physics*, vol. 134, no. 17, p. 171101, 2023. doi: [10.1063/5.0160988](https://doi.org/10.1063/5.0160988).
- [31] J. Z. Sun, “Spin-Current Interaction with a Monodomain Magnetic Body: A Model Study,” *Physical Review B*, vol. 62, no. 1, pp. 570–578, 2000. doi: [10.1103/PhysRevB.62.570](https://doi.org/10.1103/PhysRevB.62.570).
- [32] C. Petitjean, D. Luc, and X. Waintal, “Unified Drift-Diffusion Theory for Transverse Spin Currents in Spin Valves, Domain Walls, and Other Textured Magnets,” *Physical Review Letters*, vol. 109, p. 117204, 2012. doi: [10.1103/PhysRevLett.109.117204](https://doi.org/10.1103/PhysRevLett.109.117204).
- [33] M. Gajek, J. J. Nowak, J. Z. Sun, P. L. Trouilloud, E. J. O’Sullivan *et al.*, “Spin Torque Switching of 20 nm Magnetic Tunnel Junctions with Perpendicular Anisotropy,” *Applied Physics Letters*, vol. 100, no. 13, p. 132408, 2012. doi: [10.1063/1.3694270](https://doi.org/10.1063/1.3694270).
- [34] S. Bandiera, R. C. Sousa, Y. Dahmane, C. Ducruet, C. Portemont *et al.*, “Comparison of Synthetic Antiferromagnets and Hard Ferromagnets as Reference Layer in Magnetic Tunnel Junctions With Perpendicular Magnetic Anisotropy,” *IEEE Magnetics Letters*, vol. 1, pp. 3 000 204–3 000 204, 2010. doi: [10.1109/LMAG.2010.2052238](https://doi.org/10.1109/LMAG.2010.2052238).
- [35] S. Zhang, P. M. Levy, and A. Fert, “Mechanisms of Spin-Polarized Current-Driven Magnetization Switching,” *Physical Review Letters*, vol. 88, p. 236601, 2002. doi: [10.1103/PhysRevLett.88.236601](https://doi.org/10.1103/PhysRevLett.88.236601).
- [36] J.-G. J. Zhu and C. Park, “Magnetic Tunnel Junctions,” *Materials Today*, vol. 9, no. 11, pp. 36–45, 2006. doi: [10.1016/S1369-7021\(06\)71693-5](https://doi.org/10.1016/S1369-7021(06)71693-5).
- [37] O. Golonzka, J. G. Alzate, U. Arslan, M. Bohr, P. Bai *et al.*, “MRAM as Embedded Non-Volatile Memory Solution for 22FFL FinFET Technology,” in *2018 IEEE International Electron Devices Meeting (IEDM)*, 2018, pp. 18.1.1–18.1.4. doi: [10.1109/IEDM.2018.8614620](https://doi.org/10.1109/IEDM.2018.8614620).
- [38] S. Bhatti, R. Sbiaa, A. Hirohata, H. Ohno, S. Fukami *et al.*, “Spintronics Based Random Access Memory: A Review,” *Materials Today*, vol. 20, no. 9, pp. 530–548, 2017. doi: [10.1016/j.mattod.2017.07.007](https://doi.org/10.1016/j.mattod.2017.07.007).
- [39] J. A. Katine, F. J. Albert, R. A. Buhrman, E. B. Myers, and D. C. Ralph, “Current-Driven Magnetization Reversal and Spin-Wave Excitations in Co/Cu/Co Pillars,” *Physical Review Letters*, vol. 84, pp. 3149–3152, 2000. doi: [10.1103/PhysRevLett.84.3149](https://doi.org/10.1103/PhysRevLett.84.3149).
- [40] A. V. Khvalkovskiy, D. Apalkov, S. Watts, R. Chepulskii, R. S. Beach *et al.*, “Basic Principles of STT-MRAM Cell Operation in Memory Arrays,” *Journal of Physics D: Applied Physics*, vol. 46, no. 7, p. 074001, 2013. doi: [10.1088/0022-3727/46/7/074001](https://doi.org/10.1088/0022-3727/46/7/074001).

- [41] F. C. Williams and T. Kilburn, "Electronic Digital Computers," *Nature*, vol. 162, no. 4117, pp. 487–487, 1948. doi: [10.1038/162487a0](https://doi.org/10.1038/162487a0).
- [42] W. N. Papian, "A Coincident-Current Magnetic Memory Cell for the Storage of Digital Information," *Proceedings of the IRE*, vol. 40, no. 4, pp. 475–478, 1952. doi: [10.1109/JRPROC.1952.274045](https://doi.org/10.1109/JRPROC.1952.274045).
- [43] R. A. Abbott, W. M. Regitz, and J. A. Karp, "A 4K MOS Dynamic Random-Access Memory," *IEEE Journal of Solid-State Circuits*, vol. 8, no. 5, pp. 292–298, 1973. doi: [10.1109/JSSC.1973.1050406](https://doi.org/10.1109/JSSC.1973.1050406).
- [44] M. N. Baibich, J. M. Broto, A. Fert, F. Nguyen Van Dau, F. Petroff *et al.*, "Giant Magnetoresistance of (001)Fe/(001)Cr Magnetic Superlattices," *Physical Review Letters*, vol. 61, no. 21, pp. 2472–2475, 1988. doi: [10.1103/PhysRevLett.61.2472](https://doi.org/10.1103/PhysRevLett.61.2472).
- [45] G. Binasch, P. Grünberg, F. Saurenbach, and W. Zinn, "Enhanced Magnetoresistance in Layered Magnetic Structures with Antiferromagnetic Interlayer Exchange," *Physical Review B*, vol. 39, no. 7, pp. 4828–4830, 1989. doi: [10.1103/PhysRevB.39.4828](https://doi.org/10.1103/PhysRevB.39.4828).
- [46] S. A. Wolf, D. D. Awschalom, R. A. Buhrman, J. M. Daughton, S. von Molnár *et al.*, "Spintronics: A Spin-Based Electronics Vision for the Future," *Science*, vol. 294, no. 5546, pp. 1488–1495, 2001. doi: [10.1126/science.1065389](https://doi.org/10.1126/science.1065389).
- [47] T. Miyazaki, T. Yaoi, and S. Ishio, "Large Magnetoresistance Effect in 82Ni-Fe/Al-Al₂O₃/Co Magnetic Tunneling Junction," *Journal of Magnetism and Magnetic Materials*, vol. 98, no. 1, pp. L7–L9, 1991. doi: [10.1016/0304-8853\(91\)90417-9](https://doi.org/10.1016/0304-8853(91)90417-9).
- [48] J. S. Moodera, L. R. Kinder, T. M. Wong, and R. Meservey, "Large Magnetoresistance at Room Temperature in Ferromagnetic Thin Film Tunnel Junctions," *Physical Review Letters*, vol. 74, no. 16, pp. 3273–3276, 1995. doi: [10.1103/PhysRevLett.74.3273](https://doi.org/10.1103/PhysRevLett.74.3273).
- [49] W. H. Butler, X.-G. Zhang, T. C. Schulthess, and J. M. MacLaren, "Spin-Dependent Tunneling Conductance of Fe/MgO/Fe(001) Sandwiches," *Physical Review B*, vol. 63, no. 5, p. 054416, 2001. doi: [10.1103/PhysRevB.63.054416](https://doi.org/10.1103/PhysRevB.63.054416).
- [50] J. Mathon and A. Umerski, "Theory of Tunneling Magnetoresistance of Magnetic Junctions," *Physical Review B*, vol. 63, no. 22, p. 220403(R), 2001. doi: [10.1103/PhysRevB.63.220403](https://doi.org/10.1103/PhysRevB.63.220403).
- [51] T. Kawahara, K. Ito, R. Takemura, and H. Ohno, "Spin-Transfer Torque RAM Technology: Review and Prospect," *Microelectronics Reliability*, vol. 52, no. 4, pp. 613–627, 2012. doi: [10.1016/j.microrel.2011.09.028](https://doi.org/10.1016/j.microrel.2011.09.028).
- [52] W. Kim, J. H. Jeong, Y. Kim, W. C. Lim, J. H. Kim *et al.*, "Extended Scalability of Perpendicular STT-MRAM Towards Sub-20nm MTJ Node," in *2011 International Electron Devices Meeting*, 2011, pp. 24.1.1–24.1.4. doi: [10.1109/IEDM.2011.6131602](https://doi.org/10.1109/IEDM.2011.6131602).
- [53] S. Ikeda, J. Hayakawa, Y. Ashizawa, Y. M. Lee, K. Miura *et al.*, "Tunnel Magnetoresistance of 604 % at 300 K by Suppression of Ta Diffusion in CoFeB/MgO/CoFeB Pseudo-Spin-Valves Annealed at High Temperature," *Applied Physics Letters*, vol. 93, no. 8, p. 082508, 2008. doi: [10.1063/1.297643](https://doi.org/10.1063/1.297643).
- [54] J. Åkerman, "Toward a Universal Memory," *Science*, vol. 308, no. 5721, pp. 508–510, 2005. doi: [10.1126/science.1110549](https://doi.org/10.1126/science.1110549).

- [55] T. Sousa, J.-P. Wang, and T. Low, “Bidirectional Switching Assisted by Interlayer Exchange Coupling in Perpendicular MTJs,” *Physical Review B*, vol. 101, p. 081404(R), 2020. doi: [10.1103/PhysRevB.101.081404](https://doi.org/10.1103/PhysRevB.101.081404).
- [56] Y. Saito, T. Roy, S. Ikeda, M. Shirai, H. Honjo *et al.*, “Strong Antiferromagnetic Interlayer Exchange Coupling Induced by Small Additions of Re to an Ir Interlayer in Synthetic Antiferromagnetic Systems,” *Scientific Reports*, vol. 15, no. 1, p. 8977, 2025. doi: [10.1038/s41598-025-94088-w](https://doi.org/10.1038/s41598-025-94088-w).
- [57] Q. Shao, Z. Wang, Y. Zhou, S. Fukami, D. Querlioz *et al.*, “Spintronic Memristors for Computing,” *npj Spintronics*, vol. 3, no. 1, p. 16, 2025. doi: [10.1038/s44306-025-00078-z](https://doi.org/10.1038/s44306-025-00078-z).
- [58] B. Dieny, V. S. Speriosu, S. S. P. Parkin, B. A. Gurney, D. R. Wilhoit *et al.*, “Giant Magnetoresistance in Soft Ferromagnetic Multilayers,” *Physical Review B*, vol. 43, pp. 1297–1300, 1991. doi: [10.1103/PhysRevB.43.1297](https://doi.org/10.1103/PhysRevB.43.1297).
- [59] G. A. Prinz, “Spin-Polarized Transport,” *Physics Today*, vol. 48, no. 4, pp. 58–63, 1995. doi: [10.1063/1.881459](https://doi.org/10.1063/1.881459).
- [60] S. S. P. Parkin, “Spin-Dependent Tunnelling in Magnetic Nanostructures and Magnetic Tunnel Junctions for High Density Magnetic Random Access Memory,” *IBM Journal of Research and Development*, vol. 44, no. 1/2, pp. 245–262, 2000. doi: [10.1147/rd.441.0245](https://doi.org/10.1147/rd.441.0245).
- [61] M. Bowen, V. Cros, F. Petroff, A. Fert, C. Martínez Boubeta *et al.*, “Large Magnetoresistance in Fe/MgO/FeCo(001) Epitaxial Tunnel Junctions on GaAs(001),” *Applied Physics Letters*, vol. 79, no. 11, pp. 1655–1657, 2001. doi: [10.1063/1.1404125](https://doi.org/10.1063/1.1404125).
- [62] J. Faure-Vincent, C. Tiusan, E. Jouguelet, F. Canet, M. Sajieddine *et al.*, “High Tunnel Magnetoresistance in Epitaxial Fe/MgO/Fe Tunnel Junctions,” *Applied Physics Letters*, vol. 82, no. 25, pp. 4507–4509, 2003. doi: [10.1063/1.1586785](https://doi.org/10.1063/1.1586785).
- [63] J. C. Slonczewski, “Conductance and Exchange Coupling of Two Ferromagnets Separated by a Tunneling Barrier,” *Physical Review B*, vol. 39, no. 10, pp. 6995–7002, 1989. doi: [10.1103/PhysRevB.39.6995](https://doi.org/10.1103/PhysRevB.39.6995).
- [64] Y. Ji, J. Liu, and C. Yang, “Novel Modeling and Dynamic Simulation of Magnetic Tunnel Junctions for Spintronic Sensor Development,” *Journal of Physics D: Applied Physics*, vol. 50, no. 2, p. 025005, 2016. doi: [10.1088/0022-3727/50/2/025005](https://doi.org/10.1088/0022-3727/50/2/025005).
- [65] H. Jaffrès, D. Lacour, F. N. V. Dau, J. Briatico, F. Petroff *et al.*, “Angular Dependence of the Tunnel Magnetoresistance in Transition-Metal-Based Junctions,” *Physical Review B*, vol. 64, no. 6, p. 064427, 2001. doi: [10.1103/PhysRevB.64.064427](https://doi.org/10.1103/PhysRevB.64.064427).
- [66] T. Miyazaki and N. Tezuka, “Giant Magnetic Tunneling Effect in Fe/Al₂O₃/Fe Junction,” *Journal of Magnetism and Magnetic Materials*, vol. 139, no. 3, pp. L231–L234, 1995. doi: [10.1016/0304-8853\(95\)90001-2](https://doi.org/10.1016/0304-8853(95)90001-2).
- [67] D. Wang, C. Nordman, J. Daughton, Z. Qian, and J. Fink, “70% TMR at Room Temperature for SDT Sandwich Junctions with CoFeB as Free and Reference Layers,” *IEEE Transactions on Magnetics*, vol. 40, no. 4, pp. 2269–2271, 2004. doi: [10.1109/TMAG.2004.830219](https://doi.org/10.1109/TMAG.2004.830219).

- [68] S. Ikeda, J. Hayakawa, Y. M. Lee, R. Sasaki, T. Meguro *et al.*, “Dependence of Tunnel Magnetoresistance in MgO Based Magnetic Tunnel Junctions on Ar Pressure During MgO Sputtering,” *Japanese Journal of Applied Physics*, vol. 44, no. 48, pp. L1442–L1445, 2005. doi: [10.1143/JJAP.44.L1442](https://doi.org/10.1143/JJAP.44.L1442).
- [69] S. Yuasa, T. Nagahama, A. Fukushima, and Y. Suzuki, “High TMR Ratio of 410% in Epitaxial Fe/MgO/Fe Tunnel Junctions Annealed at High Temperature,” *Applied Physics Letters*, vol. 89, no. 23, p. 232510, 2006. doi: [10.1063/1.2398002](https://doi.org/10.1063/1.2398002).
- [70] Y. Huai, F. Albert, P. Nguyen, M. Pakala, and T. Valet, “Observation of Spin-Transfer Switching in Deep Submicron-Sized and Low-Resistance Magnetic Tunnel Junctions,” *Applied Physics Letters*, vol. 84, no. 16, pp. 3118–3120, 2004. doi: [10.1063/1.1707228](https://doi.org/10.1063/1.1707228).
- [71] D. C. Worledge, G. Hu, D. W. Abraham, J. Z. Sun, P. L. Trouilloud *et al.*, “Spin Torque Switching of Perpendicular Ta/CoFeB/MgO-Based Magnetic Tunnel Junctions,” *Applied Physics Letters*, vol. 98, no. 2, p. 022501, 2011. doi: [10.1063/1.3536482](https://doi.org/10.1063/1.3536482).
- [72] H. Sato, M. Yamanouchi, K. Miura, S. Ikeda, H. D. Gan *et al.*, “Junction Size Effect on Switching Current and Thermal Stability in CoFeB/MgO Perpendicular Magnetic Tunnel Junctions,” *Applied Physics Letters*, vol. 99, no. 4, p. 042501, 2011. doi: [10.1063/1.3617429](https://doi.org/10.1063/1.3617429).
- [73] T. Devolder, J.-V. Kim, F. Garcia-Sanchez, J. Swerts, W. Kim *et al.*, “Time-Resolved Spin-Torque Switching in MgO-Based Perpendicularly Magnetized Tunnel Junctions,” *Physical Review B*, vol. 93, p. 024420, 2016. doi: [10.1103/PhysRevB.93.024420](https://doi.org/10.1103/PhysRevB.93.024420).
- [74] D. C. Worledge, “Write-error-rate of Spin-Transfer-Torque MRAM (Invited),” in *2023 IEEE International Reliability Physics Symposium (IRPS)*, 2023, pp. 1–4. doi: [10.1109/IRPS48203.2023.10117666](https://doi.org/10.1109/IRPS48203.2023.10117666).
- [75] L. Wu, M. Taouil, S. Rao, E. J. Marinissen, and S. Hamdioui, “Survey on STT-MRAM Testing: Failure Mechanisms, Fault Models, and Tests,” 2020. [Online]. Available: <https://arxiv.org/abs/2001.05463>
- [76] J. C. Sankey, Y.-T. Cui, J. Z. Sun, J. C. Slonczewski, R. A. Buhrman *et al.*, “Measurement of the Spin-Transfer-Torque Vector in Magnetic Tunnel Junctions,” *Nature Physics*, vol. 4, no. 1, pp. 67–71, 2008. doi: [10.1038/nphys783](https://doi.org/10.1038/nphys783).
- [77] J. Z. Sun, “Spin Angular Momentum Transfer in Current-Perpendicular Nanomagnetic Junctions,” *IBM Journal of Research and Development*, vol. 50, no. 1, pp. 81–100, 2006. doi: [10.1147/rd.501.0081](https://doi.org/10.1147/rd.501.0081).
- [78] D. Apalkov, S. Watts, A. Driskill-Smith, E. Chen, Z. Diao *et al.*, “Comparison of Scaling of In-Plane and Perpendicular Spin Transfer Switching Technologies by Micromagnetic Simulation,” *IEEE Transactions on Magnetics*, vol. 46, no. 6, pp. 2240–2243, 2010. doi: [10.1109/TMAG.2010.2041330](https://doi.org/10.1109/TMAG.2010.2041330).
- [79] Z. Guo, J. Yin, Y. Bai, D. Zhu, K. Shi *et al.*, “Spintronics for Energy-Efficient Computing: An Overview and Outlook,” *Proceedings of the IEEE*, vol. 109, no. 8, pp. 1398–1417, 2021. doi: [10.1109/JPROC.2021.3084997](https://doi.org/10.1109/JPROC.2021.3084997).
- [80] W. Zhao, X. Zhao, B. Zhang, K. Cao, L. Wang *et al.*, “Failure Analysis in Magnetic Tunnel Junction Nanopillar with Interfacial Perpendicular Magnetic Anisotropy,” *Materials*, vol. 9, no. 1, 2016. doi: [10.3390/ma9010041](https://doi.org/10.3390/ma9010041).

- [81] R. Sbiaa, H. Meng, and S. N. Piramanayagam, “Materials with Perpendicular Magnetic Anisotropy for Magnetic Random Access Memory,” *physica status solidi (RRL) – Rapid Research Letters*, vol. 5, no. 12, pp. 413–419, 2011. doi: [10.1002/pssr.201105420](https://doi.org/10.1002/pssr.201105420).
- [82] B. Tudu and A. Tiwari, “Recent Developments in Perpendicular Magnetic Anisotropy Thin Films for Data Storage Applications,” *Vacuum*, vol. 146, pp. 329–341, 2017. doi: [10.1016/j.vacuum.2017.01.031](https://doi.org/10.1016/j.vacuum.2017.01.031).
- [83] J. Hayakawa, S. Ikeda, Y. M. Lee, R. Sasaki, T. Meguro *et al.*, “Current-Induced Magnetization Switching in MgO Barrier Based Magnetic Tunnel Junctions with CoFeB/Ru/CoFeB Synthetic Ferrimagnetic Free Layer,” *Japanese Journal of Applied Physics*, vol. 45, no. 10L, p. L1057, 2006. doi: [10.1143/JJAP.45.L1057](https://doi.org/10.1143/JJAP.45.L1057).
- [84] S. Mangin, D. Ravelosona, J. A. Katine, M. J. Carey, B. D. Terris *et al.*, “Current-Induced Magnetization Reversal in Nanopillars with Perpendicular Anisotropy,” *Nature Materials*, vol. 5, no. 3, pp. 210–215, 2006. doi: [10.1038/nmat1595](https://doi.org/10.1038/nmat1595).
- [85] L. Cuchet, B. Rodmacq, S. Auffret, S. Bandiera, B. Dieny *et al.*, “Perpendicular Magnetic Tunnel Junctions with a Synthetic Storage or Reference Layer: A New Route Towards Pt- and Pd-Free Junctions,” *Scientific Reports*, vol. 6, p. 21246, 2016. doi: [10.1038/srep21246](https://doi.org/10.1038/srep21246).
- [86] J. Igarashi, B. Jinnai, K. Watanabe, T. Shinoda, T. Funatsu *et al.*, “Single-Nanometer CoFeB/MgO Magnetic Tunnel Junctions with High-Retention and High-Speed Capabilities,” *npj Spintronics*, vol. 2, no. 1, p. 1, 2024. doi: [10.1038/s44306-023-00003-2](https://doi.org/10.1038/s44306-023-00003-2).
- [87] Q. Yang, L. Wang, Z. Zhou, L. Wang, Y. Zhang *et al.*, “Ionic Liquid Gating Control of RKKY Interaction in FeCoB/Ru/FeCoB and (Pt/Co)₂/Ru/(Co/Pt)₂ Multilayers,” *Nature Communications*, vol. 9, no. 1, p. 991, 2018. doi: [10.1038/s41467-018-03356-z](https://doi.org/10.1038/s41467-018-03356-z).
- [88] C.-L. Yang and C.-H. Lai, “High Thermal Durability of Ru-Based Synthetic Antiferromagnet by Interfacial Engineering with Re Insertion,” *Scientific Reports*, vol. 11, no. 1, p. 15214, 2021. doi: [10.1038/s41598-021-94640-4](https://doi.org/10.1038/s41598-021-94640-4).
- [89] W. F. Brown, *Micromagnetics*, ser. Interscience tracts on physics and astronomy. Interscience Publishers, 1963. [Online]. Available: <https://cir.nii.ac.jp/crid/1130282271226285568>
- [90] C. Abert, “Micromagnetics and Spintronics: Models and Numerical Methods,” *The European Physical Journal B*, vol. 92, no. 6, p. 120, 2019. doi: [10.1140/epjb/e2019-90599-6](https://doi.org/10.1140/epjb/e2019-90599-6).
- [91] H. Kronmüller, “General Micromagnetic Theory,” in *Handbook of Magnetism and Advanced Magnetic Materials*, ser. Micromagnetism, H. Kronmüller and S. S. P. Parkin, Eds. Chichester, UK: John Wiley & Sons, Ltd, 2007, vol. 2, pp. 703–741. ISBN 978-0-470-02218-4. doi: [10.1002/9780470022184.hmm201](https://doi.org/10.1002/9780470022184.hmm201).
- [92] L. D. Landau and E. M. Lifshitz, “On the Theory of the Dispersion of Magnetic Permeability in Ferromagnetic Bodies,” *Phys. Z. Sowjetunion*, vol. 8, pp. 153–164, 1935. doi: [10.1016/B978-0-08-036364-6.50008-9](https://doi.org/10.1016/B978-0-08-036364-6.50008-9).
- [93] T. Gilbert, “A Phenomenological Theory of Damping in Ferromagnetic Materials,” *IEEE Transactions on Magnetics*, vol. 40, no. 6, pp. 3443–3449, 2004. doi: [10.1109/TMAG.2004.836740](https://doi.org/10.1109/TMAG.2004.836740).

- [94] I. Cimrák, “A Survey on the Numerics and Computations for the Landau-Lifshitz Equation of Micromagnetism,” *Archives of Computational Methods in Engineering*, vol. 15, no. 3, pp. 1–37, 2007. doi: [10.1007/BF03024947](https://doi.org/10.1007/BF03024947).
- [95] J. D. Jackson, *Classical Electrodynamics*, 3rd ed. John Wiley & Sons, Ltd, 2003. doi: [10.1002/3527600434.eap109](https://doi.org/10.1002/3527600434.eap109).
- [96] W. Heisenberg, “Mehrkörperproblem und Resonanz in der Quantenmechanik,” *Zeitschrift für Physik*, vol. 38, no. 6–7, pp. 411–426, 1926. doi: [10.1007/BF01397160](https://doi.org/10.1007/BF01397160). English translation: “Many-body Problem and Resonance in Quantum Mechanics”.
- [97] D. J. Griffiths and D. F. Schroeter, *Introduction to Quantum Mechanics*, 3rd ed. Cambridge University Press, 2018. doi: [10.1017/9781316995433](https://doi.org/10.1017/9781316995433).
- [98] W. Döring, “Mikromagnetismus,” in *Handbuch der Physik*, 1st ed., ser. Encyclopedia of Physics. Berlin, Heidelberg: Springer, 1966, vol. 18, no. 2, pp. 314–437. ISBN 978-3-642-46035-7. doi: [10.1007/978-3-642-46035-7](https://doi.org/10.1007/978-3-642-46035-7).
- [99] A. Hubert and R. Schäfer, *Magnetic Domains: The Analysis of Magnetic Microstructures*. Berlin, Heidelberg: Springer, 1998. doi: [10.1007/978-3-540-85054-0](https://doi.org/10.1007/978-3-540-85054-0).
- [100] J. E. Miltat and M. J. Donahue, “Numerical Micromagnetics: Finite Difference Methods,” in *Handbook of Magnetism and Advanced Magnetic Materials*, H. Kronmüller, S. Parkin, J. E. Miltat, and M. R. Scheinfein, Eds. John Wiley & Sons, Ltd, 2007. ISBN 978-0-470-02218-4. doi: [10.1002/9780470022184.hmm202](https://doi.org/10.1002/9780470022184.hmm202).
- [101] M. Labrune and J. Miltat, “Wall Structures in Ferro/Antiferromagnetic Exchange-Coupled Bilayers: A Numerical Micromagnetic Approach,” *Journal of Magnetism and Magnetic Materials*, vol. 151, no. 1, pp. 231–245, 1995. doi: [10.1016/0304-8853\(95\)00328-2](https://doi.org/10.1016/0304-8853(95)00328-2).
- [102] A. Aharoni, “Demagnetizing Factors for Rectangular Ferromagnetic Prisms,” *Journal of Applied Physics*, vol. 83, no. 6, pp. 3432–3434, 1998. doi: [10.1063/1.367113](https://doi.org/10.1063/1.367113).
- [103] A. Makarov, “Modeling of Emerging Resistive Switching Based Memory Cells,” Ph.D. dissertation, Technische Universität Wien, 2014. doi: [10.34726/hss.2014.23875](https://doi.org/10.34726/hss.2014.23875).
- [104] J.-B. Biot and F. Savart, “Note Sur Le Magnétisme De La Pile De Volta,” *Annales de Chimie et de Physique*, vol. 15, pp. 222–223, 1820, english translation: “Note On The Magnetism Of The Voltaic Pile”.
- [105] G. Finocchio, M. Carpentieri, B. Azzerboni, L. Torres, E. Martinez *et al.*, “Micromagnetic Simulations of Nanosecond Magnetization Reversal Processes in Magnetic Nanopillar,” *Journal of Applied Physics*, vol. 99, no. 8, p. 08G522, 2006. doi: [10.1063/1.2177049](https://doi.org/10.1063/1.2177049).
- [106] J. L. García-Palacios and F. J. Lázaro, “Langevin-Dynamics Study of the Dynamical Properties of Small Magnetic Particles,” *Physical Review B*, vol. 58, pp. 14 937–14 958, 1998. doi: [10.1103/PhysRevB.58.14937](https://doi.org/10.1103/PhysRevB.58.14937).
- [107] G. Finocchio, B. Azzerboni, G. D. Fuchs, R. A. Buhrman, and L. Torres, “Micromagnetic Modeling of Magnetization Switching Driven by Spin-Polarized Current in Magnetic Tunnel Junctions,” *Journal of Applied Physics*, vol. 101, no. 6, p. 063914, 2007. doi: [10.1063/1.2496202](https://doi.org/10.1063/1.2496202).

- [108] M. A. Ruderman and C. Kittel, “Indirect Exchange Coupling of Nuclear Magnetic Moments by Conduction Electrons,” *Physical Review*, vol. 96, no. 1, pp. 99–102, 1954. doi: [10.1103/PhysRev.96.99](https://doi.org/10.1103/PhysRev.96.99).
- [109] T. Kasuya, “A Theory of Metallic Ferro- and Antiferromagnetism on Zener’s Model,” *Progress of Theoretical Physics*, vol. 16, no. 1, pp. 45–57, 1956. doi: [10.1143/PTP.16.45](https://doi.org/10.1143/PTP.16.45).
- [110] K. Yosida, “Magnetic Properties of Cu-Mn Alloys,” *Physical Review*, vol. 106, no. 5, pp. 893–898, 1957. doi: [10.1103/PhysRev.106.893](https://doi.org/10.1103/PhysRev.106.893).
- [111] P. Bruno, “Tight-Binding Approach to the Orbital Magnetic Moment and Magnetocrystalline Anisotropy of Transition-Metal Monolayers,” *Physical Review B*, vol. 52, no. 1, pp. 411–439, 1995. doi: [10.1103/PhysRevB.52.411](https://doi.org/10.1103/PhysRevB.52.411).
- [112] E. Tsymbal and D. Pettifor, “Perspectives of Giant Magnetoresistance,” *Solid State Physics*, vol. 56, pp. 113–237, 2001. doi: [10.1016/S0081-1947\(01\)80019-9](https://doi.org/10.1016/S0081-1947(01)80019-9).
- [113] S. S. P. Parkin, R. Bhadra, and K. P. Roche, “Oscillations in Exchange Coupling and Magnetoresistance in Metallic Superlattices,” *Physical Review Letters*, vol. 66, pp. 2152–2155, 1991. doi: [10.1103/PhysRevLett.66.2152](https://doi.org/10.1103/PhysRevLett.66.2152).
- [114] B. Dieny, “Giant Magnetoresistance in Spin-Valve Multilayers,” *Journal of Magnetism and Magnetic Materials*, vol. 136, no. 3, pp. 335–359, 1994. doi: [10.1016/0304-8853\(94\)90046-9](https://doi.org/10.1016/0304-8853(94)90046-9).
- [115] J. C. Slonczewski, “Fluctuation Mechanism for Biquadratic Exchange Coupling in Magnetic Multilayers,” *Physical Review Letters*, vol. 67, no. 23, pp. 3172–3175, 1991. doi: [10.1103/PhysRevLett.67.3172](https://doi.org/10.1103/PhysRevLett.67.3172).
- [116] M. Bendra, R. L. d. Orio, S. Selberherr, W. Goes, and V. Sverdlov, “Advanced Modeling and Simulation of Multilayer Spin-Transfer Torque Magnetoresistive Random Access Memory with Interface Exchange Coupling,” *Micromachines*, vol. 15, no. 5, 2024. doi: [10.3390/mi15050568](https://doi.org/10.3390/mi15050568).
- [117] J. Slonczewski, “Currents and Torques in Metallic Magnetic Multilayers,” *Journal of Magnetism and Magnetic Materials*, vol. 247, no. 3, pp. 324–338, 2002. doi: [10.1016/S0304-8853\(02\)00291-3](https://doi.org/10.1016/S0304-8853(02)00291-3).
- [118] M. Chshiev, A. Manchon, A. Kalitsov, N. Ryzhanova, A. Vedyayev *et al.*, “Analytical Description of Ballistic Spin Currents and Torques in Magnetic Tunnel Junctions,” *Physical Review B*, vol. 92, p. 104422, 2015. doi: [10.1103/PhysRevB.92.104422](https://doi.org/10.1103/PhysRevB.92.104422).
- [119] J. C. Slonczewski, “Currents, Torques, and Polarization Factors in Magnetic Tunnel Junctions,” *Physical Review B*, vol. 71, p. 024411, 2005. doi: [10.1103/PhysRevB.71.024411](https://doi.org/10.1103/PhysRevB.71.024411).
- [120] S. Lepadatu, “Unified Treatment of Spin Torques Using a Coupled Magnetisation Dynamics and Three-Dimensional Spin Current Solver,” *Scientific Reports*, vol. 7, no. 1, p. 12937, 2017. doi: [10.1038/s41598-017-13181-x](https://doi.org/10.1038/s41598-017-13181-x).
- [121] P. Graczyk and M. Krawczyk, “Nonresonant Amplification of Spin Waves Through Interface Magnetoelectric Effect and Spin-Transfer Torque,” *Scientific Reports*, vol. 11, no. 1, p. 15692, 2021. doi: [10.1038/s41598-021-95267-1](https://doi.org/10.1038/s41598-021-95267-1).
- [122] P. M. Haney, H.-W. Lee, K.-J. Lee, A. Manchon, and M. D. Stiles, “Current Induced Torques and Interfacial Spin-Orbit Coupling: Semiclassical Modeling,” *Physical Review B*, vol. 87, no. 17, p. 174411, 2013. doi: [10.1103/PhysRevB.87.174411](https://doi.org/10.1103/PhysRevB.87.174411).

- [123] S. Fiorentini, M. Bendra, J. Ender, R. L. de Orio, W. Goes *et al.*, “Spin and Charge Drift-Diffusion in Ultra-Scaled MRAM Cells,” *Scientific Reports*, vol. 12, no. 1, p. 20958, 2022. doi: [10.1038/s41598-022-25586-4](https://doi.org/10.1038/s41598-022-25586-4).
- [124] S. Fiorentini, N. P. Jørstad, J. Ender, R. L. de Orio, S. Selberherr, and M. Bendra, “Finite Element Approach for the Simulation of Modern MRAM Devices,” *Micromachines*, vol. 14, no. 5, p. 898, 2023. doi: [10.3390/mi14050898](https://doi.org/10.3390/mi14050898).
- [125] S. Fiorentini, “Computation of Torques in Magnetic Tunnel Junctions,” Ph.D. dissertation, Technische Universität Wien, 2023. doi: [10.34726/hss.2023.109800](https://doi.org/10.34726/hss.2023.109800).
- [126] D. Berkov, K. Ramstöck, and A. Hubert, “Solving Micromagnetic Problems: Towards an Optimal Numerical Method,” *physica status solidi (a)*, vol. 137, no. 1, pp. 207–225, 1993. doi: [10.1002/pssa.2211370122](https://doi.org/10.1002/pssa.2211370122).
- [127] C. Seberino and H. N. Bertram, “Concise, Efficient Three-Dimensional Fast Multipole Method for Micromagnetics,” *IEEE Transactions on Magnetism*, vol. 37, no. 2, pp. 1078–1081, 2001. doi: [10.1109/20.917510](https://doi.org/10.1109/20.917510).
- [128] C. Grossmann, H.-G. Roos, and M. Stynes, *Numerical Treatment of Partial Differential Equations*. Berlin: Springer, 2007. doi: [10.1007/978-3-540-71584-9](https://doi.org/10.1007/978-3-540-71584-9).
- [129] H. P. Langtangen and S. Linge, *Finite Difference Computing with PDEs: A Modern Software Approach*, ser. Texts in Computational Science and Engineering. Springer Cham, 2017. doi: [10.1007/978-3-319-55456-3](https://doi.org/10.1007/978-3-319-55456-3).
- [130] R. J. LeVeque, *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2007. doi: [10.1137/1.9780898717839](https://doi.org/10.1137/1.9780898717839).
- [131] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2003. doi: [10.1137/1.9780898718003](https://doi.org/10.1137/1.9780898718003).
- [132] D. M. Young, “Iterative Methods for Solving Partial Difference Equations of Elliptic Type,” *Transactions of the American Mathematical Society*, vol. 76, pp. 92–111, 1954. doi: [10.1090/S0002-9947-1954-0059635-7](https://doi.org/10.1090/S0002-9947-1954-0059635-7).
- [133] A. Kákay, “Numerical Investigations of Micromagnetic Structures,” Ph.D. dissertation, Research Institute for Solid State Physics and Optics, Hungarian Academy of Sciences, Budapest, Hungary, 2005. [Online]. Available: <https://repozitorium.omikk.bme.hu/bitstreams/58e4eeeb-0672-418d-be04-59e763811136/download>
- [134] D. Tomáš, “Modelling of Micromagnetic Structures,” Ph.D. dissertation, Paris-Sud University, Orsay and Charles University, Prague, France and Czech Republic, 1999.
- [135] F. Alouges, “A New Finite Element Scheme for Landau-Lifshitz Equations,” *Discrete and Continuous Dynamical Systems - S*, vol. 1, no. 2, pp. 187–196, 2008. doi: [10.3934/dcdss.2008.1.187](https://doi.org/10.3934/dcdss.2008.1.187).
- [136] K. Ito, T. Devolder, C. Chappert, M. J. Carey, and J. A. Katine, “Probabilistic Behavior in Subnanosecond Spin Transfer Torque Switching,” *Journal of Applied Physics*, vol. 99, no. 8, p. 08G519, 2006. doi: [10.1063/1.2176869](https://doi.org/10.1063/1.2176869).
- [137] T. J. R. Hughes, *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*, ser. Dover Civil and Mechanical Engineering. Mineola, NY: Courier Corporation, 2012. ISBN 978-0-486-13502-1

- [138] A. Ern and J.-L. Guermond, *Theory and Practice of Finite Elements*, ser. Applied Mathematical Sciences. New York, NY: Springer, 2004, vol. 159. doi: [10.1007/978-1-4757-4355-5](https://doi.org/10.1007/978-1-4757-4355-5).
- [139] S. C. Brenner and L. R. Scott, *The Mathematical Theory of Finite Element Methods*, 3rd ed., ser. Texts in Applied Mathematics. New York, NY: Springer, 2008. doi: [10.1007/978-0-387-75934-0](https://doi.org/10.1007/978-0-387-75934-0).
- [140] A. Quarteroni and A. Valli, *Numerical Approximation of Partial Differential Equations*, 1st ed., ser. Springer Series in Computational Mathematics. Springer-Verlag Berlin Heidelberg, 1994, vol. 23. doi: [10.1007/978-3-540-85268-1](https://doi.org/10.1007/978-3-540-85268-1).
- [141] J. Nitsche, “Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von Teilräumen, die keinen Randbedingungen unterworfen sind,” *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, vol. 36, no. 1, pp. 9–15, 1971. doi: [10.1007/BF02995904](https://doi.org/10.1007/BF02995904).
- [142] A. Logg, K.-A. Mardal, and G. N. Wells, Eds., *Automated Solution of Differential Equations by the Finite Element Method: The FEniCS Book*, ser. Lecture Notes in Computational Science and Engineering. Berlin, Heidelberg: Springer, 2012. doi: [10.1007/978-3-642-23099-8](https://doi.org/10.1007/978-3-642-23099-8).
- [143] ViennaSpinMag, “An Open Access Finite Element-based Application for Calculating the Magnetization Dynamics of Multi-layered Structures Composed of Ferromagnets, Metal Spacers and Tunnel Barriers,” 2025, Accessed: 19.07.2025. [Online]. Available: <https://www.iue.tuwien.ac.at/viennaspinmag>
- [144] R. Anderson, J. Andrej, A. Barker, J. Bramwell, J.-S. Camier *et al.*, “MFEM: A Modular Finite Element Methods Library,” *Computers & Mathematics with Applications*, vol. 81, pp. 42–74, 2021. doi: [10.1016/j.camwa.2020.06.009](https://doi.org/10.1016/j.camwa.2020.06.009).
- [145] T. Kolev and V. Dobrev, “Modular Finite Element Methods (MFEM),” Jun. 2010, MFEM is a modular parallel C++ library for finite element methods. Its goal is to enable high-performance scalable finite element discretization research and application development on a wide variety of platforms, ranging from laptops to supercomputers. [Online]. Available: <https://doi.org/10.11578/dc.20171025.1248>
- [146] D. Braess, *Finite Elements: Theory, Fast Solvers, and Applications in Solid Mechanics*, 3rd ed. Cambridge University Press, 2007. doi: [10.1017/CBO9780511618635](https://doi.org/10.1017/CBO9780511618635).
- [147] M. G. Larson and F. Bengzon, *Electromagnetics*, 1st ed., ser. Texts in Computational Science and Engineering. Springer, 2013, vol. 10, pp. 327–354. ISBN 978-3-642-33287-6
- [148] Q. Chen and A. Konrad, “A Review of Finite Element Open Boundary Techniques for Static and Quasi-Static Electromagnetic Field Problems,” *IEEE Transactions on Magnetics*, vol. 33, no. 1, pp. 663–676, 1997. doi: [10.1109/20.560095](https://doi.org/10.1109/20.560095).
- [149] D. Fredkin and T. Koehler, “Hybrid Method for Computing Demagnetizing Fields,” *IEEE Transactions on Magnetics*, vol. 26, no. 2, pp. 415–417, 1990. doi: [10.1109/20.106342](https://doi.org/10.1109/20.106342).
- [150] W. Hackbusch, “A Sparse Matrix Arithmetic Based on H-Matrices. Part I: Introduction to H-Matrices,” *Computing*, vol. 62, no. 2, pp. 89–108, 1999. doi: [10.1007/s006070050015](https://doi.org/10.1007/s006070050015).

- [151] S. Börm, L. Grasedyck, and W. Hackbusch, “Introduction to Hierarchical Matrices with Applications,” *Engineering Analysis with Boundary Elements*, vol. 27, no. 5, pp. 405–422, 2003. doi: [10.1016/S0955-7997\(02\)00152-2](https://doi.org/10.1016/S0955-7997(02)00152-2).
- [152] N. C. μ Mag Modeling Group, “Standard Problem #3.” [Online]. Available: <https://www.ctcms.nist.gov/~rdm/spec3.html>
- [153] C. Abert, L. Exl, G. Selke, A. Drews, and T. Schrefl, “Numerical Methods for the Stray-Field Calculation: A Comparison of Recently Developed Algorithms,” *Journal of Magnetism and Magnetic Materials*, vol. 326, pp. 176–185, 2013. doi: [10.1016/j.jmmm.2012.08.041](https://doi.org/10.1016/j.jmmm.2012.08.041).
- [154] M. Bendra, J. Ender, S. Fiorentini, T. Hadamek, R. de Orio *et al.*, “Finite Element Method Approach to MRAM Modeling,” in *2021 44th International Convention on Information, Communication and Electronic Technology (MIPRO)*, 2021, pp. 70–73. doi: [10.23919/MIPRO52101.2021.9597194](https://doi.org/10.23919/MIPRO52101.2021.9597194).
- [155] J. M. Camacho and V. Sosa, “Alternative Method to Calculate the Magnetic Field of Permanent Magnets with Azimuthal Symmetry,” *Revista mexicana de física E*, vol. 59, pp. 8–17, 06 2013. doi: [10.31349/RevMexFisE.21.020701](https://doi.org/10.31349/RevMexFisE.21.020701).
- [156] E. Hairer and G. Wanner, *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*, 2nd ed., ser. Springer Series in Computational Mathematics. Berlin, Heidelberg: Springer-Verlag, 1996, vol. 14. doi: [10.1007/978-3-642-05221-7](https://doi.org/10.1007/978-3-642-05221-7).
- [157] V. D. Tsiantos, D. Suess, T. Schrefl, and J. Fidler, “Stiffness Analysis for the Micromagnetic Standard Problem No. 4,” *Journal of Applied Physics*, vol. 89, no. 11, pp. 7600–7602, 2001. doi: [10.1063/1.1355355](https://doi.org/10.1063/1.1355355).
- [158] D. Suess, V. Tsiantos, T. Schrefl, J. Fidler, W. Scholz *et al.*, “Time Resolved Micromagnetics Using a Preconditioned Time Integration Method,” *Journal of Magnetism and Magnetic Materials*, vol. 248, no. 2, pp. 298–311, 2002. doi: [10.1016/S0304-8853\(02\)00341-4](https://doi.org/10.1016/S0304-8853(02)00341-4).
- [159] D. Shepherd, J. Miles, M. Heil, and M. Mihajlović, “Discretization-Induced Stiffness in Micromagnetic Simulations,” *IEEE Transactions on Magnetics*, vol. 50, no. 11, pp. 1–4, 2014. doi: [10.1109/TMAG.2014.2325494](https://doi.org/10.1109/TMAG.2014.2325494).
- [160] D. Praetorius, M. Ruggeri, and B. Stifftner, “Convergence of an Implicit-Explicit Midpoint Scheme for Computational Micromagnetics,” *Computers & Mathematics with Applications*, vol. 75, no. 5, pp. 1719–1738, 2018. doi: [10.1016/j.camwa.2017.11.028](https://doi.org/10.1016/j.camwa.2017.11.028).
- [161] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban *et al.*, “SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers,” *ACM Transactions on Mathematical Software*, vol. 31, no. 3, pp. 363–396, 2005. doi: [10.1145/1089014.1089020](https://doi.org/10.1145/1089014.1089020).
- [162] U. M. Ascher, S. J. Ruuth, and R. J. Spiteri, “Implicit-Explicit Runge-Kutta Methods for Time-Dependent Partial Differential Equations,” *Applied Numerical Mathematics*, vol. 25, no. 2–3, pp. 151–167, 1997. doi: [10.1016/S0168-9274\(97\)00056-1](https://doi.org/10.1016/S0168-9274(97)00056-1).
- [163] M. Benzi, G. H. Golub, and J. Liesen, “Numerical Solution of Saddle Point Problems,” *Acta Numerica*, vol. 14, pp. 1–137, 2005. doi: [10.1017/S0962492904000212](https://doi.org/10.1017/S0962492904000212).

- [164] Y. Saad and M. H. Schultz, “GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems,” *SIAM J. Sci. Stat. Comput.*, vol. 7, no. 3, pp. 856–869, 1986. doi: [10.1137/0907058](https://doi.org/10.1137/0907058).
- [165] E. Kim and J. Wilkening, “Convergence of a Mass-Lumped Finite Element Method for the Landau-Lifshitz Equation,” *Quarterly of Applied Mathematics*, vol. 76, no. 2, pp. 383–405, 2018. doi: [10.1090/qam/1485](https://doi.org/10.1090/qam/1485).
- [166] N. J. Mauser, C.-M. Pfeiler, D. Praetorius, and M. Ruggeri, “Unconditional Well-Posedness and IMEX Improvement of a Family of Predictor-Corrector Methods in Micromagnetics,” *Applied Numerical Mathematics*, vol. 180, pp. 33–54, 2022. doi: [10.1016/j.apnum.2022.05.008](https://doi.org/10.1016/j.apnum.2022.05.008).
- [167] R. E. Bank and D. J. Rose, “Some Error Estimates for the Box Method,” *SIAM Journal on Numerical Analysis*, vol. 24, no. 4, pp. 777–787, 1987. doi: [10.1137/0724050](https://doi.org/10.1137/0724050).
- [168] H. A. van der Vorst, “Bi-CGSTAB: A Fast and Smoothly Converging Variant of Bi-CG for the Solution of Nonsymmetric Linear Systems,” *SIAM J. Sci. Stat. Comput.*, vol. 13, no. 2, pp. 631–644, 1992. doi: [10.1137/0913035](https://doi.org/10.1137/0913035).
- [169] A. M. Bruaset, *Krylov Subspace Iterations for Sparse Linear Systems*. Boston, MA: Birkhäuser Boston, 1997, pp. 255–280. ISBN 978-1-4612-1984-2
- [170] A. Greenbaum, *Iterative Methods for Solving Linear Systems*. Philadelphia: Society for Industrial and Applied Mathematics, 1997. doi: [10.1137/1.9781611970937](https://doi.org/10.1137/1.9781611970937).
- [171] Lawrence Livermore National Laboratory, “HYPRE: Scalable Linear Solvers and Multigrid Methods,” <https://github.com/hypre-space/hypre>, 2024.
- [172] P. J. Roache, *Verification and Validation in Computational Science and Engineering*. Albuquerque, NM: Hermosa Publishers, 1998. ISBN 978-0913478080
- [173] Micromagnetic Modeling Activity Group, “ μ MAG Standard Problem #4,” <https://www.ctcms.nist.gov/~rdm/std4/spec4.html>, 2000, accessed: 2024.
- [174] J. Miltat, G. Albuquerque, and A. Thiaville, *An Introduction to Micromagnetics in the Dynamic Regime*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 1–33. ISBN 978-3-540-40907-6
- [175] M. d’Aquino, C. Serpico, and G. Miano, “Geometrical Integration of Landau-Lifshitz-Gilbert Equation Based on the Mid-Point Rule,” *Journal of Computational Physics*, vol. 209, no. 2, pp. 730–753, 2005. doi: <https://doi.org/10.1016/j.jcp.2005.04.001>.
- [176] G. Di Fratta, C.-M. Pfeiler, D. Praetorius, M. Ruggeri, and B. Stiftner, “Linear Second-Order IMEX-Type Integrator for the (Eddy Current) Landau-Lifshitz-Gilbert Equation,” *IMA Journal of Numerical Analysis*, vol. 40, no. 4, pp. 2802–2838, 2020. doi: [10.1093/imanum/drz046](https://doi.org/10.1093/imanum/drz046).
- [177] J. Zhan, L. Yang, R. Du, and Z. Cui, “Towards Preserving Geometric Properties of Landau-Lifshitz-Gilbert Equation Using Multistep Methods,” *Communications in Computational Physics*, vol. 36, no. 4, pp. 1024–1050, 2024. doi: [10.4208/cicp.OA-2023-0201](https://doi.org/10.4208/cicp.OA-2023-0201).
- [178] P. Chen, C. Li, Z. Liang *et al.*, “Temporal High-Order Accurate Numerical Scheme for the Landau-Lifshitz-Gilbert Equation,” *Mathematics*, vol. 12, no. 8, p. 1179, 2024. doi: [10.3390/math12081179](https://doi.org/10.3390/math12081179).

- [179] T. Valet and A. Fert, “Theory of the Perpendicular Magnetoresistance in Magnetic Multilayers,” *Physical Review B*, vol. 48, pp. 7099–7113, 1993. doi: [10.1103/PhysRevB.48.7099](https://doi.org/10.1103/PhysRevB.48.7099).
- [180] V. S. Rychkov, S. Borlenghi, H. Jaffres, A. Fert, and X. Waintal, “Spin Torque and Waviness in Magnetic Multilayers: A Bridge between Valet-Fert Theory and Quantum Approaches,” *Phys. Rev. Lett.*, vol. 103, p. 066602, 2009. doi: [10.1103/PhysRevLett.103.066602](https://doi.org/10.1103/PhysRevLett.103.066602).
- [181] S. Borlenghi, V. Rychkov, C. Petitjean, and X. Waintal, “Multiscale Approach to Spin Transport in Magnetic Multilayers,” *Phys. Rev. B*, vol. 84, p. 035412, 2011. doi: [10.1103/PhysRevB.84.035412](https://doi.org/10.1103/PhysRevB.84.035412).
- [182] M. Bendra, S. Fiorentini, J. Ender, R. L. de Orio, T. Hadámek *et al.*, “Spin Transfer Torques in Ultra-Scaled MRAM Cells,” in *2022 45th Jubilee International Convention on Information, Communication and Electronic Technology (MIPRO)*, 2022, pp. 129–132. doi: [10.23919/MIPRO55190.2022.9803736](https://doi.org/10.23919/MIPRO55190.2022.9803736).
- [183] K. Y. Camsari, S. Ganguly, D. Datta, and S. Datta, “Physics-Based Factorization of Magnetic Tunnel Junctions for Modeling and Circuit Simulation,” in *2014 IEEE International Electron Devices Meeting (IEDM)*, 2014, pp. 35.6.1–35.6.4. doi: [10.1109/IEDM.2014.7047163](https://doi.org/10.1109/IEDM.2014.7047163).
- [184] S. Zhang and Z. Li, “Roles of Nonequilibrium Conduction Electrons on the Magnetization Dynamics of Ferromagnets,” *Physical Review Letters*, vol. 93, p. 127204, 2004. doi: [10.1103/PhysRevLett.93.127204](https://doi.org/10.1103/PhysRevLett.93.127204).
- [185] M. Wang, W. Cai, K. Cao, J. Zhou, J. Wrona *et al.*, “Current-induced magnetization switching in atom-thick tungsten engineered perpendicular magnetic tunnel junctions with large tunnel magnetoresistance,” *Nature Communications*, vol. 9, p. 671, 2018. doi: [10.1038/s41467-018-03140-z](https://doi.org/10.1038/s41467-018-03140-z).
- [186] M. Wang, “Tunnel Junction with Perpendicular Magnetic Anisotropy: Status and Challenges,” *Micromachines*, vol. 6, no. 8, pp. 1023–1045, 2015. doi: [10.3390/mi6081023](https://doi.org/10.3390/mi6081023).
- [187] J. M. Iwata-Harms *et al.*, “Ultrathin perpendicular magnetic anisotropy CoFeB free layers for highly efficient, high speed writing in spin-transfer-torque MRAM,” *Scientific Reports*, vol. 9, p. 18668, 2019. doi: [10.1038/s41598-019-54466-7](https://doi.org/10.1038/s41598-019-54466-7).
- [188] M. Bendra, S. Fiorentini, J. Ender, R. L. de Orio, T. Hadámek *et al.*, “Back-Hopping in Ultra-Scaled MRAM Cells,” in *2023 46th MIPRO ICT and Electronics Convention (MIPRO)*, 2023. doi: [10.23919/MIPRO57284.2023.10159788](https://doi.org/10.23919/MIPRO57284.2023.10159788).
- [189] K. Watanabe, B. Jinnai, S. Ukami, H. Sato, and H. Ohno, “Shape Anisotropy Revisited in Single-Digit Nanometer Magnetic Tunnel Junctions,” *Nature Communications*, vol. 9, p. 663, 2018. doi: [10.1038/s41467-018-03003-7](https://doi.org/10.1038/s41467-018-03003-7).
- [190] B. Jinnai, J. Igarashi, K. Watanabe, T. Funatsu, H. Sato *et al.*, “High-Performance Shape-Anisotropy Magnetic Tunnel Junctions Down to 2.3 nm,” in *IEEE International Electron Devices Meeting (IEDM)*, 2020, pp. 24.6.1–24.6.4. doi: [10.1109/IEDM13553.2020.9371972](https://doi.org/10.1109/IEDM13553.2020.9371972).

- [191] K. Nishioka, H. Honjo, S. Ikeda, T. Watanabe, S. Miura *et al.*, “Novel Quad Interface MTJ Technology and its First Demonstration with High Thermal Stability and Switching Efficiency for STT-MRAM Beyond 2Xnm,” in *Symposium on VLSI Technology*, 2019, pp. T120–T121. doi: [10.23919/VLSIT.2019.8776499](https://doi.org/10.23919/VLSIT.2019.8776499).
- [192] K. Nishioka, H. Honjo, H. Naganuma, T. V. A. Nguyen, M. Yasuhira *et al.*, “Enhancement of Magnetic Coupling and Magnetic Anisotropy in MTJs with Multiple CoFeB/MgO Interfaces for High Thermal Stability,” *AIP Advances*, vol. 11, no. 2, p. 025231, 2021. doi: [10.1063/9.0000048](https://doi.org/10.1063/9.0000048).
- [193] M. Bendra, R. L. de Orio, S. Selberherr, W. Goes, and V. Sverdlov, “A Multi-Level Cell for Ultra-Scaled STT-MRAM Realized by Back-Hopping,” *Solid-State Electronics*, 2023. doi: [10.1016/j.sse.2023.108770](https://doi.org/10.1016/j.sse.2023.108770).
- [194] M. Bendra, R. L. de Orio, S. Selberherr, W. Goes, and V. Sverdlov, “A Multi-Level Cell for Ultra-Scaled STT-MRAM Realized by Back-Hopping,” *Solid-State Electronics*, vol. 212, p. 108830, 2024. doi: [10.1016/j.sse.2023.108830](https://doi.org/10.1016/j.sse.2023.108830). Extended version with Set B material parameters.
- [195] T. Devolder, O. Bultynck, P. Bouquin, V. D. Nguyen, S. Rao *et al.*, “Back Hopping in Spin Transfer Torque Switching of Perpendicularly Magnetized Tunnel Junctions,” *Physical Review B*, vol. 102, p. 184406, 2020. doi: [10.1103/PhysRevB.102.184406](https://doi.org/10.1103/PhysRevB.102.184406).
- [196] R. Phoomatna, S. Sampan-a pai, A. Meo, R. W. Chantrell, J. Chureemart *et al.*, “Dimensional Scaling Effects on Critical Current Density and Magnetization Switching in CoFeB-Based Magnetic Tunnel Junction,” *Journal of Physics D: Applied Physics*, vol. 57, no. 18, p. 185002, 2024. doi: [10.1088/1361-6463/ad1e53](https://doi.org/10.1088/1361-6463/ad1e53).
- [197] L. Farcis, B. M. S. Teixeira, P. Talatchian, D. Salomoni, U. Ebels, S. Auffret, B. Dieny, F. A. Mizrahi, J. Grollier, R. C. Sousa, and L. D. Buda-Prejbeanu, “Spiking Dynamics in Dual Free Layer Perpendicular Magnetic Tunnel Junctions,” *Nano Letters*, vol. 23, no. 17, pp. 7869–7875, 2023. doi: [10.1021/acs.nanolett.3c01597](https://doi.org/10.1021/acs.nanolett.3c01597).
- [198] R. Matsumoto, S. Lequeux, H. Imamura, and J. Grollier, “Chaos and Relaxation Oscillations in Spin-Torque Windmill Spiking Oscillators,” *Phys. Rev. Appl.*, vol. 11, p. 044093, 2019. doi: [10.1103/PhysRevApplied.11.044093](https://doi.org/10.1103/PhysRevApplied.11.044093).
- [199] M. Bendra, N. P. Jørstad, R. L. de Orio, S. Selberherr, W. Goes, and V. Sverdlov, “Unified Modeling of Ultra-Scaled STT-MRAM Cells: Harnessing Parasitic Effects for Enhanced Data Storage Dynamics,” in *2023 IEEE International Electron Devices Meeting (IEDM)*, 2023.
- [200] M. Bendra, S. Fiorentini, S. Selberherr, W. Goes, and V. Sverdlov, “Simulation of Spin-Torque and Magnetization Dynamics in STT-MRAM Multi-Level Cells,” in *International Symposium on Hysteresis Modeling and Micromagnetics (HMM)*, 2023.
- [201] C. Abert, H. Sepehri-Amin, F. Bruckner, C. Vogler, M. Hayashi *et al.*, “Back-Hopping in Spin-Transfer-Torque Devices: Possible Origin and Countermeasures,” *Physical Review Applied*, vol. 9, p. 054010, 2018. doi: [10.1103/PhysRevApplied.9.054010](https://doi.org/10.1103/PhysRevApplied.9.054010).
- [202] P. Bruno, “Theory of Interlayer Exchange Interactions in Magnetic Multilayers,” *Journal of Physics: Condensed Matter*, vol. 11, no. 48, 1999. doi: [10.1088/0953-8984/11/48/305](https://doi.org/10.1088/0953-8984/11/48/305).

- [203] M. Bendra, R. L. de Orio, S. Selberherr, W. Goes, and V. Sverdlov, "Influence of Interface Exchange Coupling in Multilayered Spintronic Structures," in *2024 47th MIPRO ICT and Electronics Convention (MIPRO)*, 2024, pp. 1579–1583. doi: [10.1109/MIPRO60963.2024.10569798](https://doi.org/10.1109/MIPRO60963.2024.10569798).
- [204] T. Katayama, S. Yuasa, J. Velez, M. Y. Zhuravlev, S. S. Jaswal *et al.*, "Interlayer Exchange Coupling in Fe/MgO/Fe Magnetic Tunnel Junctions," *Applied Physics Letters*, vol. 89, no. 11, p. 112503, 2006. doi: [10.1063/1.2349321](https://doi.org/10.1063/1.2349321).
- [205] Y.-C. Weng, C.-W. Cheng, and G. Chern, "Interlayer Exchange Coupling and Perpendicular Magnetic Anisotropy in Co₄₀Fe₄₀B₂₀/MgO/Co₂₀Fe₆₀B₂₀ Tunnel Junction Structures," *IEEE Transactions on Magnetics*, vol. 49, pp. 4425–4428, 07 2013. doi: [10.1109/TMAG.2013.2245308](https://doi.org/10.1109/TMAG.2013.2245308).
- [206] B. Khodadadi, J. B. Mohammadi, J. M. Jones, A. Srivastava, C. Mewes, T. Mewes, and C. Kaiser, "Interlayer Exchange Coupling in Asymmetric Co–Fe/Ru/Co–Fe Trilayers Investigated with Broadband Temperature-Dependent Ferromagnetic Resonance," *Phys. Rev. Appl.*, vol. 8, p. 014024, Jul 2017. doi: [10.1103/PhysRevApplied.8.014024](https://doi.org/10.1103/PhysRevApplied.8.014024).
- [207] A. Mouhoub, F. Millo, C. Chappert, J.-V. Kim, J. Létang, A. Solignac, and T. Devolder, "Exchange Energies in CoFeB/Ru/CoFeB Synthetic Antiferromagnets," *Phys. Rev. Mater.*, vol. 7, p. 044404, Apr 2023. doi: [10.1103/PhysRevMaterials.7.044404](https://doi.org/10.1103/PhysRevMaterials.7.044404).
- [208] M. Bendra, R. L. de Orio, S. Selberherr, W. Goes, and V. Sverdlov, "Interlayer Exchange Coupling for Enhanced Performance in Spin-Transfer Torque MRAM Devices," *Solid-State Electronics*, vol. 229, p. 109179, 2025. doi: [10.1016/j.sse.2025.109179](https://doi.org/10.1016/j.sse.2025.109179).
- [209] A. Manchon, S. Zhang, and K.-J. Lee, "Signatures of Asymmetric and Inelastic Tunneling on the Spin Torque Bias Dependence," *Phys. Rev. B*, vol. 82, p. 174420, 2010. doi: [10.1103/PhysRevB.82.174420](https://doi.org/10.1103/PhysRevB.82.174420).
- [210] S.-C. Oh, S.-Y. Park, A. Manchon, M. Chshiev, J.-H. Han *et al.*, "Bias-Voltage Dependence of Perpendicular Spin-Transfer Torque in Asymmetric MgO-Based Magnetic Tunnel Junctions," *Nature Physics*, vol. 5, no. 12, pp. 898–902, 2009. doi: [10.1038/nphys1427](https://doi.org/10.1038/nphys1427).
- [211] T. Warnatz, B. E. Skovdal, F. Magnus, H. Stopfel, D. Primetzhofer, A. Stein, R. Brucas, and B. Hjörvarsson, "The Influence of Diameter on the magnetic Saturation in Fe₈₄Cu₁₆/MgO [001] multilayered islands," *Journal of Magnetism and Magnetic Materials*, vol. 496, p. 165864, 2020. doi: <https://doi.org/10.1016/j.jmmm.2019.165864>.
- [212] S. Bin Hamid, R. Dutta, O. Hassan, and M. Z. Baten, "Implementing Bidirectional Logic with Backhopping in Magnetic Tunnel Junctions," *AIP Advances*, vol. 14, no. 2, p. 025224, 02 2024. doi: [10.1063/5.0169751](https://doi.org/10.1063/5.0169751).
- [213] T. Devolder, J.-V. Kim, F. Garcia-Sanchez, J. Swerts, W. Kim, S. Couet, G. Kar, and A. Furnemont, "Time-Resolved Spin-Torque Switching in MgO-Based Perpendicularly Magnetized Tunnel Junctions," *Phys. Rev. B*, vol. 93, p. 024420, Jan 2016. doi: [10.1103/PhysRevB.93.024420](https://doi.org/10.1103/PhysRevB.93.024420).
- [214] A. Le Goff, K. Garcia, N. Vernier, T. Tahmasebi, S. Cornelissen, T. Min, and T. Devolder, "Effect of Ta Insertion in Reference Layers of MTJs With Perpendicular Anisotropy," *IEEE Transactions on Magnetics*, vol. 50, no. 11, pp. 1–4, 2014. doi: [10.1109/TMAG.2014.2328664](https://doi.org/10.1109/TMAG.2014.2328664).

- [215] M. Bendra, R. L. de Orio, W. Goes, S. Selberherr, and V. Sverdlov, “Investigating Reliability Issues in Multi-Layered STT-MRAM with Synthetic Antiferromagnets,” in *2024 IEEE International Symposium on the Physical and Failure Analysis of Integrated Circuits (IPFA)*, 2024, pp. 1–5. doi: [10.1109/IPFA61654.2024.10690971](https://doi.org/10.1109/IPFA61654.2024.10690971).
- [216] M. Bendra, W. Goes, S. Selberherr, and V. Sverdlov, “Simulation of SAF-enhanced multilayered STT-MRAM structures,” *Microelectronic Engineering*, vol. 302, p. 112426, 2026. doi: [10.1016/j.mee.2025.112426](https://doi.org/10.1016/j.mee.2025.112426).
- [217] M. Bendra, S. Selberherr, W. Goes, and V. Sverdlov, “Enhancing Sub-Nanosecond Magnetic Tunnel Junctions with Double Spin Torque and Synthetic Antiferromagnetic Layers,” in *International Symposium on Hysteresis Modeling and Micromagnetics (HMM)*, 2025.
- [218] M. Bendra, B. Pruckner, R. L. de Orio, S. Selberherr, W. Goes *et al.*, “Advancing Nonvolatile Memory Technologies: Enhancing Reliability and Performance through Double Spin Torque Magnetic Tunnel Junctions and Interlayer Exchange Coupling,” in *2024 Device Research Conference (DRC)*, 2024, pp. 1–2. doi: [10.1109/DRC61706.2024.10605512](https://doi.org/10.1109/DRC61706.2024.10605512).
- [219] B. Pruckner, S. Fiorentini, W. Goes, S. Selberherr, and V. Sverdlov, “Micromagnetic Modeling of Double Spin-Torque Magnetic Tunnel Junction Devices,” *Physica B: Condensed Matter*, vol. 688, p. 416124, 2024. doi: [10.1016/j.physb.2024.416124](https://doi.org/10.1016/j.physb.2024.416124).
- [220] J. Chen, S. Peng, D. Xiong, H. Cheng, H. Zhou *et al.*, “Correlation of Interfacial Perpendicular Magnetic Anisotropy and Interlayer Exchange Coupling in CoFe/W/CoFe Structures,” *Journal of Physics D: Applied Physics*, vol. 53, no. 33, p. 334001, 2020. doi: [10.1088/1361-6463/ab8bfe](https://doi.org/10.1088/1361-6463/ab8bfe).

List of Publications

Journal Articles

- [1] **M. Bendra**, W. Goes, S. Selberherr, and V. Sverdlov, "Simulation of SAF-Enhanced Multilayered STT-MRAM Structures," *Microelectronic Engineering*, vol. 302, p. 112426, 2026. doi: [10.1016/j.mee.2025.112426](https://doi.org/10.1016/j.mee.2025.112426).
- [2] **M. Bendra**, de Orio, Roberto Lacerda, S. Selberherr, W. Goes, and V. Sverdlov, "Interlayer Exchange Coupling for Enhanced Performance in Spin-Transfer Torque MRAM Devices," *Solid-State Electronics*, vol. 229, p. 109179, 2025. doi: [10.1016/j.sse.2025.109179](https://doi.org/10.1016/j.sse.2025.109179).
- [3] **M. Bendra**, de Orio, Roberto Lacerda, S. Selberherr, W. Goes, and V. Sverdlov, "A Multi-Level Cell for Ultra-Scaled STT-MRAM Realized by Back-Hopping," *Solid-State Electronics*, vol. 223, p. 109027, 2025. doi: [10.1016/j.sse.2024.109027](https://doi.org/10.1016/j.sse.2024.109027). Invited.
- [4] **M. Bendra**, de Orio, Roberto Lacerda, S. Selberherr, W. Goes, and V. Sverdlov, "Advanced Modeling and Simulation of Multilayer Spin-Transfer Torque Magnetoresistive Random Access Memory with Interface Exchange Coupling," *Micromachines*, vol. 15, no. 5, 2024. doi: [10.3390/mi15050568](https://doi.org/10.3390/mi15050568). Invited.
- [5] **M. Bendra**, de Orio, Roberto Lacerda, S. Selberherr, W. Goes, and V. Sverdlov, "A Multi-Level Cell for Ultra-Scaled STT-MRAM Realized by Back-Hopping," *Solid-State Electronics*, vol. 208, p. 108738, 2023. doi: [10.1016/j.sse.2023.108738](https://doi.org/10.1016/j.sse.2023.108738).
- [6] S. Fiorentini, N. P. Jørstad, J. Ender, de Orio, Roberto Lacerda, S. Selberherr, **M. Bendra**, W. Goes, and V. Sverdlov, "Finite Element Approach for the Simulation of Modern MRAM Devices," *Micromachines*, vol. 14, no. 5, 2023. doi: [10.3390/mi14050898](https://doi.org/10.3390/mi14050898). Invited.
- [7] **M. Bendra**, S. Fiorentini, W. Goes, S. Selberherr, and V. Sverdlov, "The Influence of Interface Effects on the Switching Behavior in Ultra-Scaled MRAM Cells," *Solid-State Electronics*, vol. 201, p. 108590, 2023. doi: [10.1016/j.sse.2023.108590](https://doi.org/10.1016/j.sse.2023.108590). Invited.
- [8] **M. Bendra**, S. Fiorentini, W. Goes, S. Selberherr, and V. Sverdlov, "Interface Effects in Ultra-Scaled MRAM Cells," *Solid-State Electronics*, vol. 194, p. 108373, 2022. doi: [10.1016/j.sse.2022.108373](https://doi.org/10.1016/j.sse.2022.108373).
- [9] T. Hadámek, S. Fiorentini, **M. Bendra**, J. Ender, de Orio, Roberto Lacerda, W. Goes, S. Selberherr, and V. Sverdlov, "Temperature Increase in STT-MRAM at Writing: A Fully Three-Dimensional Finite Element Approach," *Solid-State Electronics*, vol. 193, p. 108269, 2022. doi: [10.1016/j.sse.2022.108269](https://doi.org/10.1016/j.sse.2022.108269). Invited.

- [10] S. Fiorentini, **M. Bendra**, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. Goes, S. Selberherr, and V. Sverdlov, “Spin and Charge Drift-Diffusion in Ultra-Scaled MRAM Cells,” *Scientific Reports*, vol. 12, p. 20958, 2022. doi: [10.1038/s41598-022-25586-4](https://doi.org/10.1038/s41598-022-25586-4).
- [11] J. Ender, S. Fiorentini, de Orio, Roberto Lacerda, T. Hadámek, **M. Bendra**, W. Goes, S. Selberherr, and V. Sverdlov, “Advances in Modeling Emerging Magnetoresistive Random Access Memories: From Finite Element Methods to Machine Learning Approaches,” *Proceedings of SPIE*, vol. 12157, pp. 1 215 708–1–1 215 708–14, 2022. doi: [10.1117/12.2624595](https://doi.org/10.1117/12.2624595). Invited.

Conference Contributions

- [12] V. Sverdlov, N. P. Jørstad, B. Pruckner, **M. Bendra**, W. Goes, and S. Selberherr, “Emerging Magnetoresistive Memories,” in *2025 IEEE 16th International Conference on ASIC (ASICON)*, Kunming, China, 2025, pp. 1–4. doi: [10.1109/ASICON66040.2025.11326380](https://doi.org/10.1109/ASICON66040.2025.11326380).
- [13] V. Sverdlov, B. Pruckner, N. P. Jørstad, **M. Bendra**, S. Selberherr, and W. Goes, “Modeling Advanced Magnetoresistive Memories,” in *2025 IEEE European Solid-State Electronics Research Conference (ESSERC)*, Munich, Germany, 2025, pp. 365–368. doi: [10.1109/ESSERC66193.2025.11214099](https://doi.org/10.1109/ESSERC66193.2025.11214099).
- [14] **M. Bendra**, S. Selberherr, W. Goes, and V. Sverdlov, “Enhancing Sub-Nanosecond Magnetic Tunnel Junctions with Double Spin Torque and Synthetic Antiferromagnetic Layers,” in *International Symposium on Hysteresis Modeling and Micromagnetics (HMM)*, 2025.
- [15] **M. Bendra**, de Orio, Roberto Lacerda, S. Selberherr, W. Goes, and V. Sverdlov, “Influence of Interface Exchange Coupling in Multilayered Spintronic Structures,” in *2024 47th MIPRO ICT and Electronics Convention (MIPRO)*, 2024, pp. 1579–1583. doi: [10.1109/MIPRO60963.2024.10569798](https://doi.org/10.1109/MIPRO60963.2024.10569798).
- [16] **M. Bendra**, B. Pruckner, de Orio, Roberto Lacerda, S. Selberherr, W. Goes, and V. Sverdlov, “Advancing Nonvolatile Memory Technologies: Enhancing Reliability and Performance through Double Spin Torque Magnetic Tunnel Junctions and Interlayer Exchange Coupling,” in *2024 Device Research Conference (DRC)*, 2024, pp. 1–2. doi: [10.1109/DRC61706.2024.10605512](https://doi.org/10.1109/DRC61706.2024.10605512).
- [17] **M. Bendra**, W. Goes, S. Selberherr, and V. Sverdlov, “Simulation of SAF-Enhanced Multilayered STT-MRAM Structures,” in *2024 Austrochip Workshop on Microelectronics (Austrochip)*, 2024, pp. 1–4. doi: [10.1109/Austrochip62761.2024.10716241](https://doi.org/10.1109/Austrochip62761.2024.10716241).
- [18] **M. Bendra**, de Orio, Roberto Lacerda, W. Goes, S. Selberherr, and V. Sverdlov, “Investigating Reliability Issues in Multi-Layered STT-MRAM with Synthetic Antiferromagnets,” in *2024 IEEE International Symposium on the Physical and Failure Analysis of Integrated Circuits (IPFA)*, 2024, pp. 1–5. doi: [10.1109/IPFA61654.2024.10690971](https://doi.org/10.1109/IPFA61654.2024.10690971).

- [19] B. Pruckner, N. P. Jørstad, **M. Bendra**, T. Hadámek, W. Goes, S. Selberherr, and V. Sverdlov, “Simulation of Advanced MRAM Devices for Sub-ns Switching,” in *2024 International Conference on Simulation of Semiconductor Processes and Devices (SISPAD)*, 2024, pp. 1–4. doi: [10.1109/SISPAD62626.2024.10733317](https://doi.org/10.1109/SISPAD62626.2024.10733317).
- [20] **M. Bendra**, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, N. P. Jørstad, B. Pruckner, S. Selberherr, W. Goes, and V. Sverdlov, “Back-Hopping in Ultra-Scaled MRAM Cells,” in *46th MIPRO ICT and Electronics Convention (MIPRO)*, 2023, pp. 159–162. doi: [10.23919/MIPRO57284.2023.10159764](https://doi.org/10.23919/MIPRO57284.2023.10159764).
- [21] **M. Bendra**, S. Fiorentini, S. Selberherr, W. Goes, and V. Sverdlov, “A Multi-Level Cell for Ultra-Scaled STT-MRAM Realized by Back-Hopping,” in *9th Joint International EuroSOI Workshop and International Conference on Ultimate Integration on Silicon (EuroSOI-ULIS)*, Tarragona, Spain, 2023.
- [22] **M. Bendra**, S. Fiorentini, S. Selberherr, W. Goes, and V. Sverdlov, “Simulation of Spin-Torque and Magnetization Dynamics in STT-MRAM Multi-Level Cells,” in *International Symposium on Hysteresis Modeling and Micromagnetics (HMM)*, 2023.
- [23] **M. Bendra**, de Orio, Roberto Lacerda, W. Goes, V. Sverdlov, and S. Selberherr, “Modeling of Ultra-Scaled Magnetoresistive Random Access Memory,” in *International Conference on Microelectronic Devices and Technologies (MicDAT)*, Funchal, Portugal, 2023, pp. 28–30. ISBN 978-84-09-53748-8
- [24] S. Fiorentini, de Orio, Roberto Lacerda, J. Ender, S. Selberherr, **M. Bendra**, N. P. Jørstad, W. Goes, and V. Sverdlov, “Finite Element Method for MRAM Switching Simulations,” in *Proc. of the International Conference on Mathematical Models, Computational Techniques in Science, and Engineering (MMCTSE)*, Athens, Greece, 2023, invited.
- [25] N. P. Jørstad, T. Hadámek, **M. Bendra**, J. Ender, B. Pruckner, W. Goes, and V. Sverdlov, “Numerical Simulations of Spintronic Magnetoresistive Memories,” in *Silvaco Users Global Event (SURGE)*, Santa Clara, USA, 2023, invited.
- [26] V. Sverdlov, S. Fiorentini, **M. Bendra**, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, N. P. Jørstad, W. Goes, and S. Selberherr, “Comprehensive Modeling of Nanoscale Magnetoresistive Memory,” in *3rd Global Experts Conference on Materials Science & Nanotechnology (GECMSN23)*, Dubai, UAE, 2023, invited.
- [27] V. Sverdlov, **M. Bendra**, W. Goes, S. Fiorentini, A. Garcia-Barrientos, and S. Selberherr, “Multi-Level Operation in Ultra-Scaled MRAM,” in *IEEE Latin American Electron Devices Conference (LAEDC)*, 2023, pp. 1–4. doi: [10.1109/LAEDC58183.2023.10209117](https://doi.org/10.1109/LAEDC58183.2023.10209117).
- [28] V. Sverdlov, **M. Bendra**, B. Pruckner, S. Fiorentini, W. Goes, and S. Selberherr, “Comprehensive Modeling of Advanced Composite Magnetoresistive Devices,” in *IEEE 53rd European Solid-State Device Research Conference (ESSDERC)*, 2023, pp. 93–96. doi: [10.1109/ESSDERC59256.2023.10268508](https://doi.org/10.1109/ESSDERC59256.2023.10268508).

- [29] **M. Bendra**, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. J. Loch, N. P. Jørstad, S. Selberherr, W. Goes, and V. Sverdlov, “Spin Transfer Torques in Ultra-Scaled MRAM Cells,” in *45th Jubilee International Convention on Information, Communication and Electronic Technology (MIPRO)*, 2022, pp. 129–132. doi: [10.23919/MIPRO55190.2022.9803736](https://doi.org/10.23919/MIPRO55190.2022.9803736).
- [30] **M. Bendra**, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. J. Loch, N. P. Jørstad, W. Goes, S. Selberherr, and V. Sverdlov, “Interface Effects in Ultra-Scaled MRAM Cells,” in *8th Joint International EuroSOI Workshop and International Conference on Ultimate Integration on Silicon (EuroSOI-ULIS)*, Udine, Italy, 2022, pp. 1–2.
- [31] **M. Bendra**, W. J. Loch, N. P. Jørstad, S. Fiorentini, S. Selberherr, W. Goes, and V. Sverdlov, “Modeling Ultra-Scaled Multi-Layer STT-MRAM Cells: A Unified Spin and Charge Drift-Diffusion Approach,” in *2022 IEEE International Electron Devices Meeting (IEDM)*, 2022, poster.
- [32] S. Fiorentini, **M. Bendra**, J. Ender, W. Goes, V. Sverdlov, and S. Selberherr, “Evaluating Spin Transfer Torques in Multilayered Magnetic Tunnel Junctions and Spin Valves,” in *Workshop on Innovative Nanoscale Devices and Systems (WINDS)*, Lihue, HI, USA, 2022.
- [33] S. Fiorentini, **M. Bendra**, J. Ender, T. Hadámek, W. J. Loch, N. P. Jørstad, de Orio, Roberto Lacerda, W. Goes, S. Selberherr, and V. Sverdlov, “Modeling Advanced Spintronic-Based Magnetoresistive Memory,” in *International Conference on Microwave & THz Technologies, Wireless Communications and OptoElectronics (IRPhE 2022)*, 2022, pp. 49–52. doi: [10.1049/icp.2022.2795](https://doi.org/10.1049/icp.2022.2795).
- [34] S. Fiorentini, **M. Bendra**, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. J. Loch, N. P. Jørstad, W. Goes, S. Selberherr, and V. Sverdlov, “Spin Torques in ULTRA-Scaled MRAM Devices,” in *IEEE 52nd European Solid-State Device Research Conference (ESSDERC)*, 2022, pp. 348–351. doi: [10.1109/ESSDERC55479.2022.9947196](https://doi.org/10.1109/ESSDERC55479.2022.9947196).
- [35] S. Fiorentini, W. J. Loch, **M. Bendra**, N. P. Jørstad, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. Goes, V. Sverdlov, and S. Selberherr, “Design Analysis of Ultra-Scaled MRAM Cells,” in *IEEE 16th International Conference on Solid-State & Integrated Circuit Technology (ICSICT)*, 2022, pp. 1–4. doi: [10.1109/ICSICT55466.2022.9963202](https://doi.org/10.1109/ICSICT55466.2022.9963202).
- [36] T. Hadámek, **M. Bendra**, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, W. Goes, S. Selberherr, and V. Sverdlov, “Temperature Increase in MRAM at Writing: A Finite Element Approach,” in *7th Joint International EuroSOI Workshop and International Conference on Ultimate Integration on Silicon (EuroSOI-ULIS)*, 2021, pp. 1–4. doi: [10.1109/EuroSOI-ULIS53016.2021.9560669](https://doi.org/10.1109/EuroSOI-ULIS53016.2021.9560669).
- [37] **M. Bendra**, J. Ender, S. Fiorentini, T. Hadámek, de Orio, Roberto Lacerda, W. Goes, V. Sverdlov, and S. Selberherr, “Finite Element Method Approach to MRAM Modeling,” in *44th International Convention on Information, Communication and Electronic Technology (MIPRO)*, 2021, pp. 70–73. doi: [10.23919/MIPRO52101.2021.9597194](https://doi.org/10.23919/MIPRO52101.2021.9597194).

- [38] J. Ender, S. Fiorentini, de Orio, Roberto Lacerda, T. Hadámek, **M. Bendra**, W. Goes, S. Selberherr, and V. Sverdlov, “Advanced Modeling of Emerging MRAM: From Finite Element Methods to Machine Learning Approaches,” in *Proc. of the International Conference Micro-and Nanoelectronics (ICMNE)*, 2021.
- [39] V. Sverdlov, **M. Bendra**, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. J. Loch, N. P. Jørstad, W. Goes, and S. Selberherr, “Advanced Modeling of Emerging Devices for Digital Spintronics,” in *Proc. of the International Conference on Nanoscience and Nanotechnology*, Dubai, UAE, 2022, invited.
- [40] V. Sverdlov, **M. Bendra**, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. J. Loch, N. P. Jørstad, W. Goes, and S. Selberherr, “Advanced Modeling of Emerging Magnetoresistive Memory,” in *Proc. of the International Meet & Expo on Nanotechnology (NANOMEET)*, Edinburgh, UK, 2022, invited.
- [41] V. Sverdlov, **M. Bendra**, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. J. Loch, N. P. Jørstad, W. Goes, and S. Selberherr, “Emerging Devices for Digital Spintronics,” in *2nd Global Conference & Expo on Nanotechnology & Nanoscience*, 2022, pp. 32–33, invited.
- [42] V. Sverdlov, **M. Bendra**, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. J. Loch, N. P. Jørstad, W. Goes, and S. Selberherr, “Modeling Advanced Magnetoresistive Memory: A Journey from Finite Element Methods to Machine Learning Approaches,” in *2nd Global Webinar on Nanoscience & Nanotechnology*, 2022, invited.
- [43] V. Sverdlov, W. J. Loch, **M. Bendra**, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, N. P. Jørstad, W. Goes, and S. Selberherr, “Modeling Approach to Ultra-Scaled MRAM Cells,” in *Book of Abstracts of the International Meet on Applied Science, Engineering and Technology (ASETMEET)*, 2022, pp. 7–8, invited.
- [44] W. J. Loch, S. Fiorentini, N. P. Jørstad, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, **M. Bendra**, W. Goes, S. Selberherr, and V. Sverdlov, “Double Reference Layer STT-MRAM Structures with Improved Performance,” in *8th Joint International EuroSOI Workshop and International Conference on Ultimate Integration on Silicon (EuroSOI-ULIS)*, Udine, Italy, 2022, pp. 1–2.
- [45] N. P. Jørstad, S. Fiorentini, J. Ender, de Orio, Roberto Lacerda, T. Hadámek, W. J. Loch, **M. Bendra**, W. Goes, S. Selberherr, and V. Sverdlov, “Finite Element Modeling of Spin-Orbit Torques,” in *Proc. of the 8th Joint International EuroSOI Workshop and International Conference on Ultimate Integration on Silicon (EuroSOI-ULIS)*, Udine, Italy, 2022.