

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

DynaMag: Accelerator-Portable Finite-Element Micromagnetics on Dynamic Meshes with Electromagnetic Coupling

Permalink

<https://escholarship.org/uc/item/05z7w0wt>

ISBN

9798171611149

Author

Duan, Jiawei

Publication Date

2026-10-23

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

**DynaMag: Accelerator-Portable Finite-Element Micromagnetics on Dynamic Meshes with
Electromagnetic Coupling**

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Materials Science and Engineering

by

Jiawei Duan

Committee in charge:

Professor Vitaliy Lomakin, Chair
Professor Eric Fullerton
Professor Marko Lubarda
Professor Daniel Sievenpiper
Professor Kesong Yang

Copyright

Jiawei Duan, 2026

All rights reserved.

The Dissertation of Jiawei Duan is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2026

TABLE OF CONTENTS

DISSERTATION APPROVAL NOTICE	iii
TABLE OF CONTENTS	iv
LIST OF FIGURES	vii
LIST OF TABLES	viii
ACKNOWLEDGEMENTS	x
VITA	xii
ABSTRACT OF THE DISSERTATION	xiii
1 Introduction	1
1.1 Dynamic micromagnetic devices	2
1.2 Research objectives	4
1.3 Contributions	5
1.4 Dissertation structure	5
2 Micromagnetic and Electromagnetic Foundations	8
2.1 Magnetization dynamics	8
2.2 Local, bulk, and interfacial effective fields	11
2.2.1 Bulk exchange	11
2.2.2 Surface and intergranular exchange	12
2.2.3 Uniaxial anisotropy	13
2.2.4 Cubic anisotropy	14
2.2.5 Applied field	14
2.3 Demagnetizing field: the global micromagnetic bottleneck	15
2.4 Eddy-current reaction field	17
2.5 Finite-element and time discretization	18
2.6 Rigid motion, morphing, and operator validity	20
2.7 Fast long-range evaluation and verification hierarchy	21
3 DynaMag: Accelerator-Portable Micromagnetics with Dynamic State	23
3.1 DynaMag architecture	23
3.2 The SzCompute portable numerical substrate	25
3.2.1 Matrix-free finite-element action	27
3.3 DynaMag state and operator design	30
3.4 Reference-to-current kinematics and material response	32
3.5 Accepted-state evolution on accelerators	35
3.6 Verification of dynamic geometry, frames, and materials	35

3.7	Dynamic LLG cost and scaling	38
4	Hybrid Superposition Integral–Poisson Magnetostatics	42
4.1	Hybrid integral–Poisson method	43
4.1.1	Problem formulation	43
4.1.2	Surface potential and interior Poisson equation	46
4.2	Numerical demonstration	48
4.2.1	Cylinder validation and refinement study	49
4.2.2	Multiscale write head	50
4.3	On-the-fly DynaMag extension	51
5	Periodic Green Functions and B-AIM Acceleration	55
5.1	Conventions and the static-limit construction	57
5.1.1	Normalization, displacement, and sign	57
5.1.2	Differentiation before the static limit	58
5.2	Periodic Green functions by dimensionality	58
5.2.1	1D periodicity	58
5.2.2	2D periodicity	60
5.2.3	3D periodicity	62
5.3	Numerical realization and singular interactions	64
5.3.1	Production truncation and symmetry branches	64
5.3.2	Neutrality: scalar charge, closed current, and periodic transport	65
5.3.3	Self value, panel limit, and near replacement	67
5.4	Validation hierarchy and error separation	67
5.5	B-AIM performance and storage scaling	70
5.5.1	Benchmark definition and protocol	70
5.5.2	Measured performance	71
6	Airbox-Free Eddy-Current FEM–BEM and Coupled LLG–Maxwell Dynamics	75
6.1	Airbox-free eddy-current formulation	76
6.2	Classical A– k FEM–BEM formulation	77
6.2.1	Maxwell reduction and the A– ϕ equation	78
6.2.2	Interface equations	79
6.2.3	Weak form and the two Green identities	79
6.2.4	BEM closure and trace maps	80
6.2.5	From the two weak equations to four blocks	81
6.2.6	BDF time discretization	82
6.2.7	Free-space hollow-sphere validation	83
6.3	Intermediate field-oriented H– k FEM–BEM formulation	84
6.3.1	Maxwell reduction and the H equation	84
6.3.2	Interface equations and compatible spaces	85
6.3.3	Weak form and the curl Green identity	87
6.3.4	BEM closure and boundary traces	87
6.3.5	Surface and volume discretization	89
6.3.6	From the two weak equations to four blocks	91

6.3.7	BDF time discretization	92
6.3.8	Free-space hollow-sphere validation	94
6.4	Final scalar-trace H - κ formulation	95
6.4.1	Maxwell reduction and the H equation	96
6.4.2	Interface equations and the constrained trace	97
6.4.3	Weak form and elimination of the air volume	97
6.4.4	BEM closure and the κ coordinate	99
6.4.5	From the weak equations to four blocks	101
6.4.6	BDF time discretization	102
6.4.7	Free-space operator and hollow-sphere validation	103
6.4.8	Periodic H - κ boundary operator	104
6.5	Periodic H - κ cylinder and plate validation	105
6.6	Integrated LLG–Maxwell coupling	109
6.6.1	Coupled physical model	109
6.6.2	Implemented accepted-step coupling	110
6.6.3	Integrated validation ladder	113
6.6.4	Conducting-disk coupled validation	114
6.6.5	Interpretation and remaining integration	117
7	Matrix-Free FEM–BEM Performance Optimization	119
7.1	Matrix-free BEM operators	119
7.2	Corrected and prepared matrix-free execution	121
7.2.1	Matched near replacement and linear storage	121
7.2.2	Prepared ownership, mixed precision, and time integration	122
7.3	TEAM11 runtime and memory comparison	122
7.4	Block preconditioning and solver selection	124
7.4.1	Block factorization and the operator that should be approximated	124
7.4.2	Volume edge block	125
7.4.3	Hollow-sphere evidence	126
7.4.4	Selection rules for coupled dynamics	127
8	Conclusions and Future Directions	128
8.1	Summary of results	128
8.2	Principal contributions	130
8.3	Future work	131
8.4	Summary	132
	REFERENCES	133

LIST OF FIGURES

Figure 3.1:	Applications of dynamic meshes	24
Figure 3.2:	DynaMag numerical architecture	25
Figure 3.3:	Dynamic material-parameter dependencies	34
Figure 3.4:	Non-rigid morph and material-frame transport	37
Figure 3.5:	Airy-disk morph and prescribed magnetization	40
Figure 3.6:	Dynamic LLG trajectory and cost decomposition	40
Figure 3.7:	End-to-end DynaMag cube scaling	41
Figure 4.1:	Geometry and near zones of the hybrid method	46
Figure 4.2:	Published cylinder validation	49
Figure 4.3:	Effect of a local offset on the near correction	51
Figure 4.4:	Published multiscale write-head model	52
Figure 4.5:	On-the-fly write-head timing	53
Figure 5.1:	Direct and analytical validation of the periodic $\mathbf{r}/r^3$ kernel	69
Figure 6.1:	Free-space $A-k$ hollow-sphere validation	83
Figure 6.2:	Continuous and discrete de Rham sequences for the conductor	86
Figure 6.3:	Free-space hollow-sphere $H-k$ validation	94
Figure 6.4:	Free-space $H-\kappa$ hollow-sphere validation	104
Figure 6.5:	Refined periodic $H-\kappa$ transient validation	107
Figure 6.6:	$H-\kappa$ -LLG conducting-disk switching result	115

LIST OF TABLES

Table 1.1:	Comparison of DynaMag with representative methods. The entries describe only the cited publications. Later extensions may add other capabilities.	6
Table 2.1:	Geometry updates required by the principal LLG fields.	21
Table 3.1:	Principal SzCompute services used by DynaMag.	26
Table 3.2:	Selected SzCompute backend measurements	27
Table 3.3:	Dataflow of the matrix-free P1 tetrahedral exchange-field action.	28
Table 3.4:	Matrix-free and assembled P1 tetrahedral exchange performance	29
Table 3.5:	State dependence under fixed-connectivity motion.	31
Table 3.6:	Evidence levels used for DynaMag capability claims.	36
Table 4.1:	Published cylinder refinement results	50
Table 5.1:	Zero-mode and neutrality distinction between scalar and gradient periodic kernels.	66
Table 5.2:	Independent audits of the periodic-gradient formulas.	68
Table 5.3:	Independent numerical evidence for the periodic-gradient kernel.	68
Table 5.4:	Prepared performance of the optimized scalar B-AIM path. Values after the plus-minus sign are sample standard deviations across five independent processes.	71
Table 5.5:	Prepared performance of the dedicated analytical $\mathbf{r}/r^3$ B-AIM action at 20 million vertices. Values after the plus-minus sign are sample standard deviations across five independent processes.	72
Table 5.6:	Scaling of the complete optimized free-space $1/r$ action. Setup time is excluded.	73
Table 5.7:	Device-array storage for the optimized 20-million-point scalar B-AIM calculation.	74
Table 6.1:	Free-space regularized $A-k$ hollow-sphere field errors.	84
Table 6.2:	Maxwell origin of the four time-discrete $H-k$ blocks.	92
Table 6.3:	Free-space hollow-sphere $H-k$ field errors.	95
Table 6.4:	Origin of the four $H-\kappa$ blocks.	103
Table 6.5:	$H-\kappa$ accelerated boundary action against the direct oracle.	103
Table 6.6:	$H-k$ and $H-\kappa$ hollow-sphere GMRES convergence	104

Table 6.7:	Boundary actions in the periodic $H\text{--}\kappa$ formulation.	105
Table 6.8:	Refined periodic $H\text{--}\kappa$ transient errors against analytical diffusion.	107
Table 6.9:	Validation status of the three FEM–BEM coordinates.	108
Table 6.10:	Integrated capabilities and their validation evidence.	114
Table 7.1:	Matrix-free boundary actions in the selected FEM–BEM coordinates.	120
Table 7.2:	DynaMag $A\text{--}k$ benchmark on the converted TEAM11 mesh. Timing excludes solution and axis-sample export.	123
Table 7.3:	Four-core LS-DYNA runtime and memory on the TEAM11 mesh. The 20 electromagnetic intervals cover 2 ms.	124
Table 7.4:	$A\text{--}k$ hollow-sphere preconditioner ablation.	126

ACKNOWLEDGEMENTS

First, I would like to thank my advisor, Professor Vitaliy Lomakin, for his great help and support during my entire Ph.D. He not only guided the direction of my research, but also helped me learn how to face problems better and become a better person. Through his own example, he showed me how a good attitude can help us work and live better. This work would not have been possible without his guidance and encouragement.

I would also like to thank all the members of my dissertation committee for their tremendous help and valuable feedback. The electromagnetics courses taught by Professor Vitaliy Lomakin and Professor Daniel Sievenpiper strengthened my interest in computational electromagnetics and helped me build a solid foundation for my future work. I still remember the days spent discussing research problems in their offices. I am grateful to Professor Eric Fullerton and Professor Kesong Yang, with whom I worked to improve high-performance ferrite and nanocrystalline materials. These collaborations helped me build my research skills and learn to maintain high standards in my work. Professor Marko Lubarda's extensive work on magnetic devices always gave me new perspectives when reviewing my work and strengthened my determination to pursue a better solver for demanding applications.

The work presented in this dissertation also received tremendous help from the members of our group. I would like to thank Dr. Fangzhou Ai for his help with high-performance computing and periodic structures, as well as for his suggestions on developing the next generation of solvers and pushing the boundaries of micromagnetics. I thank Egor for all our discussions about magnetic materials, including ferrites and nanocrystalline materials. I also thank Dr. Zhuonan Lin and Dr. Iana Volvach for their work, which built a solid foundation for micromagnetic modeling.

I still remember what I wrote in my application for the Ph.D. program: I wanted to create high-performance code and combine various modeling techniques to address multiphysics problems across different spectral ranges. I am happy that I kept this goal in mind and tried my best to make it happen. Through this work, I was able to bring multiphysics capabilities into the solver. Seeing

this happen after many years feels like a dream come true.

I would like to thank my parents for their trust and support while I pursued my Ph.D. I also thank my girlfriend, Yiqing Liu, for helping me through the difficult times.

Chapter 4, in part, is based on “Hybrid Superposition Integral–Poisson Solver Method for the Magnetostatic Field in Finite Element Micromagnetic Solvers” by Jiawei Duan, Xueyang Wang, and Vitaliy Lomakin, published in *IEEE Transactions on Magnetics*, vol. 59, no. 11, 2023, by IEEE, doi:10.1109/TMAG.2023.3297951. The dissertation author was the primary author of this paper.

Chapter 5, in part, is based on “Periodic Micromagnetic Finite Element Method” by Fangzhou Ai, Jiawei Duan, and Vitaliy Lomakin, published in the *Journal of Magnetism and Magnetic Materials*, vol. 615, 172776, 2025, by Elsevier, doi:10.1016/j.jmmm.2025.172776. The dissertation author was a contributing author of this paper.

VITA

- 2016 B.S. in Electrical Engineering, Xi'an Jiaotong University
- 2021 M.S. in Electrical and Computer Engineering, University of California San Diego
- 2026 Ph.D. in Materials Science and Engineering, University of California San Diego

PUBLICATIONS

- J. Duan, X. Wang, and V. Lomakin, "Hybrid superposition integral–Poisson solver method for the magnetostatic field in finite element micromagnetic solvers," *IEEE Transactions on Magnetics* 59 (11), 7002205 (2023).
- F. Ai, J. Duan, and V. Lomakin, "Periodic micromagnetic finite element method," *Journal of Magnetism and Magnetic Materials* 615, 172776 (2025).
- F. Ai, J. Duan, and V. Lomakin, "Spin-wave assisted synchronization in 2D arrays of spin torque oscillators," *Applied Physics Letters* 127 (16), 162401 (2025).
- F. Ai, Z. Lin, J. Duan, and V. Lomakin, "Periodic phase diagrams in micromagnetics with an eigenvalue solver," *IEEE Transactions on Magnetics* 61 (9), 7100205 (2025).

FIELDS OF STUDY

Major Field: Materials Science and Engineering

Studies in finite-element micromagnetics, computational electromagnetics, high-performance scientific computing, and accelerator-portable algorithms.

ABSTRACT OF THE DISSERTATION

DynaMag: Accelerator-Portable Finite-Element Micromagnetics on Dynamic Meshes with
Electromagnetic Coupling

by

Jiawei Duan

Doctor of Philosophy in Materials Science and Engineering
University of California San Diego, 2026

Professor Vitaliy Lomakin, Chair

Micromagnetic simulation supports research on magnetic materials and devices. New applications require magnetization to evolve with moving or morphing geometry, varying material properties, and induced currents. Existing solvers address only parts of this problem, limiting detailed analysis. The main obstacles are long-range magnetostatic interactions, the airbox required by eddy-current methods, and the cost of updating both operators as the system evolves.

These requirements call for a portable, memory-efficient compute foundation. SzCompute

provides common memory containers and numerical operations for multicore CPUs, NVIDIA and AMD GPUs, and Apple Metal. Its provider model keeps the simulation state on the selected device and uses unified memory when the hardware supports it. DynaMag builds its finite-element solvers on this foundation. It retains fixed topology and degree-of-freedom identity while coordinates, material frames, and properties change. Matrix-free local kernels and reusable long-range plans then avoid large geometry-dependent matrices. This design makes accelerator portability part of the solver rather than a separate implementation for each device.

On this foundation, the hybrid superposition integral–Poisson method addresses the magnetostatic field. It evaluates the free-space integral only on a truncation surface, after which a sparse Poisson solve recovers the interior field. B-AIM accelerates the nonlocal interaction, while matched local corrections retain finite-element accuracy. The point-to-point and finite-element corrections are evaluated as needed instead of being stored as matrices. The method therefore remains practical when the mesh changes and keeps the memory cost low.

The eddy-current field is treated by coupling finite elements in the conductor to boundary elements on its surface. Nédélec elements represent the conducting region, while the boundary operator represents the unbounded exterior. This construction removes the airbox, reduces the number of unknowns, and avoids exterior remeshing. The same B-AIM infrastructure applies the boundary operator without forming a dense matrix and supports both free-space and periodic Green functions. B-AIM thus provides a common long-range engine for the magnetostatic and eddy-current solvers.

Validation covers mesh motion, morphing, dynamic materials, magnetostatics, conducting transients, and magnetization–eddy-current coupling. These tests define a low-memory path toward fully dynamic finite-element micromagnetics on heterogeneous hardware.

1 Introduction

Magnetic materials are used in recording, sensing, power, and computing devices, where their behavior depends on interactions across several length and time scales. Exchange favors local alignment, while crystalline and interfacial anisotropies select preferred directions. Magnetostatics couples each part of a body to the full magnetic state, and applied fields or electric currents drive the system away from equilibrium. Micromagnetics replaces atomistic spins with a continuous magnetization field, while retaining these interactions and the structures that they produce (Brown, 1963; Aharoni, 2001).

The range of behavior covered by this continuum description is already visible in classical reduced problems. Coherent-particle hysteresis and field-driven motion of a 180° wall provide standard limits for reversal dynamics (Stoner and Wohlfarth, 1948; Schryer and Walker, 1974). At device scale, micromagnetics has been used to study domain-wall logic, perpendicular recording heads, and exchange-spring media (Allwood et al., 2005; Takano, 2005; Suess, 2007). Current-driven extensions connect spin-wave emission, spin-transfer torque, and nanowire wall motion to the broader spintronics framework (Berger, 1996; Slonczewski, 1996; Zhang and Li, 2004; Thiaville et al., 2005; Žutić et al., 2004). Thermal gradients add the related spin-caloritronic setting (Bauer et al., 2012). These examples motivate a solver architecture that can combine several excitations and material states without changing the underlying geometric representation.

Experiments increasingly probe magnetic systems whose geometry, material state, and electromagnetic environment evolve during operation. Examples include moving magnetic components, deformable structures, thermally driven media, and conducting devices with eddy-current

feedback. Periodic arrays introduce a related need to represent interactions across the surrounding lattice. Existing micromagnetic solvers are generally designed for fixed meshes and fixed material descriptions. Electromagnetic formulations often add a meshed exterior air region. These assumptions make motion, morphing, and coupled field feedback difficult to simulate efficiently. A solver for these experiments must evolve the magnetization, mesh, material data, and exterior electromagnetic field as one system. Representative published solvers and coupled formulations illustrate the fixed-grid, fixed-geometry, or bounded-domain assumptions being combined here (Chang et al., 2011; Vansteenkiste et al., 2014; Yao et al., 2022; Feischl and Tran, 2017).

This dissertation develops DynaMag, a finite-element micromagnetic and magnetoquasistatic solver for dynamic geometry and materials. Its mesh can move or deform, while material parameters and local axes can depend on time, temperature, or coupled fields. Between time steps, DynaMag retains accepted magnetic and electromagnetic states together with data that depend only on mesh topology. It then updates geometry- and material-dependent data when needed. This state model connects portable computing, low-memory magnetostatics, periodic Green functions, and airbox-free eddy-current coupling.

1.1 Dynamic micromagnetic devices

Modern micromagnetic solvers work well within their intended problem classes. Regular grids provide high throughput for simple domains, whereas finite-element methods represent curved, multimaterial, and topologically complex bodies. GPU implementations exist for both approaches: FastMag uses FEM (Chang et al., 2011), while MuMax3 uses a structured grid (Vansteenkiste et al., 2014). These solvers show how a fixed-geometry formulation can achieve high throughput. That performance does not extend automatically to dynamic geometry because motion invalidates data prepared for the original configuration. Even a rigid rotation may require updates to cached coordinates, element mappings, material directions, long-range plans, and preconditioners.

Two further tasks often surround a time-domain calculation. Coarse-graining connects the

resolved continuum scale to larger structures (Grinstein and Koch, 2003), while minimum-energy-path algorithms locate transition states between metastable configurations (Henkelman and Jónsson, 2000). Both rely on repeated, consistent evaluations of the same effective-field terms used during an LLG trajectory.

Some mesh data remain fixed during motion with fixed connectivity. These data include element incidence, degree-of-freedom identity, communication structure, and most algebraic graphs. In contrast, coordinates, Jacobians, quadrature data, surface geometry, and material frames can change. DynaMag therefore retains the fixed data and accepted history, but updates the remaining data on the active device. Before an operator runs, it also refreshes any long-range or preconditioning data that depend on geometry.

Material changes create the same issue without visible motion. Constitutive coefficients and local axes may depend on time, temperature, or coupled fields. Because the topology is unchanged, these changes do not require a new sparse graph. DynaMag instead treats constitutive data as runtime state, which local operators read at their current values. Rigid motion and general morphing then use the same fixed-connectivity state model.

Long-range fields are harder to update because, unlike exchange and anisotropy, they are not local once the mesh and material state are known. A coordinate change in a magnetostatic calculation alters source locations, surface geometry, interpolation weights, and the field throughout the domain. A dynamic solver should therefore avoid dense corrections tied to one geometry. Eddy-current calculations add both time history and an unbounded exterior domain. If that exterior is represented by an air mesh, the mesh must move or be regenerated with the device. Periodic currents may also cross an identified cell boundary. For these reasons, the demagnetizing and eddy-current fields are the hardest operators to update during motion.

1.2 Research objectives

This work extends an unstructured micromagnetic solver to dynamic geometry, materials, and Maxwell fields. To preserve physical consistency, local and long-range operators must use the same trial state, while accepted electromagnetic history must advance with the LLG state. The resulting solver must run on CPUs and accelerators without excessive memory growth. It must also preserve periodic conductor topology and represent the exterior field without an artificial airbox.

DynaMag separates invariant topology from evolving physical state. Its repeated operators use device-resident, matrix-free kernels. SzCompute provides the portable execution and memory model, while DynaMag supplies the current geometry, material frames, constitutive laws, and accepted history. Within this framework, the hybrid superposition integral–Poisson method reduces stored near-field data in finite-element magnetostatics (Duan et al., 2023). Periodic finite elements and Green functions replace a large supercell with one quotient cell (Ai et al., 2025a,b). For eddy currents, an accelerated FEM–BEM method uses finite elements only in the conductor and boundary operators for the unbounded exterior.

The periodic eddy-current solver uses the H – κ coordinate. It solves for $\mathbf{H}$ in the conductor and represents the exterior by a scalar single-layer density. In this representation, the scalar $1/r$ action reconstructs the boundary potential, while its analytical $\mathbf{r}/r^3$ gradient reconstructs the normal field without differentiating an interpolated potential. For 1D and 2D periodicity, DynaMag uses a finite-part convention that removes the divergent additive potential constant but retains the finite line or sheet field. Nédélec elements describe the conductor, and scalar boundary data describe the exterior, so the formulation requires no airbox.

DynaMag supports moving and morphing micromagnetic meshes, periodic operators, and coupled LLG–Maxwell transients. Before the solver can accept a moving-conductor state, it must refresh every geometry-dependent coefficient, trace, B-AIM object, local correction, and preconditioner.

1.3 Contributions

The first contribution is DynaMag’s compute and state model. SzCompute provides common execution paths for CPUs, NVIDIA and AMD GPUs, and Apple Metal. On this foundation, DynaMag preserves mesh and degree-of-freedom identity during rigid motion and morphing, while updating local material frames and dynamic parameters on the selected device.

The second contribution is a low-memory magnetostatic method for changing geometry. The hybrid superposition integral–Poisson method replaces a large volume correction with a boundary correction (Duan et al., 2023). DynaMag then streams the B-AIM point-to-point correction and evaluates the analytical FEM correction without storing either correction matrix.

The third contribution combines periodic finite-element micromagnetics with periodic long-range kernels. The scalar $1/r$ formulation gives periodic potentials and phase-dependent micromagnetic operators, while its Cartesian gradient gives the periodic $\mathbf{r}/r^3$ action and the required partial-period zero modes. Building on these kernels, the fourth contribution is an $H\text{--}\kappa$ FEM–BEM eddy-current solver with periodic quotient surfaces, B-AIM acceleration, and block preconditioning. The direct $H\text{--}k$ form provides an intermediate derivation, while free-space $A\text{--}k$ is a nonperiodic production alternative with the same conductor-only structure.

Table 1.1 compares DynaMag with representative methods. Each entry includes only features reported in the cited work.

The cited approaches do not combine all of these capabilities. DynaMag uses unstructured finite-element LLG with dynamic fixed-connectivity geometry and material state, together with portable computing, low-memory magnetostatics, periodic Green operators, and airbox-free eddy-current coupling.

1.4 Dissertation structure

Chapter 2 presents the physical and numerical background. After introducing the LLG equation and its effective fields, it identifies which operators change with the mesh and material

Table 1.1. Comparison of DynaMag with representative methods. The entries describe only the cited publications. Later extensions may add other capabilities.

Work	Geometry and magnetic dynamics	Compute and electromagnetic scope	Difference from DynaMag
FastMag (Chang et al., 2011)	Tetrahedral finite-element LLG for complex fixed geometry	GPU implementation; accelerated magnetostatic superposition	Unstructured GPU micromagnetics; the cited method does not combine evolving geometry with eddy-current Maxwell history
MuMax3 (Vansteenkiste et al., 2014)	Finite-difference LLG on a Cartesian grid; moving-window extension	NVIDIA-GPU micromagnetics	A moving computational window is not a moving unstructured body, and the cited solver is not an airbox-free FEM–BEM eddy solver
Yao et al. (Yao et al., 2022)	LLG coupled to a uniform Cartesian Yee grid	Portable manycore/GPU AMReX implementation of time-domain Maxwell–LLG	Scalable electromagnetic feedback, but not fixed-connectivity unstructured motion or conductor-only FEM–BEM
Feischl and Tran (Feischl and Tran, 2017)	Finite-element LLG in a bounded magnetic domain	Analytical and numerical FEM–BEM treatment of the unbounded eddy-current problem	Coupled variational formulation; no portable device-resident moving-mesh implementation
DynaMag	Unstructured fixed-connectivity LLG with rigid motion, verified morphing components, material frames, and time laws	SzCompute CPU/CUDA/HIP/Metal foundation; B-AIM open/periodic kernels; airbox-free FEM–BEM eddy coupling	Combination developed in this dissertation; moving-mesh electromagnetic coupling remains incomplete

state.

Chapter 3 presents SzCompute and DynaMag’s dynamic-state design. The chapter defines the prepare–apply–accept lifecycle, separates topology from evolving state, and measures the cost of mesh and material updates.

Chapter 4 derives the hybrid superposition integral–Poisson magnetostatic method. It evaluates memory use, accuracy, and performance for dynamic geometry.

Chapter 5 derives the periodic $1/r$ Green function and its analytical $\mathbf{r}/r^3$ gradient, then treats partial-period zero modes and validates the corrected B-AIM operators.

Chapter 6 develops airbox-free FEM–BEM eddy-current formulations. It presents the free-space $\mathbf{A}-k$ path, the intermediate $\mathbf{H}-k$ form, and the final periodic $\mathbf{H}-\kappa$ coordinate before combining the micromagnetic and eddy-current providers. The chapter ends with a fixed-geometry coupling test.

Chapter 7 describes B-AIM acceleration, local replacement, memory ownership, and preconditioning for the FEM–BEM solver, then reports solver ablations and a TEAM11 comparison with LS-DYNA.

Chapter 8 summarizes the reported solver capabilities, their present limitations, and the remaining work for moving periodic LLG–Maxwell calculations.

2 Micromagnetic and Electromagnetic Foundations

Micromagnetics describes a magnetic body through a continuous magnetization field rather than through its individual atomic spins. This continuum model retains the interactions that control magnetic structure and dynamics, including exchange, crystalline anisotropy, applied excitation, and the magnetostatic field created by the body itself (Brown, 1963; Aharoni, 2001). Conducting materials add an eddy current field through Maxwell's equations.

These fields enter the same equation of motion, but they lead to different numerical problems. Anisotropy is local, while exchange couples neighboring finite elements. In contrast, the demagnetizing field connects every magnetic source to an unbounded exterior, and the eddy current field also carries a time history. A dynamic mesh changes the geometry seen by all of these terms. Even when the connectivity is fixed, a rigid motion or a smooth morph changes coordinates, material frames, and long-range interactions. This chapter develops the physical equations first and then identifies the numerical state that must follow the moving material.

2.1 Magnetization dynamics

The magnetization at a point has both a magnitude and a direction. Let $\Omega_m(t) \subset \mathbb{R}^3$ denote the magnetic region and let $M_s(\mathbf{x}, t) > 0$ be its saturation magnetization. We write

$$\mathbf{M}(\mathbf{x}, t) = M_s(\mathbf{x}, t)\mathbf{m}(\mathbf{x}, t), \quad |\mathbf{m}| = 1. \quad (2.1)$$

The unit vector $\mathbf{m}$ is the dynamic unknown. Its motion is governed by the Landau–Lifshitz–Gilbert (LLG) equation. For a fixed body, the explicit Gilbert form is

$$\frac{\partial \mathbf{m}}{\partial t} = -\frac{\gamma}{1 + \alpha^2} [\mathbf{m} \times \mathbf{H}_{\text{eff}} + \alpha \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}})]. \quad (2.2)$$

Here $\gamma > 0$ is the magnitude of the gyromagnetic ratio, $\alpha \geq 0$ is the Gilbert damping coefficient, and $\mathbf{H}_{\text{eff}}$ is the total field that drives the LLG torque (Gilbert, 2004). The first torque causes precession around the field, whereas the second damps this motion and turns the magnetization toward the field. Because both torques are perpendicular to $\mathbf{m}$, the continuous equation preserves $|\mathbf{m}| = 1$.

The analytical literature establishes weak-solution frameworks for the Landau–Lifshitz dynamics (Visintin, 1985; Alouges and Soyeur, 1992), while finite-element analyses cover implicit, explicit, and projection-free discretizations (Alouges, 2008; Bartels and Prohl, 2006; Bartels et al., 2007; Alouges et al., 2014; Bartels, 2015). Geometric midpoint integration offers a complementary way to respect the unit-sphere constraint (d’Aquino et al., 2005), and preconditioned implicit integration addresses the stiffness of resolved micromagnetic transients (Suess et al., 2002). DynaMag uses these results as numerical context; the specific accepted-state lifecycle developed here is an implementation contract rather than a new existence result.

Equation (2.2) follows a material point. This distinction becomes important when the body moves. Let $\mathbf{x} = \chi(\mathbf{X}, t)$ map a reference point $\mathbf{X}$ to its current position. Defining $\widehat{\mathbf{m}}(\mathbf{X}, t) = \mathbf{m}(\chi(\mathbf{X}, t), t)$ gives

$$\frac{D\mathbf{m}}{Dt} = \left. \frac{\partial \widehat{\mathbf{m}}}{\partial t} \right|_{\mathbf{X}} = \frac{\partial \mathbf{m}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{m}, \quad \mathbf{v} = \frac{\partial \chi}{\partial t}. \quad (2.3)$$

The first expression follows the material, while the final expression is written at a fixed spatial point. DynaMag stores the magnetic state on material vertices and advances it at fixed $\mathbf{X}$. Rigid transport of a texture is then separated from the local rotation of its spins.

The energy-derived field collects the terms represented by the micromagnetic energy. In the

Gaussian convention used by the micromagnetic field modules,

$$\mathbf{H}_E = -\frac{1}{M_s} \frac{\delta E}{\delta \mathbf{m}} = \mathbf{H}_{\text{ex}} + \mathbf{H}_{\text{an}} + \mathbf{H}_{\text{ap}} + \mathbf{H}_{\text{d}} + \mathbf{H}_{\text{other}}. \quad (2.4)$$

The eddy-current reaction field has a different origin. It is obtained from the evolving Maxwell state, not from the variation of the instantaneous micromagnetic energy. When eddy feedback is active, the total LLG field is therefore

$$\mathbf{H}_{\text{eff}} = \mathbf{H}_E + \mathbf{H}_{\text{eddy}}. \quad (2.5)$$

An SI derivation places an additional factor of μ_0 in the denominator. The two conventions are equivalent after unit conversion, but their factors cannot be mixed within one residual. DynaMag uses nanometers, nanoseconds, emu/cm³, and Oe in the micromagnetic modules. The eddy current solver uses SI units, so Section 6.6 defines the conversions at the coupling interface. These conversions are applied while the residual is formed because the eddy reaction field acts directly on the LLG equation.

The role of damping is seen from the energy balance. For static material parameters and fixed geometry, and without electromagnetic feedback, Equation (2.2) gives

$$\frac{dE}{dt} = -\frac{\alpha\gamma}{1+\alpha^2} \int_{\Omega_m} M_s |\mathbf{m} \times \mathbf{H}_E|^2 dV \leq 0, \quad (2.6)$$

up to the selected magnetic-unit normalization. The energy decreases because the damping torque removes the component of motion transverse to the field. Applied fields, changing material laws, mechanical work, and electromagnetic loss introduce separate power terms, so the total magnetic energy may rise in a driven or deforming system. The unforced, fixed-geometry limit remains a basic check because damping alone cannot add energy.

At finite temperature, stochastic and longitudinal-relaxation models extend this zero-temperature LLG picture. Fokker–Planck and Landau–Lifshitz–Bloch descriptions, stochastic-

calculus interpretations, and near-Curie-temperature dynamics are developed in (Garanin, 1997; Berkov and Gorn, 2002; Chubykalo-Fesenko et al., 2006; Evans et al., 2012). Those models are not used in the validations below, but they explain why temperature-dependent constitutive state belongs in the runtime design.

2.2 Local, bulk, and interfacial effective fields

The local and short-range fields describe how the magnetization responds to its immediate material environment. They determine much of the stiffness of the LLG equation, but their finite-element actions involve only an element or its neighbors. Their dependence on motion is therefore visible through local geometry and constitutive data. This is simpler than the global update required by magnetostatics.

2.2.1 Bulk exchange

Neighboring atomic moments favor alignment. In the continuum model, this tendency appears as an energy penalty on spatial variation of the magnetization direction,

$$E_{\text{ex}} = \int_{\Omega_m} A_{\text{ex}} |\nabla \mathbf{m}|^2 dV, \quad (2.7)$$

where A_{ex} is the exchange stiffness. For constant A_{ex} in a homogeneous region, its strong-form field is

$$\mathbf{H}_{\text{ex}} = \frac{2A_{\text{ex}}}{M_s} \nabla^2 \mathbf{m}. \quad (2.8)$$

When A_{ex} varies in space, the conservative form is $\mathbf{H}_{\text{ex}} = (2/M_s) \nabla \cdot (A_{\text{ex}} \nabla \mathbf{m})$. Let $\mathbf{v} \in H^1(\Omega_m)^3$ be an arbitrary vector test function. The finite-element weak form then requires only first derivatives,

$$\int_{\Omega_m} M_s \mathbf{v} \cdot \mathbf{H}_{\text{ex}} dV = -2 \int_{\Omega_m} A_{\text{ex}} \nabla \mathbf{v} : \nabla \mathbf{m} dV + 2 \int_{\partial \Omega_m} A_{\text{ex}} \mathbf{v} \cdot \partial_n \mathbf{m} dS. \quad (2.9)$$

Without a separate surface energy, the boundary term gives the natural free-surface condition $\partial_n \mathbf{m} = 0$. The remaining volume integral couples only basis functions that share an element, so the discrete exchange operator is sparse.

The exchange length sets a useful scale for resolving this competition. It compares bulk exchange with magnetostatic energy and, in the Gaussian convention, is commonly written as

$$\ell_{\text{ex}} = \sqrt{\frac{A_{\text{ex}}}{2\pi M_s^2}}. \quad (2.10)$$

A mesh must resolve structures set by this and related material lengths. That condition applies to the current configuration, not only to the reference mesh. A morph can stretch an initially adequate element until it no longer resolves the same domain wall or vortex.

2.2.2 Surface and intergranular exchange

Bulk exchange acts within a continuous magnetic region. When two grains or layers meet, their coupling is instead described on the shared interface. For the traces $\mathbf{m}^+$ and $\mathbf{m}^-$ on an interface Γ_{+-} , a useful surface energy is

$$E_{\Gamma, \text{ex}} = \int_{\Gamma_{+-}} J_{\Gamma} (1 - \mathbf{m}^+ \cdot \mathbf{m}^-) dS, \quad (2.11)$$

where J_{Γ} has units of energy per area. Positive J_{Γ} favors parallel alignment and negative J_{Γ} favors antiparallel alignment. Unlike bulk exchange, this energy acts through the two traces rather than through a volume gradient. Variation with respect to $\mathbf{m}^+$ gives a term proportional to $-J_{\Gamma} \mathbf{m}^-$, while the minus trace receives the corresponding coupling. Combining this variation with Equation (2.9) gives the torque-relevant interface condition

$$(\mathbf{I} - \mathbf{m}^+ \mathbf{m}^{+T}) (2A_{\text{ex}}^+ \partial_{n^+} \mathbf{m}^+ - J_{\Gamma} \mathbf{m}^-) = \mathbf{0}, \quad (2.12)$$

and its counterpart on the minus side. The tangent projection appears because variations of a unit vector lie in the tangent plane.

In granular and multilayer structures, the interface coefficient can change reversal and loss even when the bulk material is unchanged (Suess, 2007). A dynamic material description must therefore include interfacial data as well as volume parameters. Motion changes the interface area, its normal, and the two magnetization traces. The coefficient J_{Γ} may also depend on time or on the local interface state. DynaMag stores this information with the material interface so that it can follow the evolving geometry.

2.2.3 Uniaxial anisotropy

A uniaxial crystal has one preferred material axis $\mathbf{a}$, with $|\mathbf{a}| = 1$. Let $c = \mathbf{m} \cdot \mathbf{a} = \cos \theta$, where θ is the angle between the magnetization and the axis. The first- and second-order anisotropy energy densities are $K_1 \sin^2 \theta$ and $K_2 \sin^4 \theta$ (Brown, 1963; Aharoni, 2001). Therefore,

$$E_u = \int_{\Omega_m} e_u dV, \quad e_u = K_1 (1 - c^2) + K_2 (1 - c^2)^2. \quad (2.13)$$

Differentiation with respect to $\mathbf{m}$ gives

$$\frac{\partial e_u}{\partial \mathbf{m}} = -[2K_1 + 4K_2(1 - c^2)] c \mathbf{a}, \quad \mathbf{H}_u = -\frac{1}{M_s} \frac{\partial e_u}{\partial \mathbf{m}} = \frac{[2K_1 + 4K_2(1 - c^2)] c}{M_s} \mathbf{a}. \quad (2.14)$$

This derivative confirms the sign and both numerical factors. For $K_1 > 0$ and $K_2 = 0$, the energy is lowest at $c = \pm 1$, as required for an easy axis. DynaMag removes the part of $\mathbf{H}_u$ parallel to $\mathbf{m}$ before accumulation. This projection does not change the LLG torque because $\mathbf{m} \times \mathbf{m} = \mathbf{0}$.

2.2.4 Cubic anisotropy

A cubic crystal instead has three orthonormal material axes $\{\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3\}$. Define the direction cosines $c_i = \mathbf{m} \cdot \mathbf{a}_i$. The first two cubic invariants give

$$E_c = \int_{\Omega_m} e_c dV, \quad e_c = K_{c1}(c_1^2 c_2^2 + c_2^2 c_3^2 + c_3^2 c_1^2) + K_{c2} c_1^2 c_2^2 c_3^2. \quad (2.15)$$

For each i , let j and k denote the other two crystal directions. The chain rule gives

$$\frac{\partial e_c}{\partial c_i} = 2c_i [K_{c1}(c_j^2 + c_k^2) + K_{c2}c_j^2 c_k^2]. \quad (2.16)$$

Since $\partial c_i / \partial \mathbf{m} = \mathbf{a}_i$, the complete cubic field is

$$\mathbf{H}_c = -\frac{2}{M_s} \sum_{i=1}^3 c_i [K_{c1}(c_j^2 + c_k^2) + K_{c2}c_j^2 c_k^2] \mathbf{a}_i. \quad (2.17)$$

This result confirms the negative sign and the factor of two used by the solver. With $K_{c2} = 0$, a positive K_{c1} favors the cubic axes, whereas a negative K_{c1} favors the body diagonals. As in the uniaxial case, DynaMag removes the torque-free component parallel to $\mathbf{m}$.

The anisotropy axes belong to the material rather than to the laboratory. They rotate with a rigid body and follow the local material rotation under a general deformation. If the mesh moves while these axes remain fixed in the laboratory frame, the calculation represents a different constitutive model. The configuration names the two constants K_1 and K_2 for both anisotropy types; K_{c1} and K_{c2} are used here only to distinguish the cubic constants from the uniaxial ones.

2.2.5 Applied field

An external magnetic field enters through the Zeeman energy

$$E_{ap} = - \int_{\Omega_m} \mathbf{M} \cdot \mathbf{H}_{ap} dV, \quad (2.18)$$

and appears directly in $\mathbf{H}_{\text{eff}}$. The field may vary with position, time, or an experimental control. Its reference frame must also be specified. A laboratory-fixed field does not rotate with the sample, whereas a body-fixed excitation does. A rigid rotation provides a direct test of this distinction.

2.3 Demagnetizing field: the global micromagnetic bottleneck

The magnetization creates a field that acts back on the magnetic body. This demagnetizing field is generated by both volume and surface magnetic charge. In the absence of free current, it satisfies the Gaussian-unit equations

$$\nabla \times \mathbf{H}_d = \mathbf{0}, \quad \nabla \cdot (\mathbf{H}_d + 4\pi \mathbf{M}) = 0 \quad \text{in } \mathbb{R}^3. \quad (2.19)$$

The curl-free condition permits a scalar potential $\mathbf{H}_d = -\nabla \phi$. Substitution into the divergence equation gives

$$\nabla^2 \phi = 4\pi \nabla \cdot \mathbf{M} \quad \text{in the magnetic body}, \quad (2.20)$$

together with continuity and jump conditions at material surfaces and decay at infinity. With volume charge $\rho_m = -\nabla \cdot \mathbf{M}$ and surface charge $\sigma_m = \mathbf{M} \cdot \mathbf{n}$, the isolated-body potential is

$$\phi(\mathbf{x}) = \int_{\Omega_m} \frac{\rho_m(\mathbf{y})}{|\mathbf{x} - \mathbf{y}|} d\mathbf{y} + \int_{\partial\Omega_m} \frac{\sigma_m(\mathbf{y})}{|\mathbf{x} - \mathbf{y}|} dS_{\mathbf{y}}. \quad (2.21)$$

Integrating the charge representation by parts gives the equivalent magnetization form

$$\phi(\mathbf{x}) = \int_{\Omega_m} \mathbf{M}(\mathbf{y}) \cdot \nabla_{\mathbf{y}} \frac{1}{|\mathbf{x} - \mathbf{y}|} d\mathbf{y}. \quad (2.22)$$

The corresponding field energy can be written in two equivalent ways:

$$E_d = -\frac{1}{2} \int_{\Omega_m} \mathbf{M} \cdot \mathbf{H}_d dV = \frac{1}{8\pi} \int_{\mathbb{R}^3} |\mathbf{H}_d|^2 dV. \quad (2.23)$$

The first expression measures the interaction inside the magnet. The second shows that the field stores energy throughout space, so the exterior cannot be discarded even though the magnetization is confined to Ω_m .

The Poisson and integral forms describe the same field, but they distribute the numerical cost differently. Although the Poisson matrix is sparse, a conventional finite-element treatment must surround the magnet with a large air domain and impose an artificial condition on its outer boundary. The integral form satisfies the open boundary directly, but every target then depends on every source, and direct evaluation or dense storage scales as $O(N^2)$. Fast multipole and FFT-based methods reduce this cost, although they introduce geometry-dependent grids, source projections, target interpolation, and near-field corrections.

This global dependence makes demagnetization the most difficult micromagnetic field to update during motion. A rigid translation preserves internal distances, but it changes the relation to observation points, periodic cells, and a laboratory-fixed grid. A rotation also changes vector sources, surface normals, and target frames. A morph changes the pair distances themselves. In every case, the boundary charge and the long-range operator must describe the same trial geometry. Reusing a stale normal, near list, interpolation plan, or Poisson geometry would compute the field of a different body.

Analytical cylinder and ideal-particle results provide useful limiting benchmarks for demagnetization and nucleation (Aharoni and Shtrikman, 1958; Frei et al., 1957). For spatially varying magnetization, generalized demagnetizing tensors make the cell interaction explicit (Newell et al., 1993). Early dynamic and finite-element micromagnetic calculations already identified accurate demagnetizing-field evaluation as a central repeated cost (Mansuripur and Giles, 1988; Bagnères-Viallix and Baras, 1991).

Chapter 4 resolves this difficulty by evaluating the integral representation only on selected surfaces. A sparse Dirichlet Poisson solve then recovers the interior field. Chapter 5 extends the same long-range machinery to an infinite lattice. Both methods retain the open-boundary field without introducing an airbox.

2.4 Eddy-current reaction field

A time-dependent magnetic flux induces an electric field in a conductor. The resulting current produces a magnetic reaction field that enters the LLG effective field. This feedback is described by the magnetoquasistatic form of Maxwell's equations, which becomes the following SI system after displacement current is neglected:

$$\nabla \times \mathbf{H} = \mathbf{J}_s + \mathbf{J}_e, \quad \nabla \cdot \mathbf{B} = 0, \quad (2.24)$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad \mathbf{B} = \mu_0(\mathbf{H} + \mathbf{M}), \quad (2.25)$$

$$\mathbf{J}_e = \sigma \mathbf{E} \quad \text{for a stationary isotropic conductor.} \quad (2.26)$$

Here $\mathbf{J}_s$ is an imposed source current, $\mathbf{J}_e$ is the induced current, and σ is the electrical conductivity. The part of $\mathbf{H}$ created by $\mathbf{J}_e$ is the eddy reaction field. It is converted to the micromagnetic field units before it enters the LLG torque through Equation (2.5).

The induced current introduces a second spatial scale. For a harmonic excitation of angular frequency ω , the skin depth

$$\delta = \sqrt{\frac{2}{\omega \mu \sigma}} \quad (2.27)$$

measures how far the field penetrates the conductor. A mesh that resolves the exchange length may still be too coarse at high frequency. The eddy response also carries memory because its value at t_n depends on flux changes over earlier accepted states, not only on the current magnetization $\mathbf{m}_n$. Rejected time-step trials must therefore leave the electromagnetic history unchanged.

Motion can also change the constitutive relation. For a conductor moving with laboratory velocity $\mathbf{v}$, the nonrelativistic form of Ohm's law is

$$\mathbf{J}_e = \sigma(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \quad (2.28)$$

The term $\mathbf{v} \times \mathbf{B}$ is a physical electromotive force, not a geometric correction. Its inclusion depends on

the frame and on the regime being modeled. In either case, the conductor volume and its boundary traces must be transported together. The FEM–BEM formulation itself is valid on the resulting moving domain.

The current is confined to the conductor, but its magnetic field extends through the unbounded exterior. Meshing the surrounding air would reproduce the moving-airbox problem already encountered in magnetostatics. To avoid that problem, the FEM–BEM formulation in Chapter 6 uses conforming Nédélec elements in the conductor and a boundary representation for the exterior (Nedelec, 1980; Hiptmair, 2002; Monk, 2003; Johnson and Nédélec, 1980; Feischl and Tran, 2017).

2.5 Finite-element and time discretization

Finite-element spaces are chosen from the continuity required by the field. Magnetization and scalar magnetostatic potential are represented by nodal H^1 elements. Electric and magnetic field coordinates instead require a continuous tangential trace, so they use Nédélec $H(\text{curl})$ elements. The boundary unknowns occupy the matching trace spaces on Γ and connect the conductor to the unbounded exterior. This compatibility is essential. An inappropriate basis can introduce spurious modes or break the interface identities used by FEM–BEM.

Computational micromagnetics texts and surveys place these choices in a broader numerical setting (Prohl, 2001; Kruzík and Prohl, 2006; Cimrák, 2008). Automated finite-element frameworks demonstrate how weak forms and compatible spaces can be translated into reusable discrete operators (Logg et al., 2012); the design in this dissertation instead keeps the hot actions explicitly device resident. The same finite-element LLG structure has also been extended to spin-polarized transport in magnetic multilayers (Abert et al., 2014).

For nodal basis functions N_i , write

$$\mathbf{m}_h(\mathbf{x}, t) = \sum_i N_i(\mathbf{x}) \mathbf{m}_i(t). \quad (2.29)$$

Substitution into the weak exchange form gives the sparse stiffness matrix

$$K_{ij}^{\text{ex}} = \int_{\Omega_m} 2A_{\text{ex}} \nabla N_i \cdot \nabla N_j dV, \quad (2.30)$$

because two basis functions interact only when they share an element. Anisotropy and applied fields then use local nodal or quadrature values, whereas the demagnetizing field adds a global open or periodic interaction. The eddy system further adds Nédélec mass, curl, and boundary actions. A single LLG residual therefore combines local algebra, sparse finite-element operators, and a global field with accepted time history.

These different time scales favor an implicit, adaptive integrator. Exchange and eddy diffusion can impose severe explicit stability limits, and the appropriate step changes during switching. A q th-order backward differentiation formula has the schematic residual

$$\mathbf{R}_n(\mathbf{m}_n) = \sum_{j=0}^q a_j \mathbf{m}_{n-j} - \Delta t_n \mathbf{f}(\mathbf{m}_n, \mathbf{H}_{\text{eff}}[\mathbf{m}_n, \boldsymbol{\chi}_n, \mathcal{M}_n, \mathcal{H}_{n-1}]) = \mathbf{0}, \quad (2.31)$$

where $\mathcal{M}_n$ is the material state and $\mathcal{H}_{n-1}$ is the accepted electromagnetic history. Every residual evaluation at t_n must use the same trial geometry, material data, and field operators. These values become part of the physical history only after the step is accepted. This separation between trial and accepted state leads to the prepare–apply–commit sequence developed in Chapter 3.

Implicit magnetic residuals are commonly paired with Newton–Krylov ideas (Knoll and Keyes, 2004); their nonsymmetric linear subproblems motivate the GMRES family (Saad and Schultz, 1986). Scalable parallel micromagnetic solvers and block inverse preconditioners illustrate how field structure determines practical convergence (Scholz et al., 2003; Fu et al., 2019). DynaMag follows the same principle by separating physical operator application from a fixed-cost preconditioner action.

2.6 Rigid motion, morphing, and operator validity

The preceding field equations are evaluated on the current configuration $\Omega(t)$. For a moving mesh, this configuration is related to a fixed reference domain Ω_0 through the deformation map

$$\mathbf{x} = \boldsymbol{\chi}(\mathbf{X}, t), \quad \mathbf{F} = \frac{\partial \boldsymbol{\chi}}{\partial \mathbf{X}}, \quad J = \det \mathbf{F} > 0. \quad (2.32)$$

The map $\boldsymbol{\chi}$ gives the current position, $\mathbf{F}$ gives the local deformation, and J gives the change in volume. A rigid body has $\boldsymbol{\chi}(\mathbf{X}, t) = \mathbf{Q}(t)\mathbf{X} + \mathbf{c}(t)$ with $\mathbf{Q}^T\mathbf{Q} = \mathbf{I}$, so $\mathbf{F} = \mathbf{Q}$ and $J = 1$. A morph has a spatially varying $\mathbf{F}$ and can change length, angle, area, and volume even when element connectivity remains fixed. The polar decomposition

$$\mathbf{F} = \mathbf{R}\mathbf{U} \quad (2.33)$$

separates the local material rotation $\mathbf{R}$ from the stretch $\mathbf{U}$. The rotation transports crystal axes and other material directions. The full $\mathbf{F}$ and J determine the element gradients, areas, and volumes.

For a pulled-back field $\widehat{\mathbf{m}}(\mathbf{X}) = \mathbf{m}(\boldsymbol{\chi}(\mathbf{X}))$, the physical gradient is

$$\nabla_{\mathbf{x}}\mathbf{m} = \nabla_{\mathbf{X}}\widehat{\mathbf{m}}\mathbf{F}^{-1}. \quad (2.34)$$

Under a rigid rotation, $\mathbf{F}$ is orthogonal. Exchange energy is then unchanged if the magnetization and material frame rotate together. A general morph, however, changes the metric in Equation (2.7) and therefore changes the exchange operator. Fixed connectivity does not imply fixed physics because coordinates, local frames, quadrature factors, and long-range interactions may all change while the topology remains intact.

Table 2.1 summarizes these dependencies. Low-cost rigid-body and material updates do not make all fields independent of geometry.

The table shows why the magnetostatic field dominates the geometric update. It is both long-range and open-boundary, while an implicit LLG step may evaluate it many times. A local

Table 2.1. Geometry updates required by the principal LLG fields.

Field	Discrete character	Rigid-body update	Morph update
Applied field	Pointwise	Evaluate in the laboratory or body frame	Evaluate at the current points and time
Anisotropy	Local	Rotate the crystal axes with the body	Transport the axes with local $\mathbf{R}$; update material properties
Bulk exchange	Sparse FEM	Translation leaves the metric unchanged	Recompute gradients, Jacobians, and metric factors
Surface exchange	Sparse interface	Rotate traces and normals	Update interface areas, normals, traces, and properties
Magnetostatic field	Global, open or periodic	Rotate sources; refresh coordinate data	Rebuild grid, near-field, normal, and Poisson data
Eddy-current field	Global with history	Refresh volume and boundary operators	Refresh Nédélec geometry, boundary coupling, preconditioner, and history

morph can therefore change the field at distant targets, and the cost of rebuilding acceleration data affects the entire trajectory. The eddy current problem adds a coupled boundary system and an accepted field history. The magnetostatic field remains the more universal difficulty because every magnetic body requires it, including an insulator.

In this dissertation, “fully dynamic” refers to the state used by the field operators. Moving displayed vertices is not sufficient. At each trial time, the coordinates, material frames, constitutive values, operators, and field history must describe one physical configuration. The same requirement applies to rigid motion and fixed-connectivity morphing. Remeshing and topology change are separate extensions.

2.7 Fast long-range evaluation and verification hierarchy

The open-boundary fields above are expressed through the scalar kernel $1/r$ or its gradient $\mathbf{r}/r^3$. In a periodic system, the response includes an infinite lattice of images. For periodic axes $\mathcal{P}$

and lattice translations $\mathbf{L}_{\mathbf{n}}$, the formal scalar sum is

$$G_{\mathcal{P}}(\mathbf{r}) = \sum_{\mathbf{n} \in \mathbb{Z}^{|\mathcal{P}|}} \frac{1}{|\mathbf{r} - \mathbf{L}_{\mathbf{n}}|}, \quad (2.35)$$

The meaning of this sum depends on the source mean and on the scalar gauge. Its gradient has different convergence and zero-mode behavior. In particular, it cannot be obtained safely by differentiating an arbitrarily normalized discrete scalar potential. Chapter 5 derives the periodic $\mathbf{r}/r^3$ field directly from the periodic Green function.

FFT-based acceleration methods such as B-AIM make the all-to-all interaction practical on an unstructured finite-element mesh. In B-AIM, sources are projected to a regular grid, the convolution is performed in Fourier space, and the result is interpolated back to the targets. The inaccurate local part of the grid interaction is then replaced by the physical near field (Greengard and Lee, 2004; Fessler and Sutton, 2003). Accuracy depends on the grid, projection order, kernel, panel treatment, and matched near correction. Motion adds an important distinction: topology data may be retained, but projections, near-field geometry, and target maps belong to the current configuration.

The next chapters follow this sequence. Chapter 3 develops the DynaMag state and execution model. Chapter 4 applies that model to open-boundary magnetostatics, and Chapter 5 develops the periodic Green and B-AIM foundation. Chapters 6 and 7 then use these tools for the airbox-free eddy current problem and its LLG coupling.

3 DynaMag: Accelerator-Portable Micromagnetics with Dynamic State

DynaMag is designed for magnetic bodies whose geometry and material response change during a simulation. These changes do not always require a new mesh: a rotating or smoothly deforming body can retain its element connectivity while its coordinates, element maps, material directions, and operator coefficients evolve. Likewise, a material property can change with time without changing which material belongs to an element. DynaMag therefore separates fixed topology from the evolving physical state.

Figure 3.1 shows several applications that require a dynamic mesh. These include varying gaps in granular magnetic materials such as ferrites, the magnetic response beneath a moving write head, deformation of flexible magnetic structures, and magnetically actuated microbots moving through fluid.

3.1 DynaMag architecture

Dynamic geometry is useful only if updating it does not overwhelm the field calculation. DynaMag combines unstructured finite elements, sparse linear algebra, FFT-accelerated fields, and implicit time integration. These operations share large state vectors and must form one device-resident calculation. Repeated host staging would erase much of the FFT benefit, while large temporary arrays would defeat the memory advantage of matrix-free kernels. At the same time,




Figure 3.1. Applications of dynamic meshes: varying gaps in granular magnetic materials, motion of a write head over recording media, deformation of a magnetic sheet, and magnetic actuation of a microbot.

fixed mesh data should survive each time step when their defining topology has not changed.

DynaMag follows three numerical requirements:

1. one numerical expression must cover host and accelerator execution, while large state vectors remain on the selected provider;
2. topology-dependent structure must be separated from quantities that depend on coordinates, materials, or time;
3. repeated operators must reuse prepared data and persistent workspace. Portability is verified separately at the primitive, operator, and coupled-solver levels.

Figure 3.2 shows how these requirements are divided between the software layers. SzCompute supplies the backend, memory, and numerical operations. DynaMag builds finite-element spaces, solvers, and magnetic-field models on this foundation. The field orchestrator then combines

the active terms into the effective field used by the time integrator. This separation keeps the physical model independent of the selected hardware provider.




Figure 3.2. DynaMag numerical architecture. The field orchestrator combines the applied, anisotropy, exchange, magnetostatic, eddy-current, and torque terms. These models share the finite-element spaces, solvers, preconditioners, and memory services above the CPU, CUDA, and ROCm backends. Metal support uses the same abstraction described in Section 3.2.

The separation allows coordinates, constitutive values, magnetization, and electromagnetic history to evolve without rebuilding unrelated data. It also makes the cost of a dynamic calculation easier to interpret. State updates, operator preparation, repeated solves, and output can be measured separately.

3.2 The SzCompute portable numerical substrate

The physical algorithms require the same operations on CPUs and several GPU families. SzCompute supplies this common numerical substrate in C++20. It provides explicit providers, owning containers, nonowning views, parallel algorithms, sparse algebra, and FFTs. Each finite-element vector has a known provider and defined access rules. Kernels receive their arrays and

metadata through views, so the element, interpolation, and near-field formulas do not embed a hidden memory policy.

Table 3.1. Principal SzCompute services used by DynaMag.

Provider	Principal services	Dissertation role
Multicore host	Threaded kernels, sparse/dense algebra, FFTs	Reference execution, CPU runs, and cross-provider verification
NVIDIA CUDA	Device containers, cuBLAS/cuSPARSE, cuFFT	Primary device used for the largest validation and performance cases
AMD HIP/ROCm	Device containers, rocBLAS/rocSPARSE, rocFFT	Independent GPU execution of the same primitive families
Apple Metal	Metal compute, MPS services, unified-memory integration	Portable compute foundation; full-physics coverage is stated separately from primitive support

Provider selection belongs to an operation rather than to the process as a whole. The input, output, or workspace determines where that operation runs. During an LLG solve, the magnetization, trial states, residuals, fields, and Krylov vectors remain on the selected provider. B-AIM likewise keeps its projected sources, transform grids, interpolation data, and near-field work resident. As a result, data leave the provider only at declared boundaries such as input, diagnostics, and requested output.

Most of the computational work lies in sparse local actions and grid convolutions for long-range fields. Both use the common provider interface, but their prepared objects retain backend-specific information. Sparse graphs and vendor analysis objects are reused, while FFT plans retain their backend handles and persistent workspaces own the mutable transform storage. The interface therefore unifies the mathematical operation without assuming that launch costs, transform layouts, sparse algorithms, and memory architectures are identical.

Table 3.2 confirms that the same large operation families run on the Host, CUDA, and ROCm providers. The machines differ, so the numbers should not be read as controlled speedups.

Metal is included in the support matrix because its compute and unified-memory services are part of the foundation.

Table 3.2. Selected median measurements from the tracked 22 June 2026 SzCompute Host/CUDA/ROCm snapshot. Values are milliseconds; hardware differs between columns.

Operation and size	Host	CUDA	ROCm
Double-complex FFT round trip, 128^3	3.865	0.347	0.250
Single-complex FFT round trip, 128^3	1.254	0.036	0.143
Cached sparse action, double, 10^6 rows/ 15×10^6 nonzeros	3.230	0.172	0.695
Cached sparse action, single, 10^7 rows/ 150×10^6 nonzeros	48.838	1.558	13.590
Table transpose, 5×10^5 rows/ 8×10^6 entries	26.365	2.231	6.432

3.2.1 Matrix-free finite-element action

Sparse matrices remain important for Poisson problems and preconditioners, but they are not always the best representation of a local micromagnetic field. Exchange provides a direct example. Its coefficients depend on the current element geometry and on the material ratio A_{ex}/M_s . For P1 tetrahedra, the lumped field at global degree of freedom i is

$$\mathbf{H}_{\text{ex},i} = \frac{1}{m_i^L} \sum_{e \ni i} \sum_{j \in e} \left[-\frac{2A_{\text{ex},e}(t)}{M_{s,e}(t)} \int_{\Omega_e(t)} \nabla N_i \cdot \nabla N_j dV \right] \mathbf{m}_j, \quad m_i^L = \sum_{e \ni i} \int_{\Omega_e(t)} N_i dV. \quad (3.1)$$

The bracketed coefficient is never stored in a global CSR array. Instead, one work item loads the four current vertices and their magnetization values. After evaluating the tetrahedral volume and four physical basis gradients, it applies the local 4×4 operator to all three components of $\mathbf{m}$. The resulting contributions are scattered to the four field nodes and to a lumped-volume vector, after which a final nodal kernel divides by m_i^L . This lumped-volume vector is the 8.189 MiB workspace reported below, with one single-precision scalar for each of the 2,146,689 nodes. Its allocation is retained, while its values are cleared and recomputed for the current geometry.

Table 3.3 follows the data through this action. Only topology and the reusable nodal workspace are retained. No global coefficient depending on geometry or material state is stored.

Table 3.3. Dataflow of the matrix-free P1 tetrahedral exchange-field action.

Stage	Retained input	Work performed for the current state
Topology	Element-to-DoF incidence, element and material identifiers	No coefficient matrix or CSR values are constructed
Geometry	Reference simplex and element connectivity	Load four current vertices; evaluate volume and the four constant physical gradients
Constitutive law	Material ownership and law parameters	Evaluate $A_{\text{ex}}(t)$ and $M_s(t)$ for the current trial time
Local action	Four values of each component of $\mathbf{m}$	Form the 4×4 local stiffness action in registers and accumulate three output components
Workspace	One nodal lumped-volume vector	Clear and recompute values without allocating or freeing device memory
Global accumulation	Stable global DoF indices	Scatter element field and lumped-volume contributions, then normalize nodally
Invalidation	Connectivity only	Motion or material change requires no matrix rebuild; topology change requires a new incidence structure

The benefit depends on how often the local coefficients change. To measure this trade, a 128^3 cube subdivision is used with 12,582,912 P1 tetrahedra, 2,146,689 nodal unknowns, and 31,802,497 scalar CSR entries. Table 3.4 reports single-precision CUDA measurements on an NVIDIA GeForce RTX 5090. Each result follows five warmups and twenty timed field calls. The matrix-free time includes geometry evaluation, the three-component local action, scatter, and lumped normalization. The three sparse times include only application of an assembled matrix. They exclude setup. Matrix construction received one warmup to remove the CUDA backend and JIT cold start. The timed reconstruction still allocates and clears the matrix, assembles every coefficient, and applies lumped scaling.

For a fixed operator, the assembled action is faster because it streams one stored coefficient per nonzero and does not recompute simplex geometry. The complete matrix-free field takes 2.69 times as long as the fused CSR apply, which is close to the expected factor of three for tetrahedra. Even with the geometric work included, the matrix-free kernel still processes 1.14×10^{10} element actions per second.

The comparison changes when the geometry or A_{ex}/M_s changes. If the graph is retained,

Table 3.4. P1 tetrahedral exchange-field comparison on 12,582,912 elements. Setup and repeated apply are separated so that a raw SpMV is not mistaken for a dynamic field cost.

Operation	Time (ms)	Persistent operator storage
Matrix-free three-component field	1.099	8.189 MiB nodal lumped-volume workspace; 0 MiB matrix values
Fused native three-component CSR apply	0.409	250.823 MiB CSR
Three vendor SpMV with retained caches	0.489	250.823 MiB CSR + 0.355 MiB cache scratch
Three raw native SzCompute SpMV	0.542	250.823 MiB CSR
Topology graph construction	639.991	required only after topology change
Sparse matrix construction	2.796	creates the CSR values
Vendor-cache first use, including first apply	8.860	creates cache state

matrix construction and one fused apply cost $2.796 + 0.409 = 3.205$ ms. The matrix-free path evaluates the new state in 1.099 ms, so the fused assembled path recovers its reconstruction cost only after about five applications at one state. The raw native path crosses after about six applications, while the vendor path prepares a cache and crosses after about 19. Rigid motion and morphing retain the graph, so graph construction is listed separately. If topology changes, that cost moves the fused crossover to about 932 applications. CSR is therefore useful when the operator remains fixed for many calls, whereas matrix-free evaluation avoids this amortization requirement when the operator changes with the physics.

Each parallel operation also receives a physical name. A near-field correction, element action, or preconditioner sweep can then be identified directly in a profile instead of being inferred from a compiler-generated symbol. The later performance analysis uses these names to separate setup from repeated work and to locate limits in the FFT, local correction, sparse solver, or time integrator.

3.3 DynaMag state and operator design

The matrix-free exchange example reflects a broader state model. Under prescribed motion with fixed connectivity, the discrete data can be divided into an invariant part

$$\mathcal{T} = \{\mathcal{C}, \mathcal{V}, \mathcal{M}, \mathcal{D}, \mathcal{G}\} \quad (3.2)$$

and an evolving part

$$\mathcal{S}(t) = \{\mathbf{x}(t), \mathbf{F}(t), \mathbf{R}(t), J(t), \mathbf{m}(t), \mathbf{u}_{\text{em}}(t), \mathbf{p}(t)\}. \quad (3.3)$$

Here $\mathcal{C}$ is the element connectivity; $\mathcal{V}$ and $\mathcal{M}$ identify volumes and materials; and $\mathcal{D}$ and $\mathcal{G}$ contain the degree-of-freedom numbering and reusable incidence graphs. These quantities depend on topology. In contrast, the second set contains the current coordinates $\mathbf{x}$, the deformation measures $\mathbf{F}$, $\mathbf{R}$, and J , the magnetization $\mathbf{m}$, the electromagnetic history $\mathbf{u}_{\text{em}}$, and the current material data $\mathbf{p}$. This division sets the lifetime of every prepared object used below.

This partition applies while the connectivity is fixed and each element remains valid. It preserves degree-of-freedom identity, which allows the magnetic and electromagnetic histories to remain attached to the same material points. It does not freeze the geometry; coordinates, integration factors, boundary normals, material frames, and long-range plans are refreshed from the evolving state. Table 3.5 summarizes this distinction.

A numerical operator is useful only while the data used to prepare it remain valid. DynaMag expresses this dependence through the lifecycle

$$\underbrace{\mathcal{P}(\mathcal{T}, \mathcal{S})}_{\text{prepare}} \rightarrow \underbrace{\mathbf{y} = \mathcal{A}(\mathcal{S})\mathbf{x}}_{\text{apply}} \rightarrow \underbrace{\mathcal{C}(\mathbf{y})}_{\text{commit or report}}. \quad (3.4)$$

Preparation constructs the required topology maps, transform plans, interpolation data, and work arrays. The apply stage evaluates the field or linear operator, often many times within one nonlinear solve. The commit stage advances physical history only after the time step is accepted, although

Table 3.5. State dependence under fixed-connectivity motion.

Quantity	After motion	Mathematical reason or action
Connectivity, region IDs, and DoF numbering	Retained	Incidence and ownership are unchanged
Sparse adjacency and graph structure	Retained	The nonzero pattern depends on connectivity
Coordinates, Jacobians, weights, and normals	Updated	These depend on the current geometric map
Material frames and constitutive values	Updated	Rotation, position, and time enter the constitutive law
Geometry-dependent Green and near-field data	Invalidated	Source and target positions have changed
Magnetization and accepted electromagnetic history	Preserved	These are physical state with stable DoF identity

the same stage may report an observable without changing that history. This lifecycle keeps setup outside repeated Krylov work and prevents rejected trials from altering the accepted solution.

Different prepared objects therefore have different lifetimes. A sparse graph survives fixed-connectivity motion even when its coefficient values change. A Green-function plan must refresh source and target geometry, but it can retain work arrays of the same shape. A material table remains fixed while its sampled value changes with the trial time. Time advancement alone does not trigger a rebuild; a changed dependency does.

The same separation defines the memory convention. For a nonlocal solve,

$$M_{\text{total}} = M_{\text{geometry}} + M_{\text{plan}} + M_{\text{workspace}} + M_{\text{solver}} + M_{\text{state}} + M_{\text{output}}. \quad (3.5)$$

The plan owns prepared data that are fixed during repeated applications. The workspace owns mutable scratch storage with stable allocation, while the caller owns the physical state and requested output. Later memory results use this accounting to expose any dense operator storage and to prevent caller-owned arrays from being counted twice.

3.4 Reference-to-current kinematics and material response

The evolving state is evaluated from a reference configuration rather than from the previously deformed mesh. Let $\mathbf{X} \in \Omega_0$ be a material point. Prescribed motion is represented by

$$\mathbf{x} = \boldsymbol{\chi}(\mathbf{X}, t), \quad \mathbf{F} = \frac{\partial \boldsymbol{\chi}}{\partial \mathbf{X}}, \quad J = \det \mathbf{F}. \quad (3.6)$$

Each update applies $\boldsymbol{\chi}$ to the stored reference coordinates. Accepted steps, rejected trials, and direct evaluations at the same time therefore produce the same geometry within numerical precision. No interpolation drift accumulates from one deformation to the next.

For a rigid volume v ,

$$\boldsymbol{\chi}_v(\mathbf{X}, t) = \mathbf{a}_v + \mathbf{Q}_v(t)(\mathbf{X} - \mathbf{a}_v) + \mathbf{c}_v(t), \quad \mathbf{F} = \mathbf{R} = \mathbf{Q}_v, \quad J = 1, \quad (3.7)$$

where $\mathbf{a}_v$ is an anchor, $\mathbf{Q}_v \in SO(3)$ is the prescribed rotation, and $\mathbf{c}_v$ is the translation. The rotation acts on mesh points, crystal axes, and all other vectors carried by the material frame. Translation moves the points but leaves these vectors unchanged.

A smooth morph allows each component of $\boldsymbol{\chi}_v$ to depend on reference position as well as time. The local deformation gradient is evaluated by centered differences,

$$\frac{\partial \boldsymbol{\chi}}{\partial X_i} \simeq \frac{\boldsymbol{\chi}(\mathbf{X} + h\mathbf{e}_i, t) - \boldsymbol{\chi}(\mathbf{X} - h\mathbf{e}_i, t)}{2h}, \quad h = 10^{-6} \max(1, |X|, |Y|, |Z|). \quad (3.8)$$

The local material rotation is then taken from the orthogonal polar factor

$$\mathbf{F} = \mathbf{R}\mathbf{U}, \quad \mathbf{R}_{k+1} = \frac{1}{2}(\mathbf{R}_k + \mathbf{R}_k^{-T}), \quad \mathbf{R}_0 = \mathbf{F}. \quad (3.9)$$

The polar iteration determines the rotation, but it does not test element validity. Therefore, the determinant J is checked separately to ensure that the current element remains positively oriented.

At a vertex shared by several prescribed volume maps, DynaMag uses the deterministic

conforming coordinate

$$\mathbf{x}_i(t) = \frac{1}{|\mathcal{G}_i|} \sum_{v \in \mathcal{G}_i} \chi_v(\mathbf{X}_i, t). \quad (3.10)$$

This average produces one conforming coordinate field at a shared vertex. It cannot make incompatible volume motions physically continuous, so the prescribed maps must still agree at their interface and retain positive Jacobians.

A reference material direction $\mathbf{a}_0$ is transported as

$$\mathbf{a}(\mathbf{X}, t) = \mathbf{R}(\mathbf{X}, t) \mathbf{a}_0. \quad (3.11)$$

The same material point carries its scalar properties. Quantities such as M_s , exchange, anisotropy, damping, conductivity, and permeability may be constant or tabulated in time. For a tabulated property, the value between two samples is

$$p(t) = (1 - \lambda)p_j + \lambda p_{j+1}, \quad \lambda = \frac{t - t_j}{t_{j+1} - t_j}. \quad (3.12)$$

Material identity is kept separate from geometric identity. Two bodies made from the same material can move independently, and one moving body can contain several materials.

Time is only one possible constitutive input. More generally, a material parameter can be written as

$$p = p(\mathbf{X}, t, T, \boldsymbol{\varepsilon}, \boldsymbol{\sigma}, \mathbf{J}, \mathbf{H}, \phi, c_i, \mathbf{q}_{\text{user}}), \quad (3.13)$$

where the arguments describe the local coupled state. Figure 3.3 shows possible inputs to A_{ex} , M_s , K_1 , and K_2 . The tests in this chapter cover constant and time-tabulated laws. The same material interface admits other coupled variables without changing the field interface.

The exchange energy shows how geometry and material state meet in one field operator. In the current configuration it is

$$E_{\text{ex}} = \int_{\Omega(t)} A_{\text{ex}}(\mathbf{x}, t) |\nabla_{\mathbf{x}} \mathbf{m}|^2 dV. \quad (3.14)$$




Figure 3.3. Dynamic material model in DynaMag. Material parameters such as A_{ex} , M_s , K_1 , and K_2 may depend on time, temperature, and the coupled physical state. These dependencies are evaluated during the simulation without changing the mesh connectivity.

Using the map in Equation (3.6), this integral can be pulled back to the reference mesh:

$$E_{\text{ex}} = \int_{\Omega_0} A_{\text{ex}}(\mathbf{X}, t) J (\mathbf{F}^{-T} \nabla_{\mathbf{X}} \widehat{\mathbf{m}}) : (\mathbf{F}^{-T} \nabla_{\mathbf{X}} \widehat{\mathbf{m}}) dV_0. \quad (3.15)$$

The basis incidence is unchanged because the connectivity is fixed. The current element action, however, still depends on $\mathbf{F}$, J , and A_{ex} . For a rigid motion, $\mathbf{F} = \mathbf{Q}$ and $J = 1$, so the orthogonality of $\mathbf{Q}$ preserves the exchange metric. A morph changes that metric, and its geometric coefficients must be recomputed even though the sparse graph is unchanged.

For uniaxial anisotropy,

$$E_{\text{u}} = \int_{\Omega(t)} K_1(t) [1 - (\mathbf{m} \cdot \mathbf{a})^2] dV, \quad \mathbf{a} = \mathbf{R}\mathbf{a}_0, \quad (3.16)$$

and the corresponding field is proportional to

$$\mathbf{H}_u = \frac{2K_1}{M_s} (\mathbf{m} \cdot \mathbf{a}) \mathbf{a}. \quad (3.17)$$

The anisotropy field therefore requires $\mathbf{m}$, the transported axis, K_1 , and M_s from the same trial state. Magnetostatic and eddy current operators add current normals, distances, curls, and conductivity. These mathematical dependencies determine when a prepared operator has become invalid.

3.5 Accepted-state evolution on accelerators

At a trial time t_{n+1} , DynaMag first evaluates the coordinates and local frames from the reference configuration. It then refreshes finite-element geometry, invalidates coordinate-dependent field data, and samples the material laws at the same time. Only after this state is internally consistent does it form the residual. Magnetization and electromagnetic history are committed after acceptance, while a rejected trial leaves the accepted state unchanged and requires no inverse deformation.

Repeated operators then act on this trial state. Topology-dependent graphs and finite-element incidence remain fixed, while sparse values, Green data, projections, interpolation maps, and near corrections are refreshed according to their coordinate dependence. Persistent work arrays avoid large allocations inside field evaluations and Krylov iterations.

This lifecycle is exercised by the rigid-motion LLG path and by the mesh, finite-element, and material tests for general morphs. Electromagnetic coupling introduces the same dependencies for volume traces, B-AIM geometry, near-field data, and preconditioners. These objects must be refreshed as one set before a moving-conductor calculation is advanced.

3.6 Verification of dynamic geometry, frames, and materials

Verification follows the layers in Figure 3.2. A compute primitive can agree across providers even when a complete physical workflow has not yet been tested on all of them. Table 3.6 keeps

these evidence levels separate.

Table 3.6. Evidence levels used for DynaMag capability claims.

Level	Required evidence	Example
Compute primitive	Numerical agreement and execution on the named provider	FFT round trip, sparse action, reduction
Discrete operator	Agreement for the assembled or matrix-free action	Exchange, Poisson, B-AIM potential or gradient
Physics solve	Convergence and an independent physical comparison	Demagnetizing field or eddy-current transient
Coupled workflow	Accepted-state history and end-to-end validation	Dynamic LLG with open-boundary Maxwell feedback
Performance result	Recorded inputs, timing scope, repetitions, hardware, software, and memory convention	Publication benchmark

The first tests isolate kinematics from magnetic evolution. A two-volume mesh undergoes independent translation and rotation, while the same test exercises the shared-vertex rule in Equation (3.10), refreshed bounds, and direct jumps in time. Rigid rotations and vector transport agree with their analytical values to 10^{-12} , while coordinates and shared interfaces agree to 10^{-6} . The accelerator kernel reads the same frame and coordinate views as the host calculation.

For the non-rigid map

$$\chi(X, Y, Z) = (X, Y + X^2, Z), \quad \mathbf{F} = \begin{bmatrix} 1 & 0 & 0 \\ 2X & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (3.18)$$

the polar rotation is the identity at $X = 0$ and a 45° in-plane rotation at $X = 1$. The computed position, $\mathbf{F}$, $\mathbf{J}$, $\mathbf{R}$, and transported vector agree with the analytical values to 10^{-6} . A dome and a bar twisted through π then provide larger visual checks of spatially varying frames (Figure 3.4).

The next test connects the transported frame to a magnetic field. The material frame of a one-element cube rotates by 45° while the magnetization remains fixed in the laboratory. The normalized anisotropy field should be $(0, -2, 0)$; it is recovered to 10^{-12} on the host and 2×10^{-6}




Figure 3.4. A dome-morphed sheet and a bar twisted by π . These images show the spatial structure. The analytical shear test provides the numerical acceptance criterion.

on the accelerator. Separate interpolation tests recover the exact midpoint values of tabulated M_s , surface charge, and K_1 laws. Together, these tests show that geometry, orientation, and constitutive value are evaluated as distinct parts of one trial state.

A final application-scale test combines the geometry, frame, and material checks in one prescribed morph. The reference disk body has a radius of 75 nm and a thickness of 4 nm, with 50,688 hexahedral elements and 59,815 vertices. A radial Airy-amplitude profile displaces the mid-surface. The fifth positive zero of J_1 is placed at the rim, and a nonlinear contrast map resolves the secondary rings. A rim-fixed correction gives the profile zero area mean before it is normalized to an 8 nm central height. The activation

$$s(\tau) = \sin^2(\pi\tau), \quad 0 \leq \tau \leq 1, \quad (3.19)$$

produces a 21-frame flat–morphed–flat cycle without changing connectivity.

The disk thickness follows the current surface normal rather than the laboratory z axis. At the peak morph, the mid-surface height ranges from -6.160293 to 8.000000 nm, and the maximum vertex displacement is 8.000000 nm. The minimum absolute element-center Jacobian is $7.341480 \times$

10^{-2} , compared with 9.595086×10^{-2} in the reference disk. No element reverses orientation or becomes degenerate. The last frame returns exactly to the reference coordinates, which confirms that each configuration is evaluated from the reference state.

To test state transport on this geometry, the prescribed magnetization at each vertex is aligned with the local surface normal. Figure 3.5 shows the vector field following the central lobe and secondary rings. At the peak morph, these vectors are compared with normals reconstructed independently from the top surface. The maximum angular difference is 3.0848° , and the RMS difference is 0.6429° . The largest differences occur on steep rings where the faceted surface limits the normal reconstruction.

This comparison verifies the application-scale morph and the transport of a prescribed magnetization. Because it is a kinematic test rather than an LLG trajectory, it does not claim magnetic equilibrium. An LLG calculation on the same disk would also require the geometry-dependent long-range fields to be updated at each trial state.

3.7 Dynamic LLG cost and scaling

A matched rod-switch experiment measures the cost of adding dynamic state to an LLG trajectory. The CUDA runs compare a fixed mesh with constant K_1 , the same mesh with a constant-valued three-sample $K_1(t)$ table, and a rigidly translated mesh. All three cases use the same 4,250-hexahedron mesh, fields, solver tolerances, and output schedule. Translation does not change the physical problem, so the magnetization histories provide a covariance check. The median first $\langle m_z \rangle = 0$ crossings are 0.172197, 0.172270, and 0.171977 ns. The six-run ranges overlap.

The median trajectory times are 4.3397, 4.3916, and 5.1103 s. The material-table case is 1.2% above the static median, but its bootstrap interval includes zero. Rigid translation raises the median by 17.8%, although only 0.2041 s, or 3.9% of the moving total, is spent on coordinates and frames. Most of the remaining difference comes from conservative rebuilding of the hybrid magnetostatic B-AIM plan and from variation in the adaptive solve. Material lookup and local mesh

motion are therefore small in this experiment, while magnetostatic replanning is the main dynamic overhead.

The rod experiment isolates the additional work caused by state evolution. Figure 3.7 instead measures end-to-end spatial scaling of the same architecture. The cube series spans about 1.8×10^4 to 3.09×10^6 degrees of freedom, while the nonlinear step count stays near 210. The increase in runtime is therefore caused mainly by the larger spatial operators rather than by a systematic increase in the number of time steps. Because the series is not a matched static–moving comparison, it is not used to estimate motion overhead.

Together, the tests cover rigid motion, prescribed morphing, material-frame transport, and time-dependent parameters on accelerator paths. They establish the fixed-connectivity state model used in the following chapters. Remeshing and topology change require a different update of degrees of freedom and are outside this model.




Figure 3.5. Airy-disk morph at zero, half, and peak activation. The connectivity remains fixed. The thickness and prescribed magnetization both follow the local surface normal. This test covers geometry, material frames, and state transport. It is not an LLG time integration.




Figure 3.6. Matched dynamic LLG experiment. The three trajectories overlap. Material lookup is not measurably slower than the static baseline. The mesh and frame update is a small part of the moving run. Conservative magnetostatic replanning accounts for much of the remaining cost.




Figure 3.7. Cube-simulation scaling from about 1.8×10^4 to 3.09×10^6 degrees of freedom. The step count is nearly constant. The higher cost therefore comes from the larger spatial operators.

4 Hybrid Superposition Integral–Poisson Magnetostatics

Micromagnetic simulators solve the Landau–Lifshitz–Gilbert equation and allow the static and dynamic magnetization of magnetic materials and devices to be characterized. Applications include soft magnetic materials, magnetic recording heads and media, magnetic random-access memories, and permanent magnets. Efficient simulation of large structures with complex geometry requires fast methods for both the effective fields and the time integration (Duan et al., 2023).

An important part of an LLG solver is the magnetostatic field. Finite difference methods are efficient for relatively simple structures and usually compute this field by FFT. Finite-element methods are better suited to structures with complex geometry and material distributions, but their meshes are nonuniform. The magnetostatic field is therefore more difficult to compute. Existing approaches solve a Poisson equation, evaluate the superposition integral with a fast method, or combine a finite-element Poisson problem with an integral evaluation on a truncation boundary (Chang et al., 2011; Fredkin and Koehler, 1990; Duan et al., 2023).

The ingredients of these alternatives come from several numerical lineages. Finite-element and finite-difference field discretizations were established early for Helmholtz and Maxwell problems (Arlett et al., 1968; Yee, 1966). Fast multipole, adaptive integral, and precorrected-FFT methods then reduced the cost of dense long-range interactions (Greengard and Rokhlin, 1987; Bleszynski et al., 1996; Phillips and White, 1997). Accurate local replacement can use closed-form panel and volume potential integrals (Wilton et al., 1984), while GPU implementations have shown

high-throughput Helmholtz-potential and electromagnetic integral-equation evaluation (Li et al., 2010a, 2012). Magnetostatic-specific boundary and nonuniform-grid methods apply these ideas to finite-element micromagnetics (Forster et al., 2003; Livshitz et al., 2009). The hybrid method below changes which targets require the expensive integral action; it does not change the underlying open-boundary Green function.

A Poisson method may require a large exterior domain before a zero boundary condition is accurate. The added air mesh increases the number of elements. An integral method avoids this mesh, but accurate near interactions may require large sparse matrices. This cost is severe for nonuniform meshes. A magnetic write head, for example, is much more finely resolved near its tip than in the return pole. The near-interaction matrices can then limit the largest problem that fits in GPU memory.

The hybrid superposition integral–Poisson method combines the two approaches. It evaluates the scalar potential by fast superposition only on a truncation surface. These surface values become Dirichlet data for a Poisson equation in the enclosed volume. The method reduces the near-zone matrices, gives a rapidly convergent Poisson problem, and does not require a large airbox. Additional surfaces can also divide the volume into smaller independent Poisson problems (Duan et al., 2023).

This chapter first follows the formulation and numerical demonstrations of the published method. The final section then describes the extension used by DynaMag. It evaluates both near corrections on the fly and stores neither correction matrix. This extension is important for a moving or morphing mesh, where a geometry-dependent near matrix would otherwise have to be rebuilt.

4.1 Hybrid integral–Poisson method

4.1.1 Problem formulation

Consider a computational domain that contains a magnetized region and, optionally, a nonmagnetic region. The magnetic structure is discretized by a tetrahedral finite-element mesh.

With piecewise-linear nodal basis functions, the magnetization is

$$\mathbf{M}_h(\mathbf{x}) = \sum_{j=1}^{N_V} \mathbf{M}_j \varphi_j(\mathbf{x}), \quad (4.1)$$

where N_V is the number of mesh vertices, and $\mathbf{M}_j$ and φ_j are the magnetization and basis function associated with vertex j .

The magnetostatic field is written in terms of a scalar potential,

$$\mathbf{H}_{\text{ms}} = -\nabla \phi. \quad (4.2)$$

In the Gaussian convention used by DynaMag, the potential is

$$\phi(\mathbf{r}) = - \int_{\Omega_m} \frac{\nabla' \cdot \mathbf{M}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dV' + \int_{\partial\Omega_m} \frac{\hat{\mathbf{n}}' \cdot \mathbf{M}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dS'. \quad (4.3)$$

The first term is produced by volume magnetic charge, and the second is produced by the discontinuity of magnetization at the material surface. The $1/r$ kernel is used here, so $\nabla^2(1/r) = -4\pi\delta(\mathbf{r})$. This makes the factor 4π explicit in the Poisson equation below. It can instead be absorbed into the Green function and the charge definition.

Following the integral formulation in the paper, substitution of Equation (4.1) into Equation (4.3) gives the discrete approximation

$$\boldsymbol{\phi} \simeq \mathbf{Z}_\ell \mathbf{Q} \mathbf{M} + \mathbf{Z}_0 \mathbf{M}. \quad (4.4)$$

The sparse matrix $\mathbf{Q} \in \mathbb{R}^{N_V \times 3N_V}$ projects nodal magnetization to effective nodal charges. Its nonzero entries connect nodes that belong to the same tetrahedron. The dense point-interaction matrix is

$$(\mathbf{Z}_\ell)_{ij} = \begin{cases} 1/R_{ij}, & i \neq j, \\ 0, & i = j, \end{cases} \quad R_{ij} = |\mathbf{r}_i - \mathbf{r}_j|. \quad (4.5)$$

The product $\mathbf{Z}_\ell \mathbf{Q} \mathbf{M}$ is a canonical point-to-point superposition sum. The mesh description is contained in $\mathbf{Q}$, while the point interaction is independent of the element type. A hexahedral mesh, for example, changes the charge projection but not $\mathbf{Z}_\ell$.

The matrix $\mathbf{Z}_0$ corrects the finite-element near zone. Far from an element, its source integral can be represented accurately by nodal quadrature. Close to an element, the same approximation is inaccurate, and the tetrahedral source integral is evaluated analytically. The correction replaces the nodal approximation by this accurate integral. Each element has only a finite number of nearby targets, so $\mathbf{Z}_0$ is sparse. Its size can still be large when the mesh is strongly nonuniform or contains many vertices.

Direct application of $\mathbf{Z}_\ell$ has $O(N_V^2)$ cost. Fast methods reduce this cost to $O(N_V)$ or $O(N_V \log N_V)$. The published implementation uses the FFT-based box-adaptive integral method (B-AIM). The domain is divided into boxes, and the charges are projected to nodes on a regular grid. The charges and the kernel are transformed by FFT, multiplied in Fourier space, and transformed back. The grid potential is then interpolated to the nonuniform mesh targets.

Grid interpolation is not accurate for nearby source and target points. A second sparse correction subtracts the grid-mediated near interaction and adds the direct point interaction. The fast action is therefore written as

$$\boldsymbol{\phi}_s \simeq \boldsymbol{\phi}_s^{\text{far}} + \mathbf{Z}_{s0} \mathbf{q}, \quad \mathbf{q} = \mathbf{Q} \mathbf{M}. \quad (4.6)$$

The two corrections have different purposes. The matrix $\mathbf{Z}_0$ corrects the finite-element source integral. The matrix $\mathbf{Z}_{s0}$ corrects the grid-mediated point interaction. The point correction is simpler and can be evaluated on the fly from neighboring boxes. This choice was already used in the published memory study, which therefore reports the storage of $\mathbf{Z}_0$, not $\mathbf{Z}_{s0}$ (Duan et al., 2023).

Figure 4.1 shows both near zones. An interior target is surrounded by a complete three-dimensional neighborhood. A boundary target has a smaller, one-sided neighborhood. Moving the truncation surface a short distance away from the magnetic body can eliminate the analytical

finite-element near zone completely.




Figure 4.1. Geometry, meshing, and near zones of the hybrid method. Red tetrahedra contribute to the analytical finite-element correction $\mathbf{Z}_0$. Red boxes contribute to the B-AIM point correction $\mathbf{Z}_{s0}$. The neighborhoods are smaller for a boundary target. Reproduced from Fig. 1 of Duan et al. (2023), © 2023 IEEE.

4.1.2 Surface potential and interior Poisson equation

An alternative to the full integral calculation is to solve the Poisson equation. In the present convention,

$$\nabla^2 \phi = 4\pi \nabla \cdot \mathbf{M} \quad \text{in } \Omega, \quad (4.7)$$

together with the material-interface conditions and decay at infinity. A pure finite-element treatment usually adds a large airbox so that the remote boundary potential can be approximated as zero. Classical hybrid methods avoid that large domain, but may require two Poisson solves and more complicated integral kernels (Fredkin and Koehler, 1990; Duan et al., 2023).

The method used here has two steps. First, the superposition integral is evaluated only on the truncation surface. Second, the resulting potential is used as a Dirichlet boundary condition for

the Poisson equation. If Γ denotes the truncation surface, the problem is

$$\nabla^2 \phi = 4\pi \nabla \cdot \mathbf{M} \quad \text{in } \Omega, \quad (4.8)$$

$$\phi = g \quad \text{on } \Gamma, \quad (4.9)$$

where g is obtained from Equation (4.3). The truncation surface may coincide with the magnetic surface, so an airbox is not required. It may also be placed a short distance away when that choice reduces the analytical near correction.

The boundary values determine the interior solution uniquely. The free-space superposition potential and the Dirichlet solution satisfy the same Poisson equation and have the same value on Γ . Their difference is harmonic and zero on the boundary, so it is zero throughout the connected domain. Restricting superposition to the surface therefore adds no continuum approximation.

The paper states the Poisson step in strong form. For completeness, its finite-element form is written here. Let $V_h \subset H^1(\Omega)$ be the nodal finite-element space and let $V_{h,0}$ contain functions with zero trace on Γ . The weak problem is to find ϕ_h with $\phi_h = g_h$ on Γ such that

$$\int_{\Omega} \nabla v_h \cdot \nabla \phi_h dV = 4\pi \int_{\Omega} \nabla v_h \cdot \mathbf{M} dV \quad \forall v_h \in V_{h,0}. \quad (4.10)$$

The boundary term from integration by parts vanishes because v_h has zero trace. The exterior field has not been dropped; its effect is contained in g_h .

For nodal basis functions N_i , define

$$q_i = \int_{\Omega} \nabla N_i \cdot \mathbf{M} dV, \quad K_{ij} = \int_{\Omega} \nabla N_i \cdot \nabla N_j dV. \quad (4.11)$$

Partitioning the vertices into interior and boundary sets gives

$$\mathbf{K}_{II} \boldsymbol{\phi}_I = 4\pi \mathbf{q}_I - \mathbf{K}_{IB} \mathbf{g}_B, \quad \boldsymbol{\phi}_B = \mathbf{g}_B. \quad (4.12)$$

The resulting Dirichlet system is well behaved. It can be solved by a Krylov method with an incomplete factorization, Jacobi, or another standard Poisson preconditioner. The magnetostatic field is recovered by a lumped L^2 projection of $-\nabla\phi_h$ to the nodal LLG field.

Let $\mathbf{R}_\Gamma$ select the truncation-surface targets. The boundary stage of the hybrid method is

$$\mathbf{g}_\Gamma \simeq \mathbf{R}_\Gamma \boldsymbol{\phi}_s^{\text{far}} + \mathbf{R}_\Gamma \mathbf{Z}_{s0} \mathbf{Q} \mathbf{M} + \mathbf{R}_\Gamma \mathbf{Z}_0 \mathbf{M}. \quad (4.13)$$

Only the rows of $\mathbf{Z}_0$ associated with the surface are required. This is the main memory reduction. The FFT and source projection remain global because every magnetic source contributes to each boundary value.

The hybrid approach also permits additional internal surfaces. Once the global superposition potential is known on these surfaces, the enclosed Poisson problems are independent. This allows the domain to be divided into smaller problems that can run in parallel. It can also reduce the iteration count and the memory required by each Poisson solve. A thin-film region may instead be treated entirely by the integral method when a volume Poisson solve does not provide an advantage (Duan et al., 2023).

4.2 Numerical demonstration

The published method was implemented in the FastMag micromagnetic simulator on both CPUs and GPUs. B-AIM supplied the integral evaluation and retained only the near-zone entries associated with truncation-surface targets. The Poisson equation was solved by GMRES with an ILU preconditioner and a relative tolerance of 10^{-6} . The near zone of $\mathbf{Z}_0$ was a sphere whose diameter was 2.25 times the edge length of each tetrahedron. The B-AIM grid contained about 16 mesh vertices per box and used two layers of neighboring boxes. The point correction $\mathbf{Z}_{s0}$ was evaluated on the fly. All timings in the published study were measured on an NVIDIA RTX 3080 Ti (Duan et al., 2023).

4.2.1 Cylinder validation and refinement study

The first test is a magnetic cylinder with a height of 300 nm, a diameter of 20 nm, and a saturation magnetization of 796 emu/cm^3 . A 2 nm tetrahedral mesh gives 94,833 elements and 18,722 vertices. The exchange constant is 10^{-6} erg/cm , the uniaxial anisotropy is $K_1 = 10^5 \text{ erg/cm}^3$ with its easy axis along z , and the damping constant is $\alpha = 0.2$. The applied field increases linearly from zero at 0.1 ns to 20 kOe along negative z at 0.2 ns.

Figure 4.2 compares the axial magnetostatic field from the full B-AIM integral path and the hybrid method. The cylinder has no airbox. The profiles agree at $t = 0, 0.2$, and 0.25 ns, including the strongly varying field during reversal.




Figure 4.2. Axial magnetostatic field in the switching cylinder at $t = 0, 0.2$, and 0.25 ns. The full B-AIM and hybrid results overlap. Reproduced from Fig. 2 of Duan et al. (2023), © 2023 IEEE.

A larger cylinder measures the memory and computational cost. Its height is 900 nm and its diameter is 90 nm. The mesh edge length ranges from 16 to 1 nm, which gives 2,473 to 8,378,659 vertices. Table 4.1 reports the number of nonzero entries in the full finite-element correction $\mathbf{Z}_0$ and in the boundary-restricted correction. Since the point correction is evaluated on the fly, these

counts isolate the expensive finite-element near integration.

Table 4.1. Cylinder refinement results for the full integral and hybrid methods. The time is for one complete magnetostatic-field evaluation on an NVIDIA RTX 3080 Ti. The finest-case time was not reported. Data are from Table I of Duan et al. (2023).

h (nm)	N_V	Full nnz	Hybrid nnz	Reduction	Poisson iters.	Time (s)
16	2,473	0.670 M	0.255 M	2.63×	2	0.0018
8	17,617	5.831 M	1.024 M	5.69×	3	0.012
4	137,718	51.567 M	4.420 M	11.67×	6	0.070
2	1,069,326	417.538 M	17.904 M	23.32×	12	0.54
1	8,378,659	3,186.095 M	68.059 M	46.81×	21	—

The memory required by the full correction grows linearly with the number of vertices. At 8,378,659 vertices, it occupies 46.68 GB. The boundary-restricted correction occupies 551.22 MB. The Poisson solve remains rapidly convergent: its iteration count grows from 2 to 21 while the mesh grows by more than three orders of magnitude. At 1,069,326 vertices, one field evaluation takes 0.54 s. The full integral and hybrid times are nearly the same (Duan et al., 2023).

An air layer is not required, but a small offset can reduce the correction further. Figure 4.3 shows this effect for the 4 nm cylinder mesh. The boundary correction falls as the padding distance increases and is eliminated at 8 nm. Only a small number of air vertices are added, and the field-evaluation time remains nearly unchanged.

4.2.2 Multiscale write head

The second test is a perpendicular recording write head. It contains the pole tip, wrap-around shield, main pole, yoke, return pole, and soft underlayer. The mesh has 589,494 vertices and 3,049,001 tetrahedra. The edge length is 5–25 nm near the pole tip and 35–100 nm in the coarser parts of the return path.

Without an airbox, the full B-AIM finite-element correction has 295,866,999 nonzero entries and occupies 2259.53 MB. Restricting the targets to the physical boundary reduces it to 57,620,895 entries and 440.17 MB. Most near interactions occur near the finely meshed pole tip. A small




Figure 4.3. Size of the finite-element near-correction matrix versus airbox padding distance for the 4 nm cylinder mesh. The offset is used to remove near interactions, not to approximate a zero condition at a remote boundary. Reproduced from Fig. 3 of Duan et al. (2023), © 2023 IEEE.

air layer around that region reduces the matrix to 31,853,472 entries and 245.23 MB, with only a small increase in the number of vertices. These changes give fivefold and nearly tenfold memory reductions. One complete magnetostatic-field evaluation takes 0.04 s (Duan et al., 2023).

The cylinder and write-head studies show the two main benefits reported in the paper. The integral work is restricted to a surface, so the accurate near correction uses much less memory. The remaining volume problem has a Dirichlet boundary condition and converges in a small number of iterations. The truncation surface can be the magnetic surface or a small local offset. It does not need to be a large airbox.

4.3 On-the-fly DynaMag extension

The published hybrid method reduces $\mathbf{Z}_0$ from all volume targets to the truncation surface. DynaMag removes the remaining matrix storage. Neither the B-AIM point correction $\mathbf{Z}_{s0}$ nor the finite-element analytical correction $\mathbf{Z}_0$ is stored. Both are evaluated from the current geometry during each field application. Zero correction-matrix memory does not mean zero solver memory.




Figure 4.4. Write-head model used in the memory study: (a) no airbox; (b) a small air layer around the refined tip; (c) the fine pole-tip mesh; and (d) the coarse return-pole mesh. Reproduced from Fig. 4 of Duan et al. (2023), © 2023 IEEE.

The mesh, FFT grid, Poisson matrix, preconditioner, state, and work arrays are still required.

The finite-element correction is the expensive part of this extension. Analytical tetrahedron–target integration requires more work than a direct point interaction. Applying it at every volume target would be too costly. The hybrid method makes the on-the-fly path practical because the analytical work is limited to surface targets and their one-sided neighborhoods.

Figure 4.5 measures this tradeoff for a separate write-head mesh with about 533,000 vertices. Stored all-volume superposition takes 23.50 ms per field application. The on-the-fly hybrid path takes 26.78 ms and stores neither correction matrix. The increase is 3.28 ms, or 14.0%. In the hybrid result, B-AIM takes 5.2 ms, the streamed finite-element correction takes 12.3 ms, and the Poisson solve takes 7.9 ms.




Figure 4.5. Magnetostatic-field timing for the approximately 533,000-vertex write-head comparison. The bold values are measured application totals. Both near corrections have zero persistent matrix storage in the on-the-fly hybrid path. This benchmark is separate from the published 589,494-vertex write-head study.

This change is especially useful for dynamic geometry. An assembled finite-element correction depends on the target coordinates and the current tetrahedral geometry,

$$\mathbf{Z}_0(t) = \mathbf{Z}_0[\boldsymbol{\chi}(t)]. \quad (4.14)$$

A morph changes its entries and may also change its near-pair pattern. The matrix would have to be rebuilt after each geometry update. DynaMag instead evaluates the surface correction directly

from $\chi(t)$. Fixed mesh connectivity preserves the finite-element graph and allocated state, while the current coordinates update the B-AIM plan and local integrations.

For a rigid body, internal distances and element shapes are unchanged in the body frame. Relative motion between bodies still changes the long-range interaction and their positions on the auxiliary grid. A morph also changes element Jacobians and analytical source integrals. In both cases, the hybrid method avoids a meshed exterior region that must move or be regenerated with the magnetic body.

The B-AIM source projection and FFT remain global. Their implementation and large-scale performance are examined in Chapter 5. The result here is narrower: boundary restriction makes accurate finite-element near integration small enough to evaluate on the fly. This removes the large geometry-dependent correction matrix while preserving the open-boundary magnetostatic solution.

Publication disclosure

This chapter is based in part on “Hybrid Superposition Integral–Poisson Solver Method for the Magnetostatic Field in Finite Element Micromagnetic Solvers” by Jiawei Duan, Xueyang Wang, and Vitaliy Lomakin, published in *IEEE Transactions on Magnetics*, vol. 59, no. 11, 2023, doi:10.1109/TMAG.2023.3297951. The dissertation author was the primary author of this paper.

5 Periodic Green Functions and B-AIM

Acceleration

Magnetostatic and eddy current boundary operators couple every source to every target, while in a periodic structure each source also represents an infinite lattice of images. The resulting Green function cannot be obtained by simply adding more free-space images because its convergence, scalar gauge, and zero mode depend on how many directions are periodic.

This chapter derives the periodic $1/r$ potential and its analytic $\mathbf{r}/r^3$ gradient, then shows how B-AIM evaluates these all-to-all actions without forming dense matrices. Together, these operators provide the long-range foundation for the FEM–BEM formulations in Chapter 6.

The scalar potential and its gradient require separate attention. An ordinary periodic Poisson inverse requires a neutral source, but with 1D or 2D periodicity the divergent part of a nonneutral scalar potential is only an additive constant. A finite-part convention can therefore remove this constant while retaining the finite field of the mean line or sheet source. The $H-k$ formulation uses the resulting $\mathbf{r}/r^3$ action in its magnetic trace. The $H-\kappa$ formulation uses both the scalar single layer V and its target-normal derivative K' . In neither case may a physically nonzero partial-period mode be removed by an automatic mean subtraction.

Periodic Green-function construction has a long history across several physical settings. Spectral and lattice representations appear in surface, grating, and periodic-waveguide analyses (Kalkstein and Soven, 1971; Hessel and Oliner, 1965; Peng et al., 1975), while Coulombic and Helmholtz lattice sums motivate explicit convergence prescriptions (Marshall, 2000; Capolino

et al., 2007). Rapidly convergent linear-array formulas extend to layered media (Van Orden and Lomakin, 2009, 2012); related modern approaches treat quasi-periodic scattering, periodic multiscale interactions, and arbitrarily shaped two-dimensional lattices (Barnett and Greengard, 2011; Baczewski et al., 2012; Pei et al., 2023).

Ewald methods make the summation convention explicit and expose the role of the zero and boundary modes. Boundary terms in electrostatic Ewald sums are an essential part of that interpretation (Hu, 2014). Fast Ewald methods now cover electrostatic and Stokes potentials with free-space or arbitrary periodicity (af Klinteberg et al., 2017; Shamshirgar et al., 2021; Bagge and Tornberg, 2023). The formulas below specialize these general ideas to the static scalar and Cartesian gradient kernels required by DynaMag.

Periodic unit-cell operators have also been used for transmission through aperture arrays (Lomakin and Michielssen, 2005, 2006, 2007). In magnetics, film dispersion, disordered dipolar systems, periodic demagnetization, and hard–soft nanocomposites provide distinct tests of the same long-range lattice physics (Kalinikos and Slavin, 1986; Berkov and Gorn, 1998; Lebecki et al., 2008; Wysocki and Antropov, 2017). True periodic finite elements further require the quotient topology to agree with the Green operator (Bruckner et al., 2021).

The acceleration strategy likewise draws on a broad fast-solver lineage: fast electromagnetic integral methods, sparse hierarchical matrices, polar grids, and multilevel Cartesian grids (Chew et al., 1997; Hackbusch and Khoromskij, 2000; Boag et al., 2002; Meng et al., 2010). FFT-accelerated layered and periodic interpolation methods move the nonuniform interaction to a structured convolution (Liu et al., 2020; Li et al., 2010b; Ai and Lomakin, 2024). B-AIM uses that factorization together with the matched local replacement described later in this chapter.

5.1 Conventions and the static-limit construction

5.1.1 Normalization, displacement, and sign

Let $\mathbf{x}$ be the target, $\mathbf{y}$ the source, and $\mathbf{r} = \mathbf{x} - \mathbf{y}$. Two kernel normalizations appear in the derivation. The standard fundamental solution is

$$\mathcal{G}_0(\mathbf{r}) = \frac{1}{4\pi|\mathbf{r}|}, \quad \nabla_{\mathbf{x}}\mathcal{G}_0(\mathbf{r}) = -\frac{\mathbf{r}}{4\pi|\mathbf{r}|^3}, \quad (5.1)$$

whereas the DynaMag kernel interface uses

$$G_0 = 4\pi\mathcal{G}_0 = \frac{1}{r}, \quad \mathbf{K}_0 = -4\pi\nabla\mathcal{G}_0 = \frac{\mathbf{r}}{r^3}. \quad (5.2)$$

The FEM–BEM field equation supplies the sign and the factor $1/(4\pi)$ when this public kernel is used. Equation (5.2) is therefore an interface convention rather than a different physical Green function. The distinction is important in verification: reversing $\mathbf{x} - \mathbf{y}$ changes the sign, while applying the 4π factor twice changes the magnitude. Direct tests check all three choices together.

For a lattice Λ_p that is periodic in p coordinate directions, the formal static gradient is

$$\mathbf{D}_p(\mathbf{r}) = \nabla\mathcal{G}_p(\mathbf{r}) = -\frac{1}{4\pi} \sum_{\mathbf{R} \in \Lambda_p} \frac{\mathbf{r} - \mathbf{R}}{|\mathbf{r} - \mathbf{R}|^3}. \quad (5.3)$$

Away from source images, the gradient sum for 1D periodicity is absolutely convergent. The sums for 2D and 3D periodicity are only defined after a symmetric or Ewald summation prescription is chosen for the scalar Green function. Changing the order of an infinite conditionally convergent image sum can change its value.

5.1.2 Differentiation before the static limit

The safest static construction begins before the static limit. Consider the outgoing Helmholtz fundamental solution

$$\mathcal{G}_k(\mathbf{r}) = \frac{e^{-ikr}}{4\pi r}. \quad (5.4)$$

Poisson summation converts the lattice images into Floquet modes. Away from a source, the resulting series may be differentiated term by term. The order of operations is

$$\text{image sum} \rightarrow \text{Floquet modes} \rightarrow \nabla_{\mathbf{r}} \rightarrow k \rightarrow 0. \quad (5.5)$$

If the limit $k \rightarrow 0$ is taken first, the scalar zero mode diverges for 1D and 2D periodicity. Differentiation before the limit removes the position-independent divergence and retains the finite field of the mean line or sheet source. This is why a gradient action may admit a nonzero source mean even when an ordinary periodic-Poisson scalar inverse does not.

5.2 Periodic Green functions by dimensionality

5.2.1 1D periodicity

First consider a lattice that repeats along x with period L . The distance from the periodic axis is

$$\rho = (y^2 + z^2)^{1/2}, \quad a_m = \frac{2\pi m}{L}. \quad (5.6)$$

In the $1/(4\pi r)$ normalization, Poisson summation gives the finite-part scalar Green function

$$\mathcal{G}_1(x, \rho) = -\frac{1}{2\pi L} \log \frac{\rho}{\rho_0} + \frac{1}{\pi L} \sum_{m=1}^{\infty} K_0(a_m \rho) \cos(a_m x). \quad (5.7)$$

The reference length ρ_0 fixes only an additive scalar constant and has no effect on the field. We can therefore differentiate the Fourier–Bessel series and use $K'_0(s) = -K_1(s)$ to obtain

$$D_{1,x} = -\frac{1}{\pi L} \sum_{m=1}^{\infty} a_m K_0(a_m \rho) \sin(a_m x), \quad (5.8)$$

$$D_{1,y} = -\frac{y}{2\pi L \rho^2} - \frac{y}{\pi L \rho} \sum_{m=1}^{\infty} a_m K_1(a_m \rho) \cos(a_m x), \quad (5.9)$$

$$D_{1,z} = -\frac{z}{2\pi L \rho^2} - \frac{z}{\pi L \rho} \sum_{m=1}^{\infty} a_m K_1(a_m \rho) \cos(a_m x). \quad (5.10)$$

The explicit transverse term in these equations is the zero Fourier harmonic, not a numerical correction:

$$\mathbf{D}_1^{(0)}(y, z) = -\frac{y\mathbf{e}_y + z\mathbf{e}_z}{2\pi L(y^2 + z^2)}, \quad (5.11)$$

This term is the field of the mean infinite line source. Differentiation has removed the divergent scalar constant but has preserved this finite transverse field.

For point strengths q_j , define $\Phi_1 = \sum_j q_j \mathcal{G}_1(\mathbf{r} - \mathbf{r}_j)$ and let $Q = \sum_j q_j$ be their total strength. If $Q \neq 0$, the discarded scalar constant is multiplied by Q , but the zero-harmonic derivative remains

$$\nabla_{\perp} \Phi_1^{(0)} = -\frac{Q}{2\pi L} \frac{y\mathbf{e}_y + z\mathbf{e}_z}{y^2 + z^2}. \quad (5.12)$$

In the public $G = 4\pi \mathcal{G}$ normalization, the coefficient is $-2Q/L$. A derivative-only action can therefore preserve a nonzero mean. An ordinary periodic-Poisson inverse still requires neutrality because it seeks a bounded periodic scalar potential. The partial-period finite-part action instead fixes the missing additive constant by an explicit reference and retains the nonzero source.

The Fourier–Bessel derivatives are singular term by term on the periodic axis, even though their combined limit is finite away from a source image. For $0 < \xi = x/L < 1$, the combined axial limit is

$$D_{1,x}(x, 0, 0) = -\frac{\psi_1(\xi) - \psi_1(1 - \xi)}{4\pi L^2}, \quad D_{1,y} = D_{1,z} = 0, \quad (5.13)$$

where ψ_1 is the trigamma function. The formula equals the absolutely convergent differentiated

image sum and therefore provides an independent check of sign, scale, and truncation.

5.2.2 2D periodicity

Next let x and y be periodic with periods L_x and L_y , and leave z open. The cell area is $A = L_x L_y$. Applying one-dimensional Poisson summation along x to each y image gives

$$\mathcal{G}_2(x, y, z) = -\frac{1}{2\pi L_x} \sum_{\substack{n \in \mathbb{Z} \\ \text{fp}}} \log \frac{R_n}{\rho_0} + \frac{1}{\pi L_x} \sum_{m=1}^{\infty} \sum_{n=-\infty}^{\infty} K_0(a_m R_n) \cos(a_m x), \quad (5.14)$$

where

$$a_m = \frac{2\pi m}{L_x}, \quad R_n = \sqrt{(y + nL_y)^2 + z^2}. \quad (5.15)$$

The remaining logarithmic image sum is evaluated through the product formula for $\sin(\pi(y + iz)/L_y)$:

$$\sum_{\substack{n \in \mathbb{Z} \\ \text{fp}}} \log \frac{R_n}{\rho_0} = \frac{\pi|z|}{L_y} + \frac{1}{2} \log(1 - 2q \cos \theta + q^2) + C, \quad (5.16)$$

with

$$b = \frac{2\pi}{L_y}, \quad q = e^{-b|z|}, \quad \theta = by, \quad \Delta = 1 - 2q \cos \theta + q^2. \quad (5.17)$$

The constant C fixes the scalar gauge and disappears after differentiation. Substitution into Equation (5.14) gives the Fourier–Bessel representation used for evaluation:

$$\boxed{\mathcal{G}_2(x, y, z) = -\frac{|z|}{2A} - \frac{\log \Delta}{4\pi L_x} + \frac{1}{\pi L_x} \sum_{m=1}^{\infty} \sum_{n=-\infty}^{\infty} K_0(a_m R_n) \cos(a_m x).} \quad (5.18)$$

The public $1/r$ kernel is $G_2 = 4\pi \mathcal{G}_2$.

The Cartesian gradient follows by using $K'_0(s) = -K_1(s)$ and

$$\partial_y R_n = \frac{y + nL_y}{R_n}, \quad \partial_z R_n = \frac{z}{R_n}, \quad (5.19)$$

which gives

$$D_{2,x} = -\frac{1}{\pi L_x} \sum_{m=1}^{\infty} \sum_{n=-\infty}^{\infty} a_m K_0(a_m R_n) \sin(a_m x), \quad (5.20)$$

$$D_{2,y} = -\frac{q \sin \theta}{A \Delta} - \frac{1}{\pi L_x} \sum_{m=1}^{\infty} \sum_{n=-\infty}^{\infty} a_m \frac{y + n L_y}{R_n} K_1(a_m R_n) \cos(a_m x), \quad (5.21)$$

$$D_{2,z} = -\frac{\text{sgn}(z)}{2A} - \frac{\text{sgn}(z) q (\cos \theta - q)}{A \Delta} - \frac{1}{\pi L_x} \sum_{m=1}^{\infty} \sum_{n=-\infty}^{\infty} a_m \frac{z}{R_n} K_1(a_m R_n) \cos(a_m x). \quad (5.22)$$

The periodic axes in the 2D case are mathematically equivalent, but one ordering may converge faster. DynaMag may exchange x and y to avoid a small transverse radius and reduce the number of retained terms. The gradient components are restored afterward. This permutation changes only the series representation, not the Green function.

For verification, Poisson summation may instead be applied in both periodic directions.

With

$$\mathbf{k}_{mn} = 2\pi \left(\frac{m}{L_x}, \frac{n}{L_y} \right), \quad \kappa_{mn} = |\mathbf{k}_{mn}|, \quad (5.23)$$

the scalar and differentiated forms for $z \neq 0$ are

$$\mathcal{G}_2(\boldsymbol{\rho}, z) = -\frac{|z|}{2A} + \frac{1}{2A} \sum_{(m,n) \neq (0,0)} \frac{e^{-\kappa_{mn}|z|}}{\kappa_{mn}} e^{-i\mathbf{k}_{mn} \cdot \boldsymbol{\rho}}, \quad (5.24)$$

$$\mathbf{D}_2(\boldsymbol{\rho}, z) = -\frac{\text{sgn}(z)}{2A} \mathbf{e}_z + \frac{1}{2A} \sum_{(m,n) \neq (0,0)} e^{-\kappa_{mn}|z|} e^{-i\mathbf{k}_{mn} \cdot \boldsymbol{\rho}} \begin{bmatrix} -ik_{x,m}/\kappa_{mn} \\ -ik_{y,n}/\kappa_{mn} \\ -\text{sgn}(z) \end{bmatrix}. \quad (5.25)$$

This reciprocal expression is analytically equivalent to Equations (5.20)–(5.22). Because it is evaluated by a different sum, agreement between the two forms is a strong check of the derivation.

The reciprocal form also exposes the zero harmonic:

$$\mathbf{D}_2^{(0)}(z) = -\frac{\text{sgn}(z)}{2A} \mathbf{e}_z. \quad (5.26)$$

For a source mean Q , the derivative retains

$$\nabla \Phi_2^{(0)}(z) = -\frac{Q}{2A} \operatorname{sgn}(z) \mathbf{e}_z, \quad \nabla (4\pi \Phi_2^{(0)}) = -\frac{2\pi Q}{A} \operatorname{sgn}(z) \mathbf{e}_z. \quad (5.27)$$

The two one-sided limits differ by Q/A in the fundamental-solution normalization, and their average gives the symmetric principal value on the source plane. A nonzero scalar mean still requires a finite-part convention. For the partial-period action, only the divergent additive constant is discarded. The position-dependent term $-|z|Q/(2A)$ is retained, and its derivative gives the physical sheet field. Subtracting the source mean would remove this field and fail the slab jump condition.

5.2.3 3D periodicity

A lattice with 3D periodicity has no open direction. Let its cell volume be $V = L_x L_y L_z$ and define the reciprocal vectors by

$$\mathbf{k} = 2\pi \left(\frac{m}{L_x}, \frac{n}{L_y}, \frac{\ell}{L_z} \right). \quad (5.28)$$

For a neutral cell, the mean-zero scalar Green function and its gradient are

$$\mathcal{G}_3(\mathbf{r}) = \frac{1}{V} \sum_{\mathbf{k} \neq \mathbf{0}} \frac{e^{-i\mathbf{k} \cdot \mathbf{r}}}{|\mathbf{k}|^2}, \quad (5.29)$$

$$\mathbf{D}_3(\mathbf{r}) = -\frac{i}{V} \sum_{\mathbf{k} \neq \mathbf{0}} \frac{\mathbf{k}}{|\mathbf{k}|^2} e^{-i\mathbf{k} \cdot \mathbf{r}}. \quad (5.30)$$

Pairing each reciprocal vector with its negative makes the reality and oddness of the gradient explicit:

$$\mathbf{D}_3(\mathbf{r}) = -\frac{2}{V} \sum_{\mathbf{k} \in \mathcal{H}} \frac{\mathbf{k}}{|\mathbf{k}|^2} \sin(\mathbf{k} \cdot \mathbf{r}), \quad (5.31)$$

where $\mathcal{H}$ contains one member of each reciprocal pair.

The reciprocal sum is convenient for analysis, while a mixed Fourier–Bessel representation is more efficient near the primary cell. After wrapping the displacement and choosing x as the

Fourier-expanded coordinate, define

$$a_m = \frac{2\pi m}{L_x}, \quad \theta = \frac{2\pi y}{L_y}, \quad q_k = e^{-2\pi|z+kL_z|/L_y}, \quad \Delta_k = 1 - 2q_k \cos \theta + q_k^2, \quad (5.32)$$

and

$$R_{nk} = \sqrt{(y + nL_y)^2 + (z + kL_z)^2}. \quad (5.33)$$

The corresponding scalar formula is

$$\begin{aligned} \mathcal{G}_3(\mathbf{r}) = & \frac{z^2 - |z|L_z + L_z^2/6}{2V} - \frac{1}{4\pi L_x} \sum_{k=-\infty}^{\infty} \log \Delta_k \\ & + \frac{1}{\pi L_x} \sum_{m=1}^{\infty} \sum_{n,k=-\infty}^{\infty} K_0(a_m R_{nk}) \cos(a_m x). \end{aligned} \quad (5.34)$$

The quadratic term is the periodic Bernoulli finite part of the repeated sheet mode for 2D periodicity.

Differentiating the complete expression gives

$$D_{3,x} = -\frac{1}{\pi L_x} \sum_{m=1}^{\infty} \sum_{n,k=-\infty}^{\infty} a_m K_0(a_m R_{nk}) \sin(a_m x), \quad (5.35)$$

$$D_{3,y} = -\frac{1}{L_x L_y} \sum_{k=-\infty}^{\infty} \frac{q_k \sin \theta}{\Delta_k} - \frac{1}{\pi L_x} \sum_{m=1}^{\infty} \sum_{n,k=-\infty}^{\infty} a_m \frac{y + nL_y}{R_{nk}} K_1(a_m R_{nk}) \cos(a_m x), \quad (5.36)$$

$$\begin{aligned} D_{3,z} = & \frac{2z - L_z \operatorname{sgn}(z)}{2V} - \frac{1}{L_x L_y} \sum_{k=-\infty}^{\infty} \frac{\operatorname{sgn}(z + kL_z) q_k (\cos \theta - q_k)}{\Delta_k} \\ & - \frac{1}{\pi L_x} \sum_{m=1}^{\infty} \sum_{n,k=-\infty}^{\infty} a_m \frac{z + kL_z}{R_{nk}} K_1(a_m R_{nk}) \cos(a_m x). \end{aligned} \quad (5.37)$$

As in the 2D-periodicity case, the axes may be permuted to obtain a rapidly convergent ordering.

The original component order and signs are restored after evaluation.

At the lattice origin, the singular zero image is omitted. The remaining finite periodic self

value is evaluated by an Ewald formula. For a splitting parameter $\beta > 0$,

$$\begin{aligned} \mathcal{G}_{3,\text{self}} = & -\frac{1}{4\pi\sqrt{\pi}\beta} - \frac{\beta}{V} + \sum_{\mathbf{R} \neq \mathbf{0}} \frac{\text{erfc}(|\mathbf{R}|/(2\sqrt{\beta}))}{4\pi|\mathbf{R}|} \\ & + \frac{1}{V} \sum_{\mathbf{k} \neq \mathbf{0}} \frac{e^{-\beta|\mathbf{k}|^2}}{|\mathbf{k}|^2}, \end{aligned} \quad (5.38)$$

Inversion symmetry then gives $\mathbf{D}_3(\mathbf{0}) = \mathbf{0}$. The parameter β changes the rate at which the two Ewald sums converge, but not their limiting value.

The missing open direction changes the zero-mode condition. A nonzero mean cannot discharge its flux through a transverse coordinate. For a neutral cell, omission of $\mathbf{k} = \mathbf{0}$ is automatic. For a nonneutral cell, deleting that coefficient introduces a uniform compensating background. Such a background is an additional physical model rather than a consequence of the Green-function derivation. The action for 3D periodicity therefore requires neutrality.

5.3 Numerical realization and singular interactions

5.3.1 Production truncation and symmetry branches

The analytical formulas still contain infinite sums. Their numerical evaluation uses the exponential decay of K_0 , K_1 , and the logarithmic image factors to select finite cutoffs. With tolerance parameter ϵ and

$$\epsilon = \frac{\epsilon_{\text{pi}}}{L_{\min}}, \quad (5.39)$$

the Fourier cutoff for 1D periodicity is

$$M_{\max}^{(1)} = \left\lceil \frac{L \log(1/\epsilon)}{2\pi\rho} \right\rceil. \quad (5.40)$$

For the 2D-periodicity representation after any x - y exchange,

$$M_{\max}^{(2)} = \left\lceil \frac{L_x \log(1/\epsilon)}{2\pi \sqrt{y^2 + z^2}} \right\rceil, \quad N_{\max}^{(2)} = \left\lceil \frac{L_x \log(1/\epsilon)}{2\pi L_y} \right\rceil. \quad (5.41)$$

For 3D periodicity, after the coordinate permutation used in Equation (5.34),

$$\begin{aligned} M_{\max}^{(3)} &= \left\lceil \frac{L_x \log(1/\epsilon)}{2\pi \sqrt{y^2 + z^2}} \right\rceil, \\ N_{\max}^{(3)} &= \left\lceil \frac{L_x \log(1/\epsilon)}{2\pi L_y} \right\rceil, \quad K_{\max}^{(3)} = \left\lceil \frac{L_x \log(1/\epsilon)}{2\pi L_z} \right\rceil, \\ K_{\max}^{(\log)} &= \left\lceil \frac{L_y \log(1/\epsilon)}{2\pi L_z} \right\rceil. \end{aligned} \quad (5.42)$$

The last cutoff belongs to the logarithmic image sum, while $K_{\max}^{(3)}$ belongs to the three-index Bessel sum. The scalar and gradient evaluations use the same cutoffs so that their discrete truncations remain consistent. Coordinates are signed-wrapped to the nearest image. Evenness is used for the scalar, and the signs and component order are restored for the gradient. At a source image, odd symmetry gives zero point gradient in the principal-value sense. Separate limiting formulas handle the axis for 1D periodicity, the source plane for 2D periodicity, and the scalar self value without evaluating $K_0(0)$ or $K_1(0)$ directly.

5.3.2 Neutrality: scalar charge, closed current, and periodic transport

The zero-mode rule can now be stated precisely. For point strengths q_j , an ordinary periodic-Poisson inverse requires $Q = \sum_j q_j = 0$ and returns its potential only up to a constant. The explicit finite-part Green function for 1D or 2D periodicity removes the divergent constant but retains the finite line or sheet mode. Its derivative reaches the same field directly. Neither operation subtracts the source mean. Calculations with 3D periodicity still require neutrality unless a compensating background is part of the physical model. Table 5.1 summarizes the distinction.

For a tangential surface current $\mathbf{k}$ on a closed surface Γ , the relevant vector-source moment

Table 5.1. Zero-mode and neutrality distinction between scalar and gradient periodic kernels.

Periodicity	Scalar $1/r$ potential	Cartesian gradient
1D periodicity	Ordinary inverse requires neutrality; the explicit finite part of Equation (5.7) admits a nonzero source mean	Nonzero mean allowed; finite line field, Equation (5.12)
2D periodicity	Ordinary inverse requires neutrality; the finite-part action retains the position-dependent $- z Q/(2A)$ sheet mode	Nonzero mean allowed; finite sheet field and a symmetric source-plane principal value, Equation (5.27)
3D periodicity	Neutrality or an explicit uniform background	Neutrality required; deleting a nonzero reciprocal coefficient would itself impose a background convention

is

$$\mathbf{Q}_\Gamma = \int_\Gamma \mathbf{k} dS. \quad (5.43)$$

Let $\varphi(\mathbf{x}) = \mathbf{a} \cdot \mathbf{x}$ for arbitrary constant $\mathbf{a}$. Since $\nabla_\Gamma \varphi \cdot \mathbf{k} = \mathbf{a} \cdot \mathbf{k}$, surface integration by parts yields

$$\mathbf{a} \cdot \mathbf{Q}_\Gamma = \int_\Gamma \nabla_\Gamma \varphi \cdot \mathbf{k} dS = - \int_\Gamma \varphi \nabla_\Gamma \cdot \mathbf{k} dS = 0. \quad (5.44)$$

Because $\mathbf{a}$ is arbitrary, a divergence-free current on a closed surface is neutral in every Cartesian component. A periodic surface represented by a cut cell has an additional boundary term. Fluxes on opposite cuts cancel locally, but their coordinate moments differ by a lattice vector. This remainder represents transport through the quotient cell. The loop basis must therefore contain both local contractible loops and cycles that wind around a periodic direction. Subtracting the mean from a winding cycle would change the imposed transport current.

At the implementation level, identifying vertex and edge equivalence classes across periodic cuts is a disjoint-set problem. Union–find algorithms provide an efficient way to assemble those classes (Tarjan, 1975). This operation constructs the quotient topology; it does not determine the Green-function zero-mode policy.

The retained mean of a winding current contributes directly to the H – k magnetic trace. In

the public $G = 1/r$ normalization, its zero-mode field is

$$\mathbf{H}_k^{(0)}(\mathbf{x}) = \frac{1}{4\pi} \nabla G_p^{(0)}(\mathbf{x}) \times \mathbf{Q}_\Gamma = \frac{1}{4\pi} \mathbf{Q}_\Gamma \times \mathbf{K}_p^{(0)}(\mathbf{x}), \quad p \in \{1, 2\}. \quad (5.45)$$

Equation (5.45) is lost if the source is projected to zero mean. The native periodic $\mathbf{r}/r^3$ action retains it.

5.3.3 Self value, panel limit, and near replacement

The gradient kernel is singular at a source point, but its local vector singularity is odd. A same-cloud point action uses the symmetric principal-value convention

$$\mathbf{D}_p(\mathbf{0}) = \mathbf{0}. \quad (5.46)$$

Nonzero periodic images cancel in inversion pairs at this displacement. This point convention does not define the integral over a boundary panel. For a source panel T ,

$$\nabla_{\mathbf{x}} \int_T \frac{dS_y}{4\pi|\mathbf{x} - \mathbf{y}|} = - \int_T \frac{\mathbf{x} - \mathbf{y}}{4\pi|\mathbf{x} - \mathbf{y}|^3} dS_y, \quad (5.47)$$

and the boundary limit contains both a principal-value integral and a one-sided jump. B-AIM restores this limit by subtracting the local grid approximation and adding the analytical or accurately integrated zero-image panel action. The smooth nonzero images remain in the periodic Green function and are not added again by the local correction.

5.4 Validation hierarchy and error separation

Verification begins with the Green formulas before any accelerated operator is introduced. Table 5.2 compares independent representations of the same scalar derivatives.

The first three rows compare the analytic gradient with a central difference of the truncated

Table 5.2. Independent audits of the periodic-gradient formulas.

Comparison	Maximum absolute discrepancy
Central difference of truncated scalar formula, 1D periodicity	1.1×10^{-11}
Central difference of truncated scalar formula, 2D periodicity	2.4×10^{-11}
Central difference of truncated scalar formula, 3D periodicity	2.7×10^{-11}
1D-periodicity axial trigamma formula versus 200,000 images per side	1.9×10^{-13}
2D-periodicity reciprocal formula versus real K_0/K_1 formula	3.4×10^{-16}
3D-periodicity Abel-damped reciprocal versus real K_0/K_1 formula	2.5×10^{-9}

scalar formula, while the other rows use axial-image, reciprocal, and Abel/Ewald representations that do not share the same evaluation path. Together, these comparisons test differentiation, sign, normalization, and truncation.

The next level includes projection, FFT convolution, interpolation, and near replacement. Complete B-AIM actions are compared with independent direct periodic calculations. The small-cloud relative L^2 errors are 7.897×10^{-3} , 9.397×10^{-3} , and 1.691×10^{-2} for 1D, 2D, and 3D periodicity. The first two clouds retain a nonzero source mean, whereas the fully periodic cloud is neutral. Continuum tests then add surface geometry and quadrature. For the current on the cylinder with 1D periodicity, the accelerated and direct-PGF errors are 2.340×10^{-3} and 1.965×10^{-3} . A nonneutral sheet with 2D periodicity reproduces a 7.5×10^5 A/m field jump with only 2.386×10^{-3} A/m absolute error.

Table 5.3. Independent numerical evidence for the periodic-gradient kernel.

Case	Comparison	Result
Point cloud with 1D periodicity	Complete B-AIM versus image reference	Relative L^2 error 7.897×10^{-3}
Point cloud with 2D periodicity	Complete B-AIM versus reciprocal reference	Relative L^2 error 9.397×10^{-3}
Point cloud with 3D periodicity	Complete B-AIM versus Ewald reference	Relative L^2 error 1.691×10^{-2}
Cylinder with 1D periodicity	Accelerated/direct PGF versus analytical field	$2.340 \times 10^{-3} / 1.965 \times 10^{-3}$ relative L^2
Sheet with 2D periodicity	Computed versus analytical field jump	2.386×10^{-3} A/m absolute error on 7.5×10^5 A/m




Figure 5.1. Complementary validation of the native periodic-gradient kernel. The cylinder includes a direct all-to-all periodic Green-function oracle; the sheet tests the finite nonzero derivative mode that would be removed by an inappropriate zero-mean projection.

If $\mathbf{H}_h^{\text{acc}}$, $\mathbf{H}_h^{\text{dir}}$, and $\mathbf{H}^{\text{ana}}$ denote accelerated, direct-discrete, and continuum fields, then

$$\mathbf{H}_h^{\text{acc}} - \mathbf{H}^{\text{ana}} = (\mathbf{H}_h^{\text{acc}} - \mathbf{H}_h^{\text{dir}}) + (\mathbf{H}_h^{\text{dir}} - \mathbf{H}^{\text{ana}}). \quad (5.48)$$

The accelerated and direct discrete solutions share the same continuum discretization error, so their separate errors against the analytical field do not directly give the acceleration error. For the cylinder, the triangle and reverse-triangle inequalities give the rigorous normalized interval

$$3.751 \times 10^{-4} \leq \frac{\|\mathbf{H}_h^{\text{acc}} - \mathbf{H}_h^{\text{dir}}\|_2}{\|\mathbf{H}^{\text{ana}}\|_2} \leq 4.305 \times 10^{-3}. \quad (5.49)$$

This hierarchy separates errors due to the analytic formula, PGF truncation, B-AIM acceleration, panel quadrature, and continuum discretization. It also tests the most important partial-period feature: the boundary operator retains the finite zero modes of the analytic $\mathbf{r}/r^3$ gradient.

With the formula and acceleration errors separated, the remaining question is performance. The following section measures the complete B-AIM action, including projection, convolution, interpolation, and near replacement. Chapter 7 later shows how repeated applications of this operator enter the FEM–BEM blocks.

5.5 B-AIM performance and storage scaling

The preceding tests establish the accuracy of the periodic Green functions. Their use in an implicit FEM–BEM solver also requires a fast repeated operator action. The useful measure is therefore the cost of the complete B-AIM evaluation, not the cost of the FFT alone. This section reports that cost and keeps plan preparation separate from repeated evaluation.

5.5.1 Benchmark definition and protocol

For point sources q_j and targets $\mathbf{x}_i$, the scalar B-AIM action is written as

$$\boldsymbol{\phi} \simeq \mathcal{I} \mathcal{F}^{-1} \widehat{G} \mathcal{F} \mathcal{P} \mathbf{q} + \left(\mathcal{D}_{\text{near}} - \widetilde{\mathcal{D}}_{\text{near}} \right) \mathbf{q}. \quad (5.50)$$

Here $\mathcal{P}$ reorders the sources and projects them to a Cartesian grid. The FFT applies the grid convolution, and $\mathcal{I}$ interpolates the result to the targets. The last term removes the local grid approximation and restores the physical kernel.

The periodic vector action retains its dedicated analytical gradient Green function. Its discrete form is

$$\mathbf{h} \simeq \mathcal{I} \mathcal{F}^{-1} \widehat{\nabla G_{\text{per}}} \mathcal{F} \mathcal{P} \mathbf{q} + \left(\mathcal{D}_{\text{near}}^{\nabla} - \widetilde{\mathcal{D}}_{\text{near}}^{\nabla} \right) \mathbf{q}. \quad (5.51)$$

The three components of $\widehat{\nabla G_{\text{per}}}$ are prepared from the periodic $\mathbf{r}/r^3$ kernel derived earlier in this chapter. The finite line or sheet zero mode is included analytically. This operation is not replaced by differentiating an interpolated scalar field. The two procedures are equivalent in the continuous limit, but they are different after grid projection and interpolation. Direct use of the gradient Green function preserves the validated kernel and its zero-mode treatment.

The optimized box form factors are used in both Equations (5.50) and (5.51). They reduce projection, interpolation, and local-replacement cost without changing the Green function. The periodic-gradient rate is measured with the dedicated gradient spectra. A timing obtained by differentiating the scalar interpolant would measure a different operator and is not used here.

The optimized scalar path uses trilinear projection and interpolation with about eight points per physical box. It replaces the grid interaction in a ± 2 -box cube with exact $1/r$ interactions. Particles are ordered by box, so neighboring boxes along z can be traversed as contiguous source ranges. No source–target pair list is stored. The free-space FFT dimensions are padded to nearby smooth integers, which avoids slow transforms with large prime factors.

The timed scalar action includes input mapping, projection, FFT convolution, interpolation, local replacement, output gathering, and synchronization. Point generation, plan construction, workspace allocation, and host output are excluded. Setup is measured separately. The throughput is

$$R = \frac{N_s}{T_{\text{evaluate}}}, \quad (5.52)$$

where N_s is the number of source vertices.

The tests use single precision on an NVIDIA GeForce RTX 5090 with 32 GB of memory. The driver version is 595.71.05, and the code was compiled with CUDA 12.9. The source and target sets are the same deterministic cloud in $[-10, 10]^3$. Five independent processes were measured at 5 and 20 million vertices. Each process used 8 warm-up actions and 40 timed actions.

5.5.2 Measured performance

Table 5.4 gives the replicated scalar results. The complete prepared action reaches 1379.7 ± 3.1 million vertices per second at 5 million vertices. It remains above 1.3 billion vertices per second at 20 million vertices. Neither rate includes setup time.

Table 5.4. Prepared performance of the optimized scalar B-AIM path. Values after the plus–minus sign are sample standard deviations across five independent processes.

Operator	Vertices	Setup (ms)	Evaluation (ms)	Mvertices/s
Optimized free-space $1/r$ scalar	5 million	31.384 ± 0.673	3.624 ± 0.008	1379.72 ± 3.14
Optimized free-space $1/r$ scalar	20 million	39.192 ± 0.395	15.161 ± 0.041	1319.15 ± 3.55

Against a double-precision direct sum on 103,680 points, this linear path has a normalized

root-mean-square potential error of 7.90×10^{-3} . A quadratic interpolation path lowers the error to 5.05×10^{-4} , but its 20-million-point rate is 303.0 million vertices per second. The interpolation order and local replacement radius can therefore be selected according to the required error.

The dedicated vector path uses the same local replacement idea. Its sources are first ordered by box and stored as (x, y, z, q) records. Since the box indices are z -major, the five neighboring boxes along z form one contiguous source range for each (x, y) box pair. This avoids a separate offset lookup for every neighboring box. For each source–target pair, the kernel evaluates

$$s = \text{rsqrt}(r^2), \quad c = -q_j s^3, \quad \mathbf{h}_i += c(\mathbf{x}_i - \mathbf{x}_j). \quad (5.53)$$

Thus, one reciprocal square root supplies all three vector components. The interpolation subtraction still uses the dedicated periodic-gradient Green table. Only the traversal and source layout have changed.

Table 5.5 reports the complete prepared vector action. These runs use 20 million vertices, four requested points per box, and a 171^3 physical-box grid. Five independent processes used 5 warm-up actions and 20 timed actions. The periodic axes are x for 1D periodicity and xy for 2D periodicity. The result is about 650 million vertices per second in all three cases.

Table 5.5. Prepared performance of the dedicated analytical $\mathbf{r}/r^3$ B-AIM action at 20 million vertices. Values after the plus–minus sign are sample standard deviations across five independent processes.

Boundary condition	Setup (ms)	Evaluation (ms)	Mvertices/s
Free space	36.773 ± 0.828	30.548 ± 0.023	654.71 ± 0.50
1D periodicity	37.808 ± 1.156	30.667 ± 0.008	652.18 ± 0.16
2D periodicity	41.259 ± 0.184	30.685 ± 0.009	651.79 ± 0.20

The previous dedicated-gradient benchmark used a group-based correction and eight requested points per box. It took 61.320 ms for 20 million vertices, or 326.16 million vertices per second. The new traversal, together with the four-point box setting, increases the complete vector rate by a factor of 2.01. Adding periodicity changes the measured throughput by less than 0.5%. A

sampled double-precision comparison with direct summation gave a relative L^2 gradient error of 6.98×10^{-4} at this box density.

Table 5.6 shows the scalar scaling. The rate approaches 1.4 billion vertices per second once the GPU is filled. It remains above 1.1 billion vertices per second through 300 million vertices. The latter case finishes one complete prepared action in 262.158 ms.

Table 5.6. Scaling of the complete optimized free-space $1/r$ action. Setup time is excluded.

Sources	Physical boxes	FFT grid	Throughput (Mvertices/s)
100,000	23^3	48^3	573.86
1,000,000	50^3	105^3	1312.69
2,000,000	63^3	128^3	1383.00
5,000,000	85^3	175^3	1379.72
10,000,000	108^3	224^3	1352.24
20,000,000	136^3	280^3	1319.15
50,000,000	184^3	375^3	1104.68
100,000,000	232^3	480^3	1162.81
300,000,000	335^3	672^3	1144.35

The smooth FFT dimensions matter at the largest sizes. For example, padding the 300-million-point transform from 671^3 to 672^3 reduces the convolution time from 38.709 to 13.334 ms. The logical grid and local physical correction do not change, so the direct-sum error is unchanged within single-precision roundoff.

Device storage. The storage audit uses the capacities of the actual arrays in the optimized 20-million-point scalar path. It excludes the CUDA context and library-owned allocations. Table 5.7 separates solver storage from the caller-owned coordinates, charges, and result.

No dense BEM matrix or point-pair correction matrix is stored. The local action traverses the reordered box ranges and evaluates the required pairs on the fly. The stored data therefore scale as $O(N_s + N_t + N_g)$ rather than $O(N_s N_t)$. At 300 million points, the solver internal and work storage is 17.222 GiB, and the calculation completes on the 32-GB device.

The optimized scalar action reaches nearly 1.4 billion vertices per second. The dedicated periodic-gradient action reaches about 650 million vertices per second while preserving its

Table 5.7. Device-array storage for the optimized 20-million-point scalar B-AIM calculation.

Storage group	Device storage (MiB)
Reordered particle data, particle work arrays, and maps	915.527
Box offsets	9.596
Grid values, kernel spectrum, FFT workspace, and local Green data	265.067
Solver internal and work storage	1190.190
Caller-owned coordinates, charges, and scalar result	381.470
Complete problem-array total	1571.660

analytical kernel and zero modes. The finite-element assembly, preconditioner, and Krylov convergence are measured separately in Chapter 7.

Publication disclosure

This chapter is based in part on “Periodic Micromagnetic Finite Element Method” by Fangzhou Ai, Jiawei Duan, and Vitaliy Lomakin, published in the *Journal of Magnetism and Magnetic Materials*, vol. 615, 2025, article 172776, doi:10.1016/j.jmmm.2025.172776. The dissertation author was a contributing author of this paper.

6 Airbox-Free Eddy-Current FEM–BEM and Coupled LLG–Maxwell Dynamics

Time-dependent magnetization changes the magnetic flux in a conductor, which induces a current and a reaction field. That field enters the LLG effective field, so the LLG and Maxwell states must share a time history. The conducting-disk calculation in this chapter shows how this feedback changes the switching trajectory.

To preserve the dynamic-mesh design, the insulating exterior must be represented without a moving volume mesh. DynaMag therefore meshes only the conductor and represents the unbounded exterior on its physical surface. Thus, an air mesh does not become part of the evolving geometry. Chapter 5 develops the periodic Green functions and B-AIM actions used for this representation. This chapter presents the FEM–BEM formulations and their physical validation. Repeated application, memory, and preconditioning are discussed in Chapter 7.

The variational background combines several established strands. Dual eddy-current formulations and edge-element computational electromagnetism provide conductor models that respect curl conformity (Bossavit, 1985, 1998). Johnson–Nédélec coupling and later adaptive analysis show how finite and boundary elements can share an interface without meshing the exterior (Johnson and Nédélec, 1980; Aurada et al., 2012). The required Maxwell trace spaces, integration-by-parts formula, and boundary Hodge decomposition are described in (Buffa and Ciarlet, 2001a,b).

For coupled magnetic dynamics, existence and regularity results for the Maxwell–

Landau–Lifshitz system precede convergent implicit and quasi-static finite-element schemes (Cimrák, 2007; Bañas et al., 2008; Le and Tran, 2013). Decoupled eddy-current–LLG integration and FEM–BEM analysis provide two closely related numerical reference points (Le et al., 2015; Feischl and Tran, 2017). Long-time micromagnetic analysis supplies additional context for the asymptotic state (Carbou and Fabrie, 1998). The contribution here is the device-resident, matrix-free realization and its free-space and partial-period extensions.

6.1 Airbox-free eddy-current formulation

At the device scales considered here, the displacement current is neglected but Faraday induction is retained. The magnetoquasistatic system is

$$\nabla \times \mathbf{H} = \mathbf{J}, \quad \nabla \cdot \mathbf{B} = 0, \quad (6.1)$$

$$\nabla \times \mathbf{E} = -\partial_t \mathbf{B}, \quad \mathbf{J} = \sigma \mathbf{E}, \quad (6.2)$$

$$\mathbf{B} = \mu \mathbf{H}. \quad (6.3)$$

Here conductivity is nonzero only in the conductor Ω_c . In the insulating exterior $\Omega^+ = \mathbb{R}^3 \setminus \overline{\Omega_c}$, $\nabla \times \mathbf{H}^+ = \mathbf{0}$ and $\nabla \cdot \mathbf{B}^+ = 0$, but $\mathbf{E}^+$ is generally nonzero because $\nabla \times \mathbf{E}^+ = -\partial_t \mathbf{B}^+$. The exterior is therefore a real part of the Maxwell problem even though it contains no conduction current.

Three numerical strategies are considered.

Volumetric FEM with an airbox. A conventional finite-element treatment surrounds the conductor with an air domain and truncates it at a finite boundary. The system is sparse, but it contains many nonconducting unknowns. It also replaces the condition at infinity with one imposed at a chosen distance. A nearby boundary distorts the field and its normal decay. Moving that boundary outward reduces the error but adds more air elements. The mesh must remain fine near the conductor while grading toward the distant boundary. Rigid motion or morphing also becomes an air-mesh deformation or remeshing problem.

Pure integral formulation. An integral representation enforces the open boundary at the continuum level and removes the air volume. A direct discretization couples every source point to every target point. The resulting $N_\Gamma \times N_\Gamma$ operator has quadratic storage and application cost. An implicit step applies this operator many times inside GMRES, so the all-to-all cost is repeated at every accepted time step.

Accelerated FEM–BEM. FEM–BEM keeps the sparse finite-element diffusion problem in Ω_c and uses Green’s representation to eliminate Ω^+ (Johnson and Nédélec, 1980; Feischl and Tran, 2017). It adds neither air-volume unknowns nor a finite exterior boundary. The electromagnetic operator uses the physical boundary of the LLG mesh. Prepared B-AIM and local panel replacement accelerate this dense action. The persistent state contains sparse FEM data, surface topology, Green spectra, and reusable workspaces. It contains no dense BEM matrix.

A– k gives the free-space FEM–BEM construction and its analytical hollow-sphere validation. The intermediate H– k form then exposes the vector single layer and its magnetic trace as separate boundary actions. The final H– κ form uses $\mathbf{H}$ in the conductor and a scalar single-layer coordinate on the boundary. It provides both the implemented LLG coupling and the verified transient path for 1D and 2D periodicity.

6.2 Classical A– k FEM–BEM formulation

The vector-potential coordinate extends conventional eddy-current FEM to an open-boundary solver. As the first formulation validated in DynaMag, it also provides a baseline for the field-based formulations.

6.2.1 Maxwell reduction and the $\mathbf{A}$ – ϕ equation

The derivation starts from the magnetoquasistatic Maxwell system in the conductor:

$$\begin{aligned}\nabla \times \mathbf{E} &= -\partial_t \mathbf{B}, & \nabla \cdot \mathbf{B} &= 0, \\ \nabla \times \mathbf{H} &= \mathbf{J} + \mathbf{J}_{\text{app}}, & \nabla \cdot \mathbf{J} &= 0, \\ \mathbf{B} &= \mu \mathbf{H}, & \mathbf{J} &= \sigma \mathbf{E}.\end{aligned}\tag{6.4}$$

Introduce a magnetic vector potential and an electric scalar potential:

$$\mathbf{B} = \nabla \times \mathbf{A}, \quad \mathbf{E} = -\partial_t \mathbf{A} - \nabla \phi, \quad \mathbf{J} = -\sigma (\partial_t \mathbf{A} + \nabla \phi).\tag{6.5}$$

The first definition satisfies $\nabla \cdot \mathbf{B} = 0$. The second satisfies Faraday's law because $\nabla \times \nabla \phi = \mathbf{0}$. Substitution into Ampère's law gives one equation for the two conductor potentials:

$$\sigma \partial_t \mathbf{A} + \sigma \nabla \phi + \nabla \times (\nu \nabla \times \mathbf{A}) = \mathbf{J}_{\text{app}}.\tag{6.6}$$

Charge conservation supplies the second conductor equation,

$$\nabla \cdot [\sigma (\partial_t \mathbf{A} + \nabla \phi)] = 0.\tag{6.7}$$

For a closed insulated conductor, $\mathbf{n} \cdot \mathbf{J} = 0$ on Γ . Hence ϕ solves

$$\nabla \cdot (\sigma \nabla \phi) = -\nabla \cdot (\sigma \partial_t \mathbf{A}),\tag{6.8}$$

$$\mathbf{n} \cdot \sigma \nabla \phi = -\mathbf{n} \cdot \sigma \partial_t \mathbf{A} \quad \text{on } \Gamma.\tag{6.9}$$

The $\mathbf{A}$ – k reduction is valid only if this Neumann problem has a constant solution. That property holds for the symmetric closed-current validations used here. It is not a general consequence of using a loop basis. Contacts, terminal voltages, conductivity interfaces, and general moving

conductors require the full $A\text{--}\phi\text{--}k$ form unless the same condition is proved.

6.2.2 Interface equations

Let $\mathbf{n}$ point from the conductor into the air. With no free surface current sheet and with an insulated conductor, the interface equations are

$$\begin{aligned}\mathbf{n} \times (\mathbf{H}^+ - \mathbf{H}^-) &= \mathbf{0}, & \mathbf{n} \cdot (\mathbf{B}^+ - \mathbf{B}^-) &= 0, \\ \mathbf{n} \times (\mathbf{E}^+ - \mathbf{E}^-) &= \mathbf{0}, & \mathbf{n} \cdot \mathbf{J}^- &= 0.\end{aligned}\tag{6.10}$$

The superscripts denote the air and conductor traces. These equations remain present after the air mesh is removed. BEM supplies the exterior traces that appear in them.

The vector potential also needs a gauge. DynaMag uses a compatible gauge in which the tangential trace of $\mathbf{A}$ is matched to the vector single-layer representation. This gauge trace and tangential-H continuity give the two equations later represented by the P/D and Q blocks.

6.2.3 Weak form and the two Green identities

Take $\mathbf{w} \in H(\text{curl}, \Omega_c)$. The curl Green identity is

$$\int_{\Omega_c} (\nabla \times \mathbf{u}) \cdot \mathbf{w} dV = \int_{\Omega_c} \mathbf{u} \cdot (\nabla \times \mathbf{w}) dV + \int_{\Gamma} (\mathbf{n} \times \mathbf{u}) \cdot \mathbf{w} dS.\tag{6.11}$$

Set $\mathbf{u} = \mathbf{H}^- = \nu \nabla \times \mathbf{A}$ and test Equation (6.6). This gives

$$\begin{aligned}(\sigma \partial_t \mathbf{A}, \mathbf{w})_{\Omega_c} + (\sigma \nabla \phi, \mathbf{w})_{\Omega_c} + (\nu \nabla \times \mathbf{A}, \nabla \times \mathbf{w})_{\Omega_c} \\ = (\mathbf{J}_{\text{app}}, \mathbf{w})_{\Omega_c} - \langle \mathbf{n} \times \mathbf{H}^-, \mathbf{w} \rangle_{\Gamma}.\end{aligned}\tag{6.12}$$

Tangential-H continuity changes the last trace from $\mathbf{H}^-$ to the exterior field $\mathbf{H}^+$. Define

$$f_{\text{bem}}(\mathbf{w}) := -\langle \mathbf{n} \times \mathbf{H}^+, \mathbf{w} \rangle_{\Gamma}.\tag{6.13}$$

This is the only change made by eliminating the air volume from the A row.

The charge equation uses the scalar divergence identity. For $q \in H^1(\Omega_c)$,

$$0 = \int_{\Omega_c} q \nabla \cdot \mathbf{J} dV = -(\mathbf{J}, \nabla q)_{\Omega_c} + \langle \mathbf{n} \cdot \mathbf{J}, q \rangle_{\Gamma}. \quad (6.14)$$

The boundary term vanishes by insulation. Using Equation (6.5) leaves

$$(\sigma(\partial_t \mathbf{A} + \nabla \phi), \nabla q)_{\Omega_c} = 0. \quad (6.15)$$

Equations (6.12) and (6.15) are the complete conductor weak problem before the A- k reduction.

6.2.4 BEM closure and trace maps

Let $\mathbf{r}$ be a tangential, surface-divergence-free density. The exterior vector single layer is

$$\mathbf{A}^+(\mathbf{x}) = \frac{\mu_0}{4\pi} \int_{\Gamma} \frac{\mathbf{r}(\mathbf{y})}{|\mathbf{x} - \mathbf{y}|} dS_y, \quad \mathbf{H}^+ = \mu_0^{-1} \nabla \times \mathbf{A}^+. \quad (6.16)$$

It is harmonic away from Γ and has the required decay. Its exterior magnetic trace contains a Cauchy-principal-value integral and a one-sided half jump. Thus Green's representation supplies both the potential trace and the magnetic trace without a finite outer boundary.

The raw boundary density is parameterized by loop coordinates,

$$\mathbf{r} = L\mathbf{k}, \quad B_{\Gamma}L = 0. \quad (6.17)$$

Here B_{Γ} is the surface incidence operator. L makes the boundary density divergence free before it enters the BEM operator.

Three trace maps connect the boundary representation to FEM. $P\mathbf{r}$ is the weak tangential trace of the $1/r$ single layer. $D\mathbf{a}$ is the same trace taken from the FEM vector potential. Their

gauge-compatible matching equation is

$$-D\mathbf{a} + P\mathbf{r} = \mathbf{0}. \quad (6.18)$$

$Q\mathbf{r}$ is the one-sided exterior magnetic trace tested in the FEM edge rows. It evaluates the derivative kernel $\mathbf{r}/r^3$ and adds the half jump. It is the discrete form of f_{bem} in Equation (6.13). The P/D equation fixes the potential trace. The Q action enforces the magnetic interface condition in the volume row.

Neither P nor Q is stored as a dense matrix. B-AIM spreads the boundary sources to a regular grid, applies the Green convolution by FFT, and interpolates back to the boundary or FEM targets. Exact or graded local panel replacement restores the singular near field. The same operation is used in every Krylov iteration. Chapter 7 gives its storage and timing details.

6.2.5 From the two weak equations to four blocks

With edge basis functions $\{\mathbf{N}_i\}$, define

$$(M_\sigma)_{ij} = \int_{\Omega_c} \sigma \mathbf{N}_j \cdot \mathbf{N}_i dV, \quad (6.19)$$

$$(C_\nu)_{ij} = \int_{\Omega_c} \nu (\nabla \times \mathbf{N}_j) \cdot (\nabla \times \mathbf{N}_i) dV. \quad (6.20)$$

Before reduction, the two conductor equations have the algebraic form

$$\begin{aligned} M_\sigma \dot{\mathbf{a}} + G_\sigma \mathbf{p} + C_\nu \mathbf{a} &= \mathbf{f}_J + \mathbf{f}_{\text{bem}}, \\ G_\sigma^T \dot{\mathbf{a}} + K_\phi \mathbf{p} &= \mathbf{0}. \end{aligned} \quad (6.21)$$

G_σ comes from $(\sigma \nabla \phi, \mathbf{w})$, and K_ϕ comes from $(\sigma \nabla \phi, \nabla q)$. When Equations (6.8)–(6.9) admit a constant ϕ , the $\mathbf{p}$ column vanishes. Combining the remaining A row with Equation (6.18),

substituting $\mathbf{r} = L\mathbf{k}$, and testing the boundary row by $L^\top$ gives

$$M_\sigma \dot{\mathbf{a}} + C_\nu \mathbf{a} - QL\mathbf{k} = \mathbf{g}_A(t), \quad (6.22)$$

$$-L^\top D\mathbf{a} + L^\top PL\mathbf{k} = \mathbf{0}. \quad (6.23)$$

For an applied magnetic vector-potential load, $\mathbf{g}_A = -\partial_t \mathbf{f}_{\text{app}}$. The three coupling blocks are

$$A_{Ak} = QL, \quad A_{kA} = L^\top D, \quad A_{kk} = L^\top PL. \quad (6.24)$$

The four rows and columns now have a direct origin. M_σ and C_ν come from Ohm's and Ampère's laws. A_{Ak} returns the exterior magnetic trace to the A row. A_{kA} and A_{kk} are the two sides of the boundary trace equation.

6.2.6 BDF time discretization

For a BDF formula of order p , write

$$D_t \mathbf{a}^n = \frac{1}{\Delta t_n} \sum_{j=0}^p \alpha_j \mathbf{a}^{n-j}. \quad (6.25)$$

Only the first row of Equations (6.22)–(6.23) is differential. One BDF solve is therefore

$$\begin{bmatrix} \frac{\alpha_0}{\Delta t_n} M_\sigma + C_\nu + \epsilon G_A & -A_{Ak} \\ -A_{kA} & A_{kk} \end{bmatrix} \begin{bmatrix} \mathbf{a}^n \\ \mathbf{k}^n \end{bmatrix} = \begin{bmatrix} \mathbf{g}_A^n - \frac{1}{\Delta t_n} M_\sigma \sum_{j=1}^p \alpha_j \mathbf{a}^{n-j} \\ \mathbf{0} \end{bmatrix}. \quad (6.26)$$

If $\mathbf{g}_A = -\partial_t \mathbf{f}_{\text{app}}$, the same BDF stencil is used for that derivative. The history is formed before GMRES starts. The algebraic boundary equation has no history term.

For backward Euler, $\alpha_0 = 1$ and $\alpha_1 = -1$. Multiplying the first row by Δt gives the form

used in the free-space validation:

$$\begin{bmatrix} M_\sigma + \Delta t C_\nu + \epsilon G_A & -\Delta t A_{Ak} \\ -A_{kA} & A_{kk} \end{bmatrix} \begin{bmatrix} \mathbf{a}^n \\ \mathbf{k}^n \end{bmatrix} = \begin{bmatrix} M_\sigma \mathbf{a}^{n-1} - \Delta \mathbf{f}_{\text{app}}^n \\ \mathbf{0} \end{bmatrix}. \quad (6.27)$$

Here $\Delta \mathbf{f}_{\text{app}}^n = \mathbf{f}_{\text{app}}^n - \mathbf{f}_{\text{app}}^{n-1}$. The small gauge term removes the vector-potential nullspace. It does not approximate the eliminated air domain.

6.2.7 Free-space hollow-sphere validation

Figure 6.1 shows the nonperiodic A - k hollow-sphere result. The regularized run used 185 output-clipped accepted steps and no rejected attempts. It required 864 Krylov iterations and 7.024 s of solve time. Across 606 samples at 1, 5, and 20 ms, the combined maximum and RMS errors were 0.193151 T and 0.012014 T. The largest error occurs at the conductor interface. The other profiles follow the expected transient. H - k and H - κ give better boundary accuracy on this mesh. The final periodic path therefore uses the magnetic field as its primary volume unknown.




Figure 6.1. Validated free-space A - k alternative using the regularized $1/r$ vector potential, FEM curl, and matched near correction. The localized interface maximum is reported rather than hidden; the path is not used as evidence for a periodic A - k kernel.

Table 6.1. Free-space regularized A– k hollow-sphere field errors.

Time	Maximum error (T)	RMS error (T)
1 ms	0.193151	0.015670
5 ms	0.099538	0.011265
20 ms	0.071901	0.007781
Combined	0.193151	0.012014

6.3 Intermediate field-oriented H– k FEM–BEM formulation

Let Ω_c be the conducting part of the quotient cell and $\Gamma_{\mathcal{C}}$ its physical conductor–insulator boundary after periodic cuts have been identified and removed. Displacement current is neglected and $\rho = \sigma^{-1}$.

6.3.1 Maxwell reduction and the H equation

Start again from Maxwell’s equations rather than from a diffusion model:

$$\begin{aligned}
 \nabla \times \mathbf{E} &= -\partial_t \mathbf{B}, & \nabla \cdot \mathbf{B} &= 0, \\
 \nabla \times \mathbf{H} &= \mathbf{J}, & \nabla \cdot \mathbf{J} &= 0, \\
 \mathbf{B} &= \mu_0(\mathbf{H} + \mathbf{M}), & \mathbf{J} &= \sigma \mathbf{E}.
 \end{aligned} \tag{6.28}$$

Write the total field as $\mathbf{H} = \mathbf{H}_a + \mathbf{h}$. The known applied field is current free in the conductor. Ampère’s law and Ohm’s law then give

$$\mathbf{J} = \nabla \times \mathbf{h}, \quad \mathbf{E} = \rho \nabla \times \mathbf{h}. \tag{6.29}$$

Insert these relations and $\mathbf{B} = \mu_0(\mathbf{H}_a + \mathbf{h} + \mathbf{M})$ into Faraday’s law. The result is the single conductor equation for $\mathbf{h}$:

$$\mu_0 \partial_t \mathbf{h} + \nabla \times (\rho \nabla \times \mathbf{h}) = -\mu_0 \partial_t (\mathbf{H}_a + \mathbf{M}) \quad \text{in } \Omega_c. \tag{6.30}$$

Setting $\mathbf{M} = \mathbf{0}$ gives the standalone problem. In the coupled problem, accepted magnetization and applied-field increments use the same time stencil as $\mathbf{h}$.

6.3.2 Interface equations and compatible spaces

Let $\mathbf{n}$ point from conductor to air. In the absence of a free surface current sheet, the interface equations are

$$\begin{aligned} \mathbf{n} \times (\mathbf{H}^+ - \mathbf{H}^-) &= \mathbf{0}, & \mathbf{n} \cdot (\mathbf{B}^+ - \mathbf{B}^-) &= 0, \\ \mathbf{n} \times (\mathbf{E}^+ - \mathbf{E}^-) &= \mathbf{0}, & \mathbf{n} \cdot \mathbf{J}^- &= 0. \end{aligned} \quad (6.31)$$

The first equation couples the magnetic trace, while the third couples the electric trace in the weak form. The last equation is the insulation condition. If the known incident field satisfies the homogeneous interface equations, the same relations hold for the reaction field.

The field energy determines the volume space. The reaction field must satisfy

$$\mathbf{h} \in H(\text{curl}, \Omega_c) := \{\mathbf{u} \in L^2(\Omega_c)^3 : \nabla \times \mathbf{u} \in L^2(\Omega_c)^3\}, \quad (6.32)$$

because both the magnetic-energy term $\|\mathbf{h}\|_{L^2}^2$ and the Ohmic term $\|\rho^{1/2} \nabla \times \mathbf{h}\|_{L^2}^2$ must be finite. Ampère's law then gives $\mathbf{J} = \nabla \times \mathbf{h} \in L^2(\Omega_c)^3$. This is the middle space in the de Rham sequence shown in Figure 6.2 (Hiptmair, 2002; Monk, 2003).

This structure explains the use of Nédélec elements. Their degrees of freedom are oriented edge circulations,

$$h_e = \int_e \mathbf{h} \cdot \mathbf{t}_e ds, \quad (6.33)$$

which appear directly in Ampère's circulation law. They enforce the tangential continuity required by $H(\text{curl})$ without forcing every vector component to be continuous. The incidence matrices also satisfy $C_h G_h = 0$ and $D_h C_h = 0$ exactly. Discrete gradients lie in the curl nullspace, and a discrete curl has no spurious divergence. On a multiply connected or periodic quotient domain, harmonic




Figure 6.2. Continuous and discrete de Rham sequences used by the conducting FEM problem. The vertical arrows denote the canonical finite-element projections Π_h^q . Nédélec elements preserve the curl sequence and pass a compatible tangential trace to the boundary discretization; the following face space carries the normal flux of $\mathbf{J} = \nabla \times \mathbf{h}$.

cycles remain. They are physical modes and must be retained in L_{harm} .

The same choice makes the FEM and BEM spaces compatible. The twisted tangential trace $\gamma_t \mathbf{h} = \mathbf{n} \times \mathbf{h}$ of an $H(\text{curl})$ field belongs to the Maxwell trace space $H^{-1/2}(\text{div}_\Gamma, \Gamma_{\mathcal{C}})$, and obeys

$$\nabla_\Gamma \cdot (\gamma_t \mathbf{h}) = -\mathbf{n} \cdot (\nabla \times \mathbf{h}). \quad (6.34)$$

The boundary edges of V_h^1 induce the same oriented tangential quantities used by the surface-current discretization. The loop map L selects the surface-divergence-free local and harmonic subspace. The lower FEM–BEM block therefore enforces a conforming Maxwell trace relation. It is not an interpolation between unrelated variables. Equation (6.34) shows why the physical boundary trace cannot contain arbitrary edge modes: it must be surface-divergence free for a closed insulated conductor. The Nédélec volume field and loop-compatible boundary trace preserve this constraint under discretization (Nedelec, 1980).

6.3.3 Weak form and the curl Green identity

Set $\mathbf{u} = \rho \nabla \times \mathbf{h} = \mathbf{E}^-$ in the curl Green identity

$$\int_{\Omega_c} (\nabla \times \mathbf{u}) \cdot \mathbf{v} dV = \int_{\Omega_c} \mathbf{u} \cdot (\nabla \times \mathbf{v}) dV + \int_{\Gamma} (\mathbf{n} \times \mathbf{u}) \cdot \mathbf{v} dS. \quad (6.35)$$

Testing Equation (6.30) with $\mathbf{v} \in H(\text{curl}, \Omega_c)$ then gives

$$(\mu_0 \partial_t \mathbf{h}, \mathbf{v})_{\Omega_c} + (\rho \nabla \times \mathbf{h}, \nabla \times \mathbf{v})_{\Omega_c} + \langle \mathbf{n} \times \mathbf{E}^-, \mathbf{v} \rangle_{\Gamma} = -(\mu_0 \partial_t (\mathbf{H}_a + \mathbf{M}), \mathbf{v})_{\Omega_c}. \quad (6.36)$$

Tangential-E continuity replaces $\mathbf{n} \times \mathbf{E}^-$ by $\mathbf{n} \times \mathbf{E}^+$. The boundary term is therefore not dropped. BEM must evaluate it from the exterior field.

6.3.4 BEM closure and boundary traces

The boundary representation follows from Green's second identity. Define the scalar layer potentials

$$\begin{aligned} (\mathcal{S}q)(\mathbf{x}) &= \int_{\Gamma} G(\mathbf{x}, \mathbf{y}) q(\mathbf{y}) dS_y, \\ (\mathcal{D}\lambda)(\mathbf{x}) &= \int_{\Gamma} \partial_{n_y} G(\mathbf{x}, \mathbf{y}) \lambda(\mathbf{y}) dS_y. \end{aligned} \quad (6.37)$$

Here $\mathbf{y}$ is the source, $\mathbf{x}$ is the target, and n_y is the source normal. Apply Green's second identity in a truncated exterior domain. The term at infinity vanishes. The outward normal of the air domain is $-\mathbf{n}$ on Γ . A harmonic scalar component u therefore satisfies

$$u(\mathbf{x}) = \mathcal{D}\lambda(\mathbf{x}) - \mathcal{S}q(\mathbf{x}), \quad \lambda = \gamma^+ u, \quad q = \partial_n u, \quad (6.38)$$

This formula is Green's representation of the same exterior solution. It is not an added boundary model.

The H- k form uses the Maxwell-compatible vector counterpart. A tangential density with

$\nabla_\Gamma \cdot \mathbf{k} = 0$ gives

$$\mathbf{H}^+ = \nabla \times \mathcal{S}_\Lambda \mathbf{k}. \quad (6.39)$$

The vector single layer is divergence free because

$$\nabla \cdot \mathcal{S}_\Lambda \mathbf{k} = \int_\Gamma G_\Lambda(\mathbf{x}, \mathbf{y}) \nabla_\Gamma \cdot \mathbf{k}(\mathbf{y}) dS_y = 0. \quad (6.40)$$

This step uses $\nabla_x G = -\nabla_y G$ and surface integration by parts on the closed quotient surface. Since the layer is harmonic away from Γ , $\nabla \times \mathbf{H}^+ = \nabla \nabla \cdot \mathcal{S}_\Lambda \mathbf{k} - \Delta \mathcal{S}_\Lambda \mathbf{k} = \mathbf{0}$. Its curl is also divergence free. Thus the representation satisfies both source-free Maxwell equations in air.

The exterior representation also supplies the electric trace required by the boundary term. If $\mathcal{S}_\Lambda$ is the componentwise periodic single-layer potential and $\mathbf{H}^+ = \nabla \times \mathcal{S}_\Lambda \mathbf{k}$, then a compatible exterior electric field is

$$\mathbf{E}^+ = -\mu_0 \mathcal{S}_\Lambda \partial_t \mathbf{k} - \nabla \psi. \quad (6.41)$$

Let $\boldsymbol{\tau} = \mathbf{v} \times \mathbf{n}$ be the corresponding surface test field. Its loop-space part is surface-divergence free. Surface integration by parts then gives

$$\langle \nabla_\Gamma \psi, \boldsymbol{\tau} \rangle_\Gamma = -\langle \psi, \nabla_\Gamma \cdot \boldsymbol{\tau} \rangle_\Gamma = 0. \quad (6.42)$$

The electric scalar potential has not been assumed to vanish. Its gradient is orthogonal to the loop test space. The remaining term is the time derivative of the vector single layer. It produces the upper-right mass-like block, with its sign fixed by the stored trace orientation.

6.3.5 Surface and volume discretization

Let $\{\mathbf{N}_i\}$ be the volume edge basis. Split the conductor state into interior edge coefficients $\mathbf{h}_I$ and compatible boundary-trace coefficients $\mathbf{b}$,

$$\mathbf{z} = \begin{bmatrix} \mathbf{h}_I \\ \mathbf{b} \end{bmatrix}, \quad \mathbf{h}_h = T_H \mathbf{z}, \quad T_H = \begin{bmatrix} R_I & R_\Gamma L \end{bmatrix}. \quad (6.43)$$

R_I injects the interior edges. R_Γ fixes the oriented boundary-edge trace, and L maps loop and harmonic coefficients into the divergence-free surface space. The physical trace $\mathbf{b}$ and representation density $\mathbf{k}$ are different unknowns. Identifying them would omit the one-sided magnetic trace operator.

The unreduced local volume matrices are

$$(M_H)_{ij} = \int_{\Omega_c} \mu_0 \mathbf{N}_j \cdot \mathbf{N}_i dV, \quad (6.44)$$

$$(K_\rho)_{ij} = \int_{\Omega_c} \rho (\nabla \times \mathbf{N}_j) \cdot (\nabla \times \mathbf{N}_i) dV. \quad (6.45)$$

and the matrices acting on the constrained state are

$$M_c = T_H^T M_H T_H, \quad K_c = T_H^T K_\rho T_H. \quad (6.46)$$

The edge space has a well-defined tangential trace and a curl represented by the discrete de Rham complex. Its nullspace can be handled without imposing spurious nodal continuity on the vector field.

The exterior reaction field is represented by a tangential surface current $\mathbf{k}$ rather than by an exterior volume mesh. Its vector potential and magnetic field are

$$\mathbf{A}_k(\mathbf{x}) = \frac{\mu_0}{4\pi} \int_{\Gamma_{\mathcal{C}}} G_\Lambda(\mathbf{x} - \mathbf{y}) \mathbf{k}(\mathbf{y}) dS_y, \quad (6.47)$$

$$\mathbf{H}_k(\mathbf{x}) = \frac{1}{4\pi} \int_{\Gamma_{\mathcal{C}}} \nabla_x G_\Lambda(\mathbf{x} - \mathbf{y}) \times \mathbf{k}(\mathbf{y}) dS_y. \quad (6.48)$$

In free space, where $G(\mathbf{r}) = 1/|\mathbf{r}|$, the second equation becomes

$$\mathbf{H}_k(\mathbf{x}) = \frac{1}{4\pi} \int_{\Gamma} \frac{\mathbf{k}(\mathbf{y}) \times (\mathbf{x} - \mathbf{y})}{|\mathbf{x} - \mathbf{y}|^3} dS_y. \quad (6.49)$$

Equation (6.49) shows directly why the accuracy of the periodic gradient kernel, rather than only the scalar potential, determines the exterior magnetic field.

Let $\{\boldsymbol{\psi}_m\}$ be tangential surface-current basis functions and let L map loop and harmonic coefficients to a discrete divergence-free current,

$$\mathbf{k}_h = \sum_m (L\mathbf{k})_m \boldsymbol{\psi}_m, \quad B_{\Gamma} L = 0. \quad (6.50)$$

Three raw surface pairings are required:

$$(G_{\Gamma})_{mn} = \langle \boldsymbol{\psi}_m, \boldsymbol{\psi}_n \rangle_{\Gamma_{\mathcal{C}}}, \quad (6.51)$$

$$(Q_{\Lambda})_{mn} = \langle \boldsymbol{\psi}_m, \gamma_t^+ \nabla \times \mathcal{S}_{\Lambda} \boldsymbol{\psi}_n \rangle_{\Gamma_{\mathcal{C}}}, \quad (6.52)$$

$$(P_{\Lambda})_{mn} = \langle \boldsymbol{\psi}_m, \gamma \mathcal{S}_{\Lambda} \boldsymbol{\psi}_n \rangle_{\Gamma_{\mathcal{C}}}. \quad (6.53)$$

Their loop restrictions are

$$G_L = L^T G_{\Gamma} L, \quad Q_L = L^T Q_{\Lambda} L, \quad P_L = L^T P_{\Lambda} L. \quad (6.54)$$

G_L is sparse. Q_L is the second-kind magnetic trace and contains the principal value and half jump. P_L is the vector $1/r$ single-layer pairing. Only Q_L and P_L are long range.

Their application follows the same boundary-to-grid-to-boundary map used by the other BEM paths. Surface basis coefficients are evaluated at quadrature sources and spread to the B-AIM grid. FFT convolution produces the potential or analytic gradient. Interpolation and surface testing return the raw boundary residual, which L^T maps to loop space. When a regularized grid kernel is used, matched panel replacement removes its local error and restores the singular BEM operator.

No dense P_L or Q_L matrix is formed.

For a cell-contained surface, the columns of L are ordinary contractible loops. On a quotient surface the coordinate map separates them from harmonic cycles that wind around periodic directions,

$$L = \begin{bmatrix} L_{\text{loc}} & L_{\text{harm}} \end{bmatrix}. \quad (6.55)$$

The harmonic columns carry the nonzero equivalent-current modes discussed in Chapter 5. The exterior tangential trace obeys the jump relation

$$\gamma_t \mathbf{H}_k = \left(\frac{1}{2} I + \mathcal{H}_\Lambda \right) \mathbf{k} \quad \text{on } \Gamma_{\mathcal{C}}, \quad (6.56)$$

with the sign of the half jump fixed by the chosen normal and trace orientation. Testing the transmission condition in the loop space gives

$$G_L \mathbf{b} - Q_L \mathbf{k} = \mathbf{0}. \quad (6.57)$$

This equation is the discrete distinction between the physical conductor trace $\mathbf{b}$ and the surface density $\mathbf{k}$ that generates the exterior.

6.3.6 From the two weak equations to four blocks

Let B extract $\mathbf{b}$ from $\mathbf{z}$, and define the source functional by

$$\delta \mathbf{z}^T F_s(\mathbf{H}_a + \mathbf{M}) = (\mathbf{H}_a + \mathbf{M}, T_H \delta \mathbf{z})_{\Omega_c}. \quad (6.58)$$

The exterior electric field in Equation (6.41), tested by the physical boundary trace $L\mathbf{b}$, contributes $\mu_0 B^T P_L \dot{\mathbf{k}}$ to the conductor row. Combining that term with the instantaneous trace constraint gives

the semidiscrete four-block system

$$\begin{bmatrix} M_c & \mu_0 B^T P_L \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \dot{\mathbf{z}} \\ \dot{\mathbf{k}} \end{bmatrix} + \begin{bmatrix} K_c & 0 \\ G_L B & -Q_L \end{bmatrix} \begin{bmatrix} \mathbf{z} \\ \mathbf{k} \end{bmatrix} = \begin{bmatrix} -\mu_0 \frac{d}{dt} F_s(\mathbf{H}_a + \mathbf{M}) + \mathbf{f}_{\text{inc}} \\ \mathbf{0} \end{bmatrix}. \quad (6.59)$$

Each block now has a defined origin. M_c and K_c are the constrained conductor mass and resistive curl–curl actions. The time derivative of the vector single layer P_L supplies the exterior electric coupling. $G_L B$ tests the physical FEM trace, while Q_L applies the one-sided exterior magnetic trace. The lower row is a transmission constraint, not an evolution equation. The system has no exterior mass or air stiffness matrix. Neither long-range block is stored densely.

Table 6.2. Maxwell origin of the four time-discrete H–k blocks.

Block	Action	Physical origin
A_{HH}	$M_c / \Delta t + K_c$	Faraday induction with Ampere’s law and Ohm’s law in the conductor
A_{Hk}	$\mu_0 B^T P_L / \Delta t$	Exterior Faraday law and tangential-E continuity; vector $1/r$ single layer
A_{kH}	$G_L B$	Physical conductor side of tangential-H continuity
A_{kk}	$-Q_L$	One-sided exterior magnetic trace; principal-value $\mathbf{r}/r^3$ action plus the half jump

6.3.7 BDF time discretization

Let a variable-step BDF formula at t_n be written as

$$D_t \mathbf{x}^n = \frac{1}{\Delta t_n} \sum_{j=0}^p \alpha_j \mathbf{x}^{n-j}. \quad (6.60)$$

Apply this stencil to both time-dependent terms in the first row of Equation (6.59). The algebraic trace row is evaluated only at t_n . One BDF solve is

$$\begin{bmatrix} \frac{\alpha_0}{\Delta t_n} M_c + K_c & \frac{\alpha_0}{\Delta t_n} \mu_0 B^T P_L \\ G_L B & -Q_L \end{bmatrix} \begin{bmatrix} \mathbf{z}^n \\ \mathbf{k}^n \end{bmatrix} = \begin{bmatrix} \mathbf{r}_H^n \\ \mathbf{0} \end{bmatrix}, \quad (6.61)$$

where

$$\begin{aligned} \mathbf{r}_H^n = & \mathbf{f}_{\text{inc}}^n - \frac{\mu_0}{\Delta t_n} \sum_{j=0}^p \alpha_j F_s(\mathbf{H}_a^{n-j} + \mathbf{M}^{n-j}) \\ & - \frac{1}{\Delta t_n} \sum_{j=1}^p \alpha_j (M_c \mathbf{z}^{n-j} + \mu_0 B^T P_L \mathbf{k}^{n-j}). \end{aligned} \quad (6.62)$$

The first sum differentiates the known magnetic drive, while the second carries the accepted H and boundary-current states. A rejected step changes neither history. Changing the BDF order affects only the coefficients and the number of stored states.

For a backward-Euler step from t_n to t_{n+1} , multiply the evolution row by Δt and keep the algebraic trace row unscaled. The result is

$$\begin{bmatrix} M_c + \Delta t K_c & \mu_0 B^T P_L \\ G_L B & -Q_L \end{bmatrix} \begin{bmatrix} \mathbf{z}^{n+1} \\ \mathbf{k}^{n+1} \end{bmatrix} = \begin{bmatrix} M_c \mathbf{z}^n + \mu_0 B^T P_L \mathbf{k}^n - \mu_0 \Delta F_s^{n+1} + \Delta t \mathbf{f}_{\text{inc}}^{n+1} \\ \mathbf{0} \end{bmatrix}, \quad (6.63)$$

where $\Delta F_s^{n+1} = F_s(\mathbf{H}_a^{n+1} + \mathbf{M}^{n+1}) - F_s(\mathbf{H}_a^n + \mathbf{M}^n)$. The previous surface current appears on the right-hand side because $P_L \mathbf{k}$ is an exterior magnetic-energy term. Dropping it would change the transient, even if every instantaneous boundary action were correct.

The energy structure provides a compact sign and adjointness check. From Equation (6.57), $\mathbf{b} = C_L \mathbf{k}$ with $C_L = G_L^{-1} Q_L$. The exterior energy matrix is

$$E_L = C_L^T P_L, \quad \mathbf{k}^T E_L \mathbf{k} = \|\mathbf{H}^+(L\mathbf{k})\|_{\Omega^+}^2 \geq 0. \quad (6.64)$$

Loss of symmetry or nonnegativity indicates an error in the trace sign, quadrature, or adjoint pairing. This check is performed before a transient is run.

At this stage, the formulation is independent of the boundary kernel. Periodicity changes the surface topology and long-range BEM action. It does not change the conductor de Rham complex. The quotient-cell construction and periodic Green functions are introduced after the free-space validation.

6.3.8 Free-space hollow-sphere validation

The free-space test is a hollow conducting sphere with conductivity 5×10^8 S/m and permeability μ_0 . A uniform applied B_z ramps from zero to 1 T in 0.01 ms. The cubed-sphere mesh has 16 divisions along each face direction and 10 radial layers. It contains 15,360 hexahedra and 16,918 vertices. The final double-precision CUDA calculation uses 100 fixed backward-Euler steps of 0.2 ms. It uses an analytic-gradient B-AIM trace, the vector single-layer time action, and a regularization length of 1 mm. All 100 steps converged. The linear solves required 1,693 iterations in total, with a maximum of 29 in one step and 15 at the final step.

Figure 6.3 compares 606 axial samples with the analytical transient at 1, 5, and 20 ms. The maximum discrepancy across both axis segments is 0.04022 T. The combined RMS discrepancy is 0.00447 T. The individual RMS errors fall from about 6–7 mT at 1 ms to 2–3 mT at 20 ms. The profiles check the reaction-field sign, vector-potential coupling, direct $\mathbf{r}/r^3$ field, and time history without an air mesh.




Figure 6.3. Free-space H–k validation for the hollow conducting sphere. Lines are the analytical transient and markers are the accelerated FEM–BEM result. Color, line pattern, and marker shape distinguish the three sample times. Shaded intervals indicate the conducting shell.

Table 6.3. Free-space hollow-sphere H- k field errors.

Time	Sample axis	Maximum error (T)	RMS error (T)
1 ms	x	0.039500	0.007279
1 ms	z	0.040217	0.005955
5 ms	x	0.010156	0.002916
5 ms	z	0.026317	0.003718
20 ms	x	0.008372	0.001625
20 ms	z	0.019274	0.002595

6.4 Final scalar-trace H- κ formulation

The final field coordinate is derived first as an H- ψ formulation. Here $\mathbf{h}$ is the reaction field in the conductor and ψ is a harmonic scalar potential in the current-free exterior. The implemented Krylov coordinate replaces the boundary value of ψ by a single-layer density κ . For this reason, the implementation is called H- κ , even though its physical fields are H- ψ .

The direct H- k formulation has two important advantages. Its volume unknown is the magnetic field, so the eddy-current field can be passed directly to LLG without taking an additional curl of a vector potential. It also separates the vector-potential and magnetic-trace kernels. However, the coupled H- k system showed slower iterative convergence than the A- k system in the hollow-sphere calculations. This convergence cost was the main reason for developing the H- ψ formulation.

H- ψ keeps the magnetic field as the conductor unknown but represents the current-free exterior by a harmonic scalar potential. In the implemented system, the boundary value of this potential is replaced by the single-layer density κ , which gives the H- κ coordinate. This path retains direct coupling to LLG and gives much better convergence than direct H- k . It is therefore used as the final path and has been verified on transient problems with 1D and 2D periodicity.

6.4.1 Maxwell reduction and the $\mathbf{H}$ equation

The starting equations are

$$\begin{aligned}\nabla \times \mathbf{E} &= -\partial_t \mathbf{B}, & \nabla \cdot \mathbf{B} &= 0, \\ \nabla \times \mathbf{H} &= \mathbf{J} + \partial_t \mathbf{D}, & \mathbf{J} &= \sigma \mathbf{E}, \\ \mathbf{B} &= \mu_0(\mathbf{H} + \mathbf{M}).\end{aligned}\tag{6.65}$$

Magnetoquasistatics neglects $\partial_t \mathbf{D}$ but retains Faraday induction. Let $\mathbf{h}$ be the eddy reaction field.

Collect the known induction-driving terms in

$$\mathbf{s}_\Omega = \mathbf{H}_{\text{app}} + \mathbf{H}_{\text{demag}} + \mathbf{M}.\tag{6.66}$$

In the conductor, Ampère's and Ohm's laws give

$$\mathbf{J}_{\text{eddy}} = \nabla \times \mathbf{h}, \quad \mathbf{E} = \rho \nabla \times \mathbf{h}.\tag{6.67}$$

Substitution into Faraday's law leaves one equation for the unknown field:

$$\mu_0 \partial_t \mathbf{h} + \nabla \times (\rho \nabla \times \mathbf{h}) = -\mu_0 \partial_t \mathbf{s}_\Omega \quad \text{in } \Omega_c.\tag{6.68}$$

In the current-free exterior, $\nabla \times \mathbf{h}^+ = \mathbf{0}$ and $\nabla \cdot \mathbf{h}^+ = 0$. After the required periodic or topological modes are included, the exterior field can be written as

$$\mathbf{h}^+ = \nabla \psi, \quad \Delta \psi = 0.\tag{6.69}$$

6.4.2 Interface equations and the constrained trace

Let $\mathbf{n}$ point from conductor to air. The interface equations are

$$\begin{aligned}\mathbf{n} \times (\mathbf{H}^+ - \mathbf{H}^-) &= \mathbf{0}, & \mathbf{n} \cdot (\mathbf{B}^+ - \mathbf{B}^-) &= 0, \\ \mathbf{n} \times (\mathbf{E}^+ - \mathbf{E}^-) &= \mathbf{0}, & \mathbf{n} \cdot \mathbf{J}^- &= 0.\end{aligned}\tag{6.70}$$

Define the exterior Dirichlet and Neumann traces

$$\lambda = \gamma^+ \psi, \quad q = \partial_n^+ \psi.\tag{6.71}$$

Tangential-H continuity becomes the interface equation

$$\mathbf{n} \times \mathbf{h}^- = \mathbf{n} \times \nabla_\Gamma \lambda.\tag{6.72}$$

This equation is imposed by the trial space. It is not imposed by a penalty. For an oriented boundary edge $e = (a, b)$,

$$\int_e \nabla_\Gamma \lambda_h \cdot \boldsymbol{\tau}_e ds = \lambda_h(b) - \lambda_h(a).\tag{6.73}$$

Hence the full Nédélec edge vector is

$$\mathbf{h}_{\text{full}} = T_H \mathbf{h}_I + T_\lambda \boldsymbol{\lambda}.\tag{6.74}$$

T_H inserts interior edge moments. T_λ forms boundary edge moments from nodal differences. An additive constant is harmless because $T_\lambda \mathbf{1} = 0$.

6.4.3 Weak form and elimination of the air volume

Applying the conductor equation alone would leave an unknown electric trace. The clean derivation returns to Faraday's law on all of space. Let $\tilde{\mathbf{h}}$ equal $\mathbf{h}$ in the conductor and $\nabla \psi$ in the

air. After moving the known drive to the right,

$$\mu_0 \partial_t \tilde{\mathbf{h}} + \nabla \times \mathbf{E} = -\mu_0 \partial_t \mathbf{s}. \quad (6.75)$$

Choose a test pair (ξ, ζ) satisfying the same tangential constraint as Equation (6.72). Let Z_ζ be the decaying harmonic exterior function with trace ζ . Define

$$\tilde{\xi} = \begin{cases} \xi, & \mathbf{x} \in \Omega_c, \\ \nabla Z_\zeta, & \mathbf{x} \in \Omega^+. \end{cases} \quad (6.76)$$

The matching tangential traces make this a global $H(\text{curl})$ test field. Its exterior curl is zero.

The curl Green identity is

$$\int_{\Omega} (\nabla \times \mathbf{a}) \cdot \mathbf{b} \, dV = \int_{\Omega} \mathbf{a} \cdot (\nabla \times \mathbf{b}) \, dV + \int_{\partial\Omega} (\mathbf{n}_{\Omega} \times \mathbf{a}) \cdot \mathbf{b} \, dS. \quad (6.77)$$

Apply it separately in the conductor and air. The interface terms cancel because the two outward normals have opposite signs and tangential $\mathbf{E}$ is continuous. The term at infinity vanishes. Since $\nabla \times \nabla Z_\zeta = \mathbf{0}$, the curl term becomes

$$(\nabla \times \mathbf{E}, \tilde{\xi})_{\mathbb{R}^3} = (\rho \nabla \times \mathbf{h}, \nabla \times \xi)_{\Omega_c}. \quad (6.78)$$

The magnetic time term splits into conductor and air parts:

$$(\partial_t \tilde{\mathbf{h}}, \tilde{\xi})_{\mathbb{R}^3} = (\partial_t \mathbf{h}, \xi)_{\Omega_c} + (\nabla \partial_t \psi, \nabla Z_\zeta)_{\Omega^+}. \quad (6.79)$$

To remove the air integral, truncate the air by a large sphere. Green's first identity gives

$$\begin{aligned} (\nabla \partial_t \psi, \nabla Z_\zeta)_{\Omega_R^+} &= - \int_{\Omega_R^+} (\Delta \partial_t \psi) Z_\zeta dV \\ &\quad + \int_{\partial \Omega_R^+} \partial_{n_a} (\partial_t \psi) Z_\zeta dS. \end{aligned} \quad (6.80)$$

The volume term is zero because ψ is harmonic. The outer-sphere term vanishes by decay. At the conductor boundary the outward air normal is $\mathbf{n}_a = -\mathbf{n}$, so $\partial_{n_a} \psi = -q$. Taking the sphere to infinity therefore gives

$$(\nabla \partial_t \psi, \nabla Z_\zeta)_{\Omega^+} = -\langle \partial_t q, \zeta \rangle_\Gamma. \quad (6.81)$$

Substitution into the tested global Faraday equation gives

$$\mu_0(\partial_t \mathbf{h}, \boldsymbol{\xi})_{\Omega_c} - \mu_0 \langle \partial_t q, \zeta \rangle_\Gamma + (\rho \nabla \times \mathbf{h}, \nabla \times \boldsymbol{\xi})_{\Omega_c} = -\mu_0 \frac{d}{dt} \mathcal{G}_s(\boldsymbol{\xi}, \zeta). \quad (6.82)$$

$\mathcal{G}_s$ is the pairing of the known driving field with the global test field. The same exterior Green step reduces any curl-free air part to its boundary trace. Its Galerkin coefficient vector is denoted by $\mathbf{g}_s$ below. The exterior has no resistive term because its test field is a gradient. It does have magnetic energy. That energy is exactly the boundary term with q . No air contribution has been discarded.

6.4.4 BEM closure and the κ coordinate

The free-space fundamental solution is $G_0(\mathbf{x}, \mathbf{y}) = 1/(4\pi|\mathbf{x}-\mathbf{y}|)$. Define the layer potentials

$$\begin{aligned} (\mathcal{V} q)(\mathbf{x}) &= \int_\Gamma G_0(\mathbf{x}, \mathbf{y}) q(\mathbf{y}) dS_y, \\ (\mathcal{K} \lambda)(\mathbf{x}) &= \int_\Gamma \partial_{n_y} G_0(\mathbf{x}, \mathbf{y}) \lambda(\mathbf{y}) dS_y. \end{aligned} \quad (6.83)$$

The normal in $\mathcal{H}$ is the source normal. Apply Green's second identity in the truncated exterior. The outer term vanishes, and the inner air normal is $-\mathbf{n}$. The result is

$$\psi = \mathcal{H} \lambda - \mathcal{V} q. \quad (6.84)$$

The single layer is continuous. The exterior double-layer trace is $(\frac{1}{2}I + K) \lambda$. Taking the exterior trace therefore gives

$$\lambda = (\frac{1}{2}I + K) \lambda - Vq, \quad Vq + (\frac{1}{2}I - K) \lambda = 0. \quad (6.85)$$

This boundary equation is Green's representation written in trace variables. It replaces the unbounded air domain.

Solving Equation (6.85) for q would require a V solve inside every outer Krylov application. H- κ instead chooses a single-layer density as the trial coordinate:

$$\psi = \mathcal{V} \kappa, \quad (\mathcal{V} \kappa)(\mathbf{x}) = \frac{1}{4\pi} \int_{\Gamma} \frac{\kappa(\mathbf{y})}{|\mathbf{x} - \mathbf{y}|} dS_y. \quad (6.86)$$

Its two physical traces are

$$\lambda = V\kappa, \quad q = \left(-\frac{1}{2}I + K'\right) \kappa, \quad -q = J\kappa := \left(\frac{1}{2}I - K'\right) \kappa. \quad (6.87)$$

K' contains the target-normal derivative $\mathbf{n}_x \cdot \nabla_x G_0$ and uses the $\mathbf{r}/r^3$ kernel. The half identity is the one-sided jump. The Calderón identity $KV = VK'$ verifies that these traces satisfy Equation (6.85). Thus κ , λ , and q describe one harmonic exterior field. They are not the same quantity.

With continuous Q1 basis functions, Galerkin projection gives

$$M_{\Gamma} \lambda = V_{11} \kappa, \quad \lambda = L\kappa, \quad L := M_{\Gamma}^{-1} V_{11}, \quad (6.88)$$

and

$$J = \frac{1}{2}M_\Gamma - K'_{11}. \quad (6.89)$$

$V_{11}\boldsymbol{\kappa}$ is a weak load, and the surface mass inverse converts it to nodal trace coefficients. DynaMag uses a fixed degree-24 Chebyshev map for this operation.

V_{11} and K'_{11} are dense only as mathematical operators; DynaMag applies them by B-AIM. Sources are spread to a grid, the Green convolution is performed by FFT, and the result is interpolated to test points. A regularized grid kernel can make this convolution smooth; matched local panel replacement then restores the singular BEM integral. The periodic H- κ path instead uses the analytic unregularized periodic kernel with explicit self and shared-edge treatment. Neither path stores a dense matrix.

6.4.5 From the weak equations to four blocks

Let M_H be the full-edge magnetic mass, including μ_0 , and let C_H be the resistive curl-curl matrix. After substituting $\boldsymbol{\lambda} = L\boldsymbol{\kappa}$, set

$$\mathbf{x} = \begin{bmatrix} \mathbf{h}_I \\ \boldsymbol{\kappa} \end{bmatrix}. \quad (6.90)$$

The time operator obtained from Equation (6.82) is

$$\mathcal{B}_{H\kappa} = \begin{bmatrix} T_H^\top M_H T_H & T_H^\top M_H T_\lambda L \\ T_\lambda^\top M_H T_H & T_\lambda^\top M_H T_\lambda L + \mu_0 J \end{bmatrix}. \quad (6.91)$$

The J term is the discrete form of $-\mu_0 \langle \partial_t q, \zeta \rangle_\Gamma$. The resistive operator is

$$\mathcal{C}_{H\kappa} = \begin{bmatrix} T_H^\top C_H T_H & T_H^\top C_H T_\lambda L \\ T_\lambda^\top C_H T_H & T_\lambda^\top C_H T_\lambda L \end{bmatrix}. \quad (6.92)$$

The semidiscrete equation is

$$\mathcal{B}_{H\kappa} \dot{\mathbf{x}} + \mathcal{C}_{H\kappa} \mathbf{x} = -\mu_0 \dot{\mathbf{g}}_s. \quad (6.93)$$

The first row is obtained with an interior-edge test. The second uses a boundary-potential test and its harmonic exterior lift. These are the two variational rows behind the four algebraic blocks.

6.4.6 BDF time discretization

Use

$$D_t \mathbf{x}^n = \frac{1}{\Delta t_n} \sum_{j=0}^p \alpha_j \mathbf{x}^{n-j}. \quad (6.94)$$

Equation (6.93) becomes

$$\left(\frac{\alpha_0}{\Delta t_n} \mathcal{B}_{H\kappa} + \mathcal{C}_{H\kappa} \right) \mathbf{x}^n = -\frac{1}{\Delta t_n} \sum_{j=1}^p \alpha_j \mathcal{B}_{H\kappa} \mathbf{x}^{n-j} - \frac{\mu_0}{\Delta t_n} \sum_{j=0}^p \alpha_j \mathbf{g}_s^{n-j}. \quad (6.95)$$

Only the complete accepted state $[\mathbf{h}_I, \boldsymbol{\kappa}]$ enters the BDF history; rejected states do not. Variable order changes the coefficients and history length, but not the spatial operators.

Divide by $\alpha_0/\Delta t_n$ and define

$$\gamma = \frac{\Delta t_n}{\alpha_0}, \quad F_\gamma = M_H + \gamma C_H. \quad (6.96)$$

The normalized four-block matrix is

$$\mathcal{A}_{H\kappa} = \begin{bmatrix} T_H^\top F_\gamma T_H & T_H^\top F_\gamma T_\lambda L \\ T_\lambda^\top F_\gamma T_H & T_\lambda^\top F_\gamma T_\lambda L + \mu_0 J \end{bmatrix}. \quad (6.97)$$

The lower row remains in the physical λ test coordinate. Only the trial coordinate is changed from λ to κ . There is therefore no $L^\top$ on the lower row.

One matrix-free application evaluates $V\boldsymbol{\kappa}$ and $K'\boldsymbol{\kappa}$ with a shared B-AIM source spectrum. It applies the surface Riesz map and reconstructs the complete $H(\text{curl})$ edge field. The sparse volume operators then form the two test rows, and $\mu_0 J\boldsymbol{\kappa}$ is added to the lower row. The time history stores only $[\mathbf{h}_I, \boldsymbol{\kappa}]$. The reconstructed $\boldsymbol{\lambda}$ and $J\boldsymbol{\kappa}$ are workspaces.

Table 6.4. Origin of the four H- κ blocks.

Block	Algebra	Maxwell role
A_{HH}	$T_H^T F_\gamma T_H$	Interior magnetic mass and resistive diffusion
$A_{H\kappa}$	$T_H^T F_\gamma T_\lambda L$	Boundary potential reconstructed by $V\kappa$ and inserted into the volume field
$A_{\lambda H}$	$T_\lambda^T F_\gamma T_H$	Interior field tested by admissible boundary-potential variations
$A_{\lambda\kappa}$	$T_\lambda^T F_\gamma T_\lambda L + \mu_0 J$	Boundary-gradient volume response plus the exterior Neumann trace

6.4.7 Free-space operator and hollow-sphere validation

The fixed Chebyshev inverse uses the measured scaled interval $[0.105, 1.001]$. It requires 23 sparse surface-mass products. Against a sparse direct solve, the relative error is 1.58×10^{-7} in λ and 2.18×10^{-7} in the recovered boundary field. Table 6.5 compares the complete accelerated boundary action with a streamed direct oracle.

Table 6.5. H- κ accelerated boundary action against the direct oracle.

Quantity	Analytic-gradient relative error	Centered-gradient error
Weak single layer V	0.2463%	Not applicable
Weak adjoint double layer K'	0.6537%	4.6502%
Strong trace λ	0.2403%	Not applicable
Complete J action	0.2439%	1.7687%

The hollow-sphere runs at 1, 5, and 20 ms accepted 12, 32, and 107 adaptive BDF steps. They required 123, 256, and 637 GMRES iterations, with no more than 16 in one solve. No step was rejected. A warm 20-ms repeat used 0.4555 s for setup and 3.6088 s for integration. Its maximum/RMS field error was 25.624/2.002 mT.

Table 6.6 places the two direct-field paths side by side. H- k required 1,693 iterations over 100 fixed backward-Euler steps. H- κ required 637 iterations over 107 adaptive BDF steps. The mean counts were 16.93 and 5.95 iterations per accepted step. The maximum counts were 29 and 16.

The measured H- κ path therefore used about 38% as many total iterations and 35% as

Table 6.6. GMRES convergence of the H- k and H- κ hollow-sphere calculations on the cubed-sphere mesh with 15,360 hexahedra.

Path	Time scheme	Steps	Total iterations	Mean/step	Maximum/step
H- k	Backward Euler, $\Delta t = 0.2$ ms	100	1,693	16.93	29
H- κ	Adaptive BDF1-3	107	637	5.95	16

many iterations per accepted step. This is a comparison of the complete tested paths. The time integrators and preconditioners differ, so the ratio is not a coordinate-only convergence factor.




Figure 6.4. Free-space H- κ validation with the analytic-gradient boundary operator and variable-order BDF integration. The figure is used for physical accuracy; current kernel timing uses the later shared-spectrum implementation.

Using $\mathbf{H}$ as the electromagnetic unknown also lets DynaMag insert the reaction field directly into the LLG effective field. Section 6.6 gives the accepted-step coupling and the conducting-disk switching result.

6.4.8 Periodic H- κ boundary operator

Chapter 5 derives and validates the periodic $1/r$ Green function and its analytic gradient. For the FEM-BEM solve, paired entities are identified on one quotient cell, and artificial periodic

caps are removed from the physical conductor surface. The conductor matrices and the four-block structure in Equation (6.97) do not change; only V_{11} and K'_{11} use periodic boundary actions.

The periodic single layer still constructs one harmonic exterior field,

$$\psi = \mathcal{V}_\Lambda \kappa, \quad \lambda = V_\Lambda \kappa, \quad -q = J_\Lambda \kappa = \left(\frac{1}{2}I - K'_\Lambda\right) \kappa. \quad (6.98)$$

V_Λ uses the finite-part periodic $1/r$ kernel, while K'_Λ uses its target-normal analytic gradient. For 1D or 2D periodicity, the finite-part convention removes the divergent additive potential constant but retains the finite line or sheet field in the gradient. This is consistent with the trace map because $T_\lambda \mathbf{1} = 0$: a potential gauge does not change the reconstructed magnetic field.

DynaMag does not force a nonzero-mean κ to zero in the partial-periodic path. It applies the finite-part policy to V_Λ and retains the native zero mode in K'_Λ . The periodic path currently uses zero kernel regularization. Exact self-panel replacement treats the singular part of V_Λ , while graded shared-edge replacement treats the singular and jump contributions in K'_Λ . Smooth periodic images remain in the global B-AIM convolution.

Table 6.7. Boundary actions in the periodic H- κ formulation.

Action	Periodic kernel	Physical role
$V_\Lambda \kappa$	Finite-part scalar $1/r$	Reconstructs the exterior potential trace λ
$K'_\Lambda \kappa$	Target-normal analytic $\mathbf{r}/r^3$	Reconstructs the exterior Neumann trace and its line or sheet zero mode

6.5 Periodic H- κ cylinder and plate validation

The next test places the periodic V_Λ/K'_Λ boundary operator inside the full quotient-cell H- κ transient. The cylinder and plate/slab contain conducting volume unknowns and are driven by a time-dependent applied field.

The end-to-end transient uses a 50 Hz applied field

$$\mathbf{B}_a(t) = B_0 \sin(\omega t) \mathbf{d}, \quad B_0 = 1 \text{ T}, \quad \omega = 2\pi(50 \text{ Hz}), \quad (6.99)$$

and conductivity $5 \times 10^7 \text{ S/m}$. The cylinder with 1D periodicity has $R = 10 \text{ mm}$ and $L_z = 25 \text{ mm}$. The slab with 2D periodicity has half-thickness $a = 10 \text{ mm}$, with $L_x = 18 \text{ mm}$ and $L_y = 12 \text{ mm}$. Both cases start from zero field and take 200 steps with $\Delta t = 0.1 \text{ ms}$. The first two steps use BDF1 and BDF2. The remaining steps use BDF3. The cylinder mesh has 5,120 hexahedra, divided 64 times azimuthally, 16 times radially, and 4 times axially. It has 15,108 H(curl) unknowns and 256 κ unknowns. The slab mesh has $24^3 = 13,824$ hexahedra, 40,320 H(curl) unknowns, and 1,152 κ unknowns.

The tests use CUDA, the production periodic B-AIM implementation, AMS for the H block, and a fixed AMGX cycle for the κ surrogate. Every linear solve converged. The cylinder required 7,808 GMRES iterations over 200 steps, with a maximum of 60 in one step. The slab required 2,982 iterations, with a maximum of 20. For reference, the earlier H- k runs on the same spatial and temporal discretizations required 21,968 and 4,595 iterations. The preconditioners are not identical, so these counts compare the complete tested solver paths. They do not isolate the coordinate transformation alone.

For comparison with the zero-initial-condition diffusion solution, define

$$q(\lambda, t) = \frac{\omega [\lambda \cos(\omega t) + \omega \sin(\omega t) - \lambda e^{-\lambda t}]}{\lambda^2 + \omega^2}. \quad (6.100)$$

If α_n are the positive zeros of J_0 , the cylinder field is

$$\frac{B_z(r, t)}{B_0} = \sin(\omega t) - \sum_{n=1}^{\infty} \frac{2J_0(\alpha_n r/R)}{\alpha_n J_1(\alpha_n)} q\left(\frac{\alpha_n^2}{\mu_0 \sigma R^2}, t\right). \quad (6.101)$$

For the slab, with $\beta_m = (2m + 1)\pi/2$,

$$\frac{B_x(z, t)}{B_0} = \sin(\omega t) - \sum_{m=0}^{\infty} \frac{4(-1)^m}{(2m + 1)\pi} \cos(\beta_m z/a) q\left(\frac{\beta_m^2}{\mu_0 \sigma a^2}, t\right). \quad (6.102)$$

These are startup transients, not steady-state phasor curves.

Table 6.8. Refined periodic H- κ transient errors against analytical diffusion.

Geometry	Time	Maximum error (mT)	RMS error (mT)
Cylinder	1 ms	1.223	0.639
Cylinder	5 ms	0.805	0.382
Cylinder	20 ms	1.124	0.494
Slab	1 ms	1.510	0.538
Slab	5 ms	0.928	0.389
Slab	20 ms	1.530	0.441




Figure 6.5. Refined periodic H- κ FEM-BEM fields and analytical diffusion solutions under the sinusoidal source. Lines show the analytical transient, while markers show the FEM-BEM result. Color, line pattern, and marker shape distinguish the three sample times. Shading marks the conductor. At 20 ms the applied sine has returned to zero, while the conductor retains a phase-lagged diffusive field.

Table 6.8 and Figure 6.5 show combined maximum/RMS errors of 1.223/0.516 mT for the cylinder and 1.530/0.460 mT for the slab. On the matched BDF3 runs, the earlier H- κ values

were 4.280/1.477 mT and 1.539/0.453 mT. $H-\kappa$ is more accurate for the refined cylinder and effectively tied for the slab. At 20 ms, the applied field has completed one period and returned to zero. The conductor still carries a finite diffusive field.

The profiles use direct periodic evaluation outside the conductor. For targets within roughly one surface-panel width, the local replacement in Equation (7.2) removes the inaccurate central-cell grid contribution and restores the panel action. This output treatment does not change the time-dependent state. It separates near-surface field evaluation from the accuracy of the coupled solve.

The validation has three levels. First, the hollow sphere checks the free-space $H-\kappa$ transient. Next, Chapter 5 tests the periodic scalar kernel, its analytic gradient, and their zero modes. Finally, the sinusoidal cylinder and slab exercise the complete quotient-cell evolution. Each level has an independent analytical reference, although the current results do not establish an asymptotic convergence order.

Table 6.9 compares the three formulations. All three avoid an airbox: Nédélec elements resolve the conductor and an accelerated boundary action resolves the exterior. $A-k$ is the production baseline for the TEAM11 comparison. $H-k$ is the intermediate direct-field form. $H-\kappa$ is the final path for periodic and LLG-coupled work.

Table 6.9. Validation status of the three FEM–BEM coordinates.

Coordinate	Validation evidence	Current limitation
$A-k$	Regularized free-space $1/r$ potential, FEM curl, hollow-sphere transient, and block preconditioner	No validated regularized periodic Green convention in the production path
$H-k$	Free-space sphere and matched periodic cylinder/slab reference runs	Intermediate path; periodic Krylov cost is higher in the reported tests and LLG coupling is not integrated
$H-\kappa$	Direct $V/K'/J$ oracles, analytical free-space sphere, periodic cylinder/slab transients, and integrated disk	Coupled nonneutral periodic transient and full refinement studies remain

Moving geometry adds a lifecycle dependency, so DynaMag currently guards simultaneous

mesh motion with $H-k/H-\kappa$. The evolving coordinates have not yet been connected atomically to every geometry-dependent operator. The runtime still needs a coordinated refresh and a moving-conductor validation, but the FEM–BEM formulation itself does not change.

6.6 Integrated LLG–Maxwell coupling

DynaMag couples the LLG and FEM–BEM solvers on a fixed conducting mesh. Its dynamic-state layer separates topology, coordinates, material frames, and time-dependent properties. Within this structure, hybrid superposition computes the demagnetizing field without exterior air elements, while FEM–BEM computes the induced-current field without a dense boundary matrix. Both solvers exchange state and history within one time-integration procedure.

On a fixed conducting mesh, DynaMag implements accepted-step LLG–eddy coupling with one Maxwell update per accepted LLG step. In the reported switching calculation, the eddy reaction changes the reversal time. This result is reported for the dense mesh described below. Simultaneous mesh motion and eddy-current evolution also remain incomplete because all geometry- dependent data and history transformations must be connected to the mesh- update path.

6.6.1 Coupled physical model

Let $\mathbf{m} = \mathbf{M}/M_s$ be normalized magnetization. The LLG equation is

$$\frac{\partial \mathbf{m}}{\partial t} = -\frac{\gamma}{1 + \alpha^2} [\mathbf{m} \times \mathbf{H}_{\text{eff}} + \alpha \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}})], \quad |\mathbf{m}| = 1. \quad (6.103)$$

In the coupled runtime the effective field is assembled as

$$\mathbf{H}_{\text{eff}} = \mathbf{H}_{\text{an}} + \mathbf{H}_{\text{app}} + \mathbf{H}_{\text{ex}} + \mathbf{H}_{\text{d}} + \mathbf{H}_{\text{eddy}}. \quad (6.104)$$

The first four terms are the anisotropy, applied, exchange, and demagnetizing fields. The fifth is the reaction field recovered from the FEM–BEM Maxwell state. It contains only the induced-

current response; it does not duplicate the applied or demagnetizing fields already present in Equation (6.104).

The Maxwell source is the physical induction associated with magnetization and the non-eddy magnetic field,

$$\mathbf{B}_{\text{src}} = \mu_0 (\mathbf{H}_{\text{app}} + \mathbf{H}_{\text{d}} + \mathbf{M}). \quad (6.105)$$

The LLG solver uses Oersted for field, emu/cm³ for M_s , nanometers for length, and nanoseconds for time, whereas the Maxwell solver uses SI. Each source sample is therefore converted as

$$\mathbf{B}_{\text{src}}[\text{T}] = 10^{-4} [\mathbf{H}_{\text{app}}[\text{Oe}] + \mathbf{H}_{\text{d}}[\text{Oe}] + 4\pi M_s[\text{emu/cm}^3] \mathbf{m}]. \quad (6.106)$$

The anisotropy and exchange fields do not appear in Equation (6.106); they are thermodynamic effective fields in LLG rather than independent sources in Maxwell's induction.

Jacobian-weighted element contributions project the recovered edge field to LLG vertices. The field is then converted using

$$1 \text{ A/m} = \frac{4\pi}{1000} \text{ Oe}. \quad (6.107)$$

The same field state also gives the current-density output,

$$\mathbf{J}_{\text{eddy}} = \nabla \times \mathbf{H}_{\text{eddy}} \quad [\text{A/m}^2], \quad (6.108)$$

using the physical curl of the edge basis. Both recovery paths run on the selected provider, and the current-density output reuses the edge-field reconstruction and nodal weights. It therefore does not create a second persistent electromagnetic state.

6.6.2 Implemented accepted-step coupling

The implemented coupling is partitioned. At accepted time t_n , the LLG state $\mathbf{m}_n$, Maxwell history $\mathcal{E}_n$, and cached reaction $\mathbf{H}_{\text{eddy},n}$ are consistent. An attempted step to $t_{n+1} = t_n + \Delta t_n$ follows four stages.

1. The implicit LLG solver attempts $\mathbf{m}_{n+1}$ using $\mathbf{H}_{\text{eddy},n}$. If LLG rejects the attempt, the accepted Maxwell source and $\mathcal{E}_n$ remain unchanged.
2. If LLG accepts, the runtime reevaluates $\mathbf{H}_{\text{app}}(\mathbf{m}_{n+1}, t_{n+1})$ and $\mathbf{H}_{\text{d}}(\mathbf{m}_{n+1}, t_{n+1})$ and forms Equation (6.106).
3. The new source is projected into the FEM–BEM test coordinate. The Maxwell BDF step is solved once, while the source remains a trial load.
4. On convergence, the Maxwell state and source history are committed at t_{n+1} . The reaction field is recovered, and all outputs are interpolated to the same scheduled times.

In compact notation, the completed one-pass update is

$$\mathbf{m}_{n+1} = \mathcal{L}_{\Delta t_n}(\mathbf{m}_n; \mathbf{H}_{\text{eddy},n}), \quad (6.109)$$

$$\mathbf{s}_{n+1} = \mathcal{P}(\mathbf{m}_{n+1}, \mathbf{H}_{\text{app},n+1}, \mathbf{H}_{\text{d},n+1}), \quad (6.110)$$

$$(\mathcal{E}_{n+1}, \mathbf{H}_{\text{eddy},n+1}) = \mathcal{M}_{\Delta t_n}(\mathcal{E}_n; \mathbf{s}_{n+1}). \quad (6.111)$$

Here $\mathcal{L}$ is the implicit LLG step, $\mathcal{P}$ is the device-side source transfer, and $\mathcal{M}$ is the transient FEM–BEM solve.

With this ordering, a rejected LLG attempt cannot advance the electromagnetic BDF stencil or consume a future applied field. The ordering also introduces one explicit lag because LLG at t_{n+1} uses the reaction from t_n . Both subsystems may use BDF3 internally, but the order of the coupled method must still be measured directly.

Two extensions can reduce this lag. A predictor can extrapolate $\mathbf{H}_{\text{eddy}}$ to the trial time before LLG. Alternatively, partitioned fixed-point or Newton iterations can repeat the coupled updates until both fields agree at t_{n+1} . A monolithic solve is also possible, but it requires both cross-Jacobian actions. The first integration study should measure the one-pass method under time-step refinement.

Provider-resident transfer. The coupling does not stage full fields through the host. Accepted magnetization is converted from the outer double-precision state to the field precision on the selected provider, where the applied and demagnetizing fields are also evaluated. The Maxwell right-hand side remains on the provider through source projection, solution, and field recovery. Only scheduled output and explicit diagnostics cross to the host.

Both solvers use the SzCompute view and provider model of Chapter 3. The eddy reaction is a normal fifth LLG field, not a special application callback. It uses the same configuration, construction, accumulation, output, and test paths as the other effective-field terms.

The two solvers use different precision for different purposes. LLG and the outer Maxwell Krylov state are retained in double precision. The large physical FEM/BEM/B-AIM data and recovered fields use single precision. The unit conversion in Chapter 6.106 occurs before physical projection, so a mixed cgs/SI value is never stored as a nominal Maxwell state.

Geometry-consistent coupling. For a moving or morphing calculation, all fields at one trial time must refer to one geometry state. At minimum this includes

$$\mathcal{G}(t) = \{\mathbf{x}, J_e, \mathbf{R}, M_s, \sigma, \mu, \mathcal{Q}_{\text{FEM}}, \mathcal{Q}_{\text{BEM}}, \mathcal{N}_{\text{near}}, \mathcal{P}_{\text{B-AIM}}, \mathcal{P}_{\text{lin}}\}_t, \quad (6.112)$$

This set contains the geometry, material state, boundary operators, and linear solver state at time t . Mixing times within this set does not define a consistent discrete Maxwell problem. One example is using current volume Jacobians with old boundary normals.

Fixed connectivity permits incidence, degree-of-freedom numbering, and sparse graphs to remain stable. It does not make all numerical values geometry invariant. A complete accepted geometry update must classify each object as one of

$$\text{retain,} \quad \text{refresh values,} \quad \text{replan,} \quad \text{rebuild.} \quad (6.113)$$

The choice must follow an invariance argument rather than the storage type. A sparse matrix can

have a fixed graph but geometry-dependent values. A B-AIM plan may keep its topology but still need new bounds, bins, and interpolation data.

Electromagnetic history also needs a defined coordinate frame and transport rule. The frame may be laboratory, material, or arbitrary-Lagrangian–Eulerian. The moving-conductor weak equation must also include the intended motional electromotive term. These are integration choices within a valid FEM–BEM framework.

The runtime therefore guards against simultaneous mesh motion and Heddy. That guard exists because Equation (6.112) is not yet updated atomically across every eddy-current object. It prevents stale geometry from producing a plausible but incorrect field. The moving-mesh H – k algorithm remains valid, but its runtime lifecycle and end-to-end validation are incomplete.

6.6.3 Integrated validation ladder

A coupled result is credible only when its component and interface errors can be separated. This dissertation uses four validation levels.

1. *Components and prepared operators*: discrete basis and unit checks, host/device agreement, direct-versus-accelerated BEM actions, and sparse solver residuals.
2. *Standalone transients*: LLG benchmarks and the analytical conducting-sphere eddy transient.
3. *Static-geometry coupling*: accepted-step history, source/reaction consistency, aligned output, and a matched no-eddy control.
4. *General dynamic coupling*: rigid and non-rigid motion, moving-conductor history, evolving materials, and periodicity.

Table 6.10 summarizes the implemented capabilities and their validation status.

Table 6.10. Integrated capabilities and their validation evidence.

Capability	Implementation	Validation evidence
Dynamic materials and frames	Constant and table models with rigid and polar frames	Unit, field, and matched-rod tests
Fixed-connectivity motion	Rigid motion and expression-driven morphing	Analytical checks and a translated-rod benchmark
Hybrid magnetostatic field	Boundary B-AIM with an interior Poisson solve	Published accuracy, memory, and timing tests
Free-space FEM–BEM	Matrix-free volume–boundary transient	Analytical sphere and direct-operator tests
LLG–eddy coupling	Accepted-step coupling on a fixed mesh	Disk reversal and a matched no-eddy control
Moving FEM–BEM	Moving-interface formulation	Geometry and history integration remain to be completed
Periodic magnetostatics	Periodic FEM and Green kernels	Published validation and current kernel tests
Periodic H– κ	Quotient-cell system with finite-part V/K' actions	Analytical cylinder and slab transients

6.6.4 Conducting-disk coupled validation

The integrated test is a conducting magnetic disk initialized along $-\mathbf{e}_z$. A positive z -directed field drives the reversal. The disk has radius $2.0\ \mu\text{m}$ and thickness $0.5\ \mu\text{m}$. Its mesh contains 145,152 hexahedra and 159,133 vertices. It uses 24 radial, 288 azimuthal, and 12 thickness segments. The mesh has 449,280 faces and 463,260 edges. Its nominal outer-rim, radial-shell, and thickness spacings are all about 42–44 nm. The Gaussian exchange length from Equation (2.10) is 44.6 nm for this material. The complete 0.7 ns calculation runs on CUDA.

The material parameters are

$$\begin{aligned}
 K_1 &= 5.0 \times 10^5 \text{ erg/cm}^3, & M_s &= 200 \text{ emu/cm}^3, & A_{\text{ex}} &= 5.0 \times 10^{-6} \text{ erg/cm}, \\
 \alpha &= 0.1, & \sigma &= 10^7 \text{ S/m}.
 \end{aligned}
 \tag{6.114}$$

The applied field rises linearly from zero to 14,000 Oe during the first 0.05 ns and then remains fixed. The initial LLG step is 0.0005 ns, and the maximum is 0.002 ns. Both solvers allow BDF order three. The Maxwell relative tolerance is 10^{-6} with a 500-iteration limit. Its B-AIM boundary

action uses second-order quadrature, a near-distance factor of two, automatic regularization, and eight targets per box.

The control calculation uses exactly the same mesh, material, applied field, LLG tolerances, and output schedule but omits $\mathbf{H}_{\text{eddy}}$ from Equation (6.104). This matched pair isolates the effect of the induced-current reaction under the stated numerical settings.

Dense-mesh coupled switching trajectory




Figure 6.6. Coupled H- κ -LLG conducting-disk result on the 145,152-element mesh. The volume-average magnetization is shown with eddy-current feedback and for the matched control.

Figure 6.6 shows the volume-average magnetization. Linear interpolation between accepted states gives the first zero crossing

$$t_{\text{switch}}^{\text{eddy}} = 0.20324210 \text{ ns}, \quad t_{\text{switch}}^{\text{control}} = 0.32023295 \text{ ns}. \quad (6.115)$$

At 0.7 ns, the averages are

$$\langle m_z \rangle_{\text{eddy}} = 0.98987206, \quad \langle m_z \rangle_{\text{control}} = 0.91265298. \quad (6.116)$$

The eddy-enabled and control runs accept 631 and 548 LLG steps, respectively.

Figure 6.6 is evaluated from the accepted-step simulation outputs. The following discussion provides its physical interpretation.

The direction of the change cannot be inferred from the statement that *eddy currents add damping*. Here, both the reaction field and reversal are strongly nonuniform. The reaction can change the nucleation site and the motion of the reversal front. Its effect on the volume-average crossing therefore depends on the full trajectory. For this case, the computed induced-current feedback advances the reported crossing.

Cost of the electromagnetic solve

The 0.7 ns eddy-enabled run on the 145,152-element mesh used 49,389 Maxwell Krylov iterations over 631 accepted steps, or 78.27 iterations per accepted step. The maximum was 136 iterations in one step. GMRES restarted four times and applied the operator 50,024 times. The trajectory used 49,389 preconditioner applications. The complete timed command took 925.74 s, including setup and finalization but excluding field export. The sampled peak GPU allocation was 7189 MiB for the whole process. This includes CUDA and library pools; it is not the storage of a dense BEM matrix. The statistics characterize the calculation at the reported resolution.

Reproducibility record. The dense-mesh disk calculations reach the requested final time and retain the accepted-step average magnetization and solver summary. The reported values in Equations (6.115) and (6.116) are evaluated from the accepted-step output, not transcribed from the plotted curves.

6.6.5 Interpretation and remaining integration

The integrated disk case supports three conclusions.

1. Heddy is a device-executed LLG field. Its Maxwell source includes the accepted magnetic state with an explicit cgs-to-SI conversion.
2. Magnetostatics and eddy current both represent the open exterior without a volumetric airbox or a dense BEM matrix.
3. The induced-current reaction materially changes the switching trajectory in the disk problem.

Path to the fully dynamic solver. The component architecture removes both exterior airboxes. Hybrid superposition handles magnetostatics, while FEM–BEM handles eddy current. Fixed-connectivity geometry and material laws preserve stable state identities. Three steps remain for full end-to-end validation.

1. Finish the atomic geometry refresh in Equation (6.112), including the $H(\text{curl})$ space, boundary geometry, B-AIM plan, near corrections, source projection, and preconditioners.
2. Define electromagnetic history transport for the moving-conductor weak equation. Verify rigid covariance, then validate a smooth non-rigid case with positive Jacobians.
3. Connect periodic $H-\kappa$ to the LLG accepted-state lifecycle. Report coordinate-update and operator-rebuild costs separately.

The periodic $H-\kappa$ formulation is validated, but its moving-mesh path is not yet fully integrated.

The implemented solver performs a device-resident, airbox-free LLG–Maxwell transient. The electromagnetic reaction changes the magnetic switching result. The calculation uses explicit source units, accepted-step history, a matrix-free volume–boundary operator, and aligned output. In the disk, the first average-magnetization zero crossings on the 145,152-element mesh are 0.20324 ns with eddy reaction and 0.32023 ns in the matched control.

The current implementation validates fixed-geometry partitioned coupling. Moving and periodic LLG integration still need operator updates, history transport, and coupling-convergence tests. These tasks do not change the underlying FEM–BEM or B-AIM formulations. Both open-boundary solvers already avoid an airbox and a dense long-range matrix.

This chapter presents airbox-free eddy-current calculations in free space and in quotient cells with 1D and 2D periodicity. It also presents fixed-geometry LLG–Maxwell coupling. Chapter 7 treats repeated application, memory ownership, preconditioning, and the TEAM11 comparison.

7 Matrix-Free FEM–BEM Performance Optimization

The FEM–BEM formulations in Chapter 6 remove the airbox, but their unaccelerated boundary operators are still dense. These operators must also be applied many times inside a Krylov solve. DynaMag addresses both costs with prepared B-AIM pipelines and fixed-cost block preconditioners. It stores neither a dense BEM matrix nor a point-pair table. Hollow-sphere tests compare the preconditioners, and the converted TEAM11 problem provides an external LS-DYNA comparison.

The implementation uses sparse local operators and a fixed number of multilevel operations. It has no tolerance-driven inner BEM solve. Its setup, memory, and repeated applications can therefore be managed explicitly when geometry changes.

7.1 Matrix-free BEM operators

The three formulations use different boundary maps but share one prepared long-range pattern. Table 7.1 separates each mathematical operator from the value returned to the finite-element row. A scalar potential alone does not supply the target-normal derivative in K' . Differentiating an interpolated grid potential is also not equivalent to using the analytical gradient kernel.

General-purpose boundary-element software demonstrates reusable Galerkin assembly and

solution of boundary integral equations (Śmigaj et al., 2015). DynaMag specializes the repeated operator to the Green kernels and trace maps in this chapter so that projection, FFT workspaces, and local replacement remain resident on the selected device.

Table 7.1. Matrix-free boundary actions in the selected FEM–BEM coordinates.

Coordinate	Block	Green action	Quantity returned
A– k	P, Q	Regularized vector $1/r$, followed by FEM curl	Exterior vector-potential trace and magnetic trace
H– k	P_L, Q_L	Vector $1/r$ and analytic $\mathbf{r}/r^3$	Exterior time-mass action and magnetic transmission trace
H– κ	V, K', J	Scalar $1/r$ plus analytic Cartesian gradient dotted with target normals	Potential trace and one-sided normal derivative

For a generic point or panel source vector $\mathbf{q}$, the accelerated action is written

$$\mathcal{B}\mathbf{q} \simeq \mathcal{I}_h \mathcal{F}^{-1} [\widehat{G}_h \odot \mathcal{F}(\mathcal{P}_h \mathbf{q})] + (\mathcal{D}_{\text{near}} - \widetilde{\mathcal{D}}_{\text{near}}) \mathbf{q}. \quad (7.1)$$

$\mathcal{P}_h$ projects nonuniform sources to the auxiliary grid, while $\widehat{G}_h$ is the prepared Green spectrum. The operator then applies the FFT convolution, and $\mathcal{I}_h$ interpolates the result to the targets. Finally, the parenthesized term replaces the inaccurate local grid interaction. The cost is about $O(N_\Gamma + N_g \log N_g)$ rather than $O(N_\Gamma^2)$. Chapter 5 validates this corrected action as a standalone operator. Here it is one block inside flexible GMRES.

When scalar and gradient outputs use the same source density, the projected charge and forward spectrum are computed once. The scalar and three derivatives then use one forward and four inverse transforms. The intermediate H– k form reuses the source preparation while keeping its $1/r$ and $\mathbf{r}/r^3$ equations distinct. The final H– κ form shares the spectrum between V and K' .

7.2 Corrected and prepared matrix-free execution

7.2.1 Matched near replacement and linear storage

The grid convolution is a smooth far-field approximation; it is not a singular-panel quadrature rule. DynaMag therefore uses the algebraically matched replacement

$$\boxed{\mathcal{B}_{\text{corrected}} = \mathcal{B}_{\text{B-AIM},\varepsilon} - \mathcal{B}_{\text{local-grid},\varepsilon} + \mathcal{B}_{\text{local-panel}}.} \quad (7.2)$$

The all-interaction and subtraction terms must use the same regularization, periodic displacement, source weights, and grid interpolant. Otherwise their difference leaves an uncontrolled local residual. The added term supplies the required singular self action, one-sided jump, or nearby panel integral. Only the central-cell contribution is replaced. Smooth periodic images remain in the global convolution.

The production A- k path evaluates the regularized potential

$$\mathbf{A}_k(\mathbf{x}_i) = \frac{\mu_0}{4\pi} \int_{\Gamma} \frac{\mathbf{k}(\mathbf{y})}{\sqrt{|\mathbf{x}_i - \mathbf{y}|^2 + \varepsilon^2}} dS_{\mathbf{y}}, \quad \mathbf{B}_k = \nabla_h \times \mathbf{A}_k, \quad (7.3)$$

at volume vertices and takes the finite-element curl. H- κ instead uses the analytic gradient of the same regularized scalar kernel for K' ; the three interpolated gradient components are dotted with the target normal. Direct H- k uses the validated free-space or periodic analytic-gradient kernel for Q_L . These paths share the replacement principle but not the same physical trace.

No dense BEM matrix is assembled. The plan retains geometry, box offsets, interpolation data, and Green spectra. Nearby nonshared panels are streamed from the box structure instead of stored in a pair table. Shared-edge corrections are local blocks built during setup. They scale linearly with the boundary mesh and occupy 0.797 MiB for the cubed-sphere mesh with 15,360 hexahedra. The persistent state is $O(N_{\Gamma} + N_g)$, with no B-AIM near-pair list.

7.2.2 Prepared ownership, mixed precision, and time integration

Preparation creates the geometry-dependent data and reusable work grids. The plan owns immutable reordered geometry and spectra. The workspace owns the grids, cached source spectrum, reorder buffers, and shared FFT work area. Repeated applications preserve these addresses. They perform no SzCompute allocation, resize, Green-table upload, or host/device round trip. A geometry change triggers explicit preparation and cannot reuse stale local data.

Physical FEM/BEM/B-AIM arrays and preconditioner actions use single precision. The outer GMRES state, residual norms, BDF history, adaptive estimator, and time use double precision. For variable-step BDF order $q_n \leq 3$,

$$\dot{\mathbf{x}}_n \simeq \frac{1}{\Delta t_n} \sum_{j=0}^{q_n} \alpha_j^{(n)} \mathbf{x}_{n-j}. \quad (7.4)$$

History states are combined before one time-mass action. A higher BDF order therefore does not multiply the number of BEM evaluations. Trial state and history are committed only after step acceptance.

The H- κ hollow sphere shows the combined effect. A warm 20-ms repeat used 0.4555 s for setup and 3.6088 s for integration. Linear solves accounted for 2.8885 s. The run accepted 107 steps with no rejection and used 637 outer iterations with no restart. Its maximum/RMS field error was 25.624/2.002 mT. Reporting time, iterations, and error together separates kernel fusion from a change in the physical discretization.

7.3 TEAM11 runtime and memory comparison

The largest external comparison uses a direct conversion of the LS-DYNA TEAM11 hollow-sphere mesh. The electromagnetic volume has 49,168 nodes and 43,008 hexahedra, while its discrete de Rham complex has 141,326 $H(\text{curl})$ edge unknowns and 12,290 boundary-loop unknowns. The production A- k system therefore has 153,616 unknowns. The benchmark

uses regularized free-space $1/r$ B-AIM, matched near replacement, adaptive BDF1–3, and the mixed-precision policy above.

Table 7.2. DynaMag A– k benchmark on the converted TEAM11 mesh. Timing excludes solution and axis-sample export.

Metric	Measured value
Simulated time	20 ms
Accepted/rejected adaptive steps	185/0
Total GMRES iterations/restarts	929/0
Integration time, excluding export	7.51 s
Warm-repeat integration time	7.56 s
Setup, startup, and teardown bound	0.93 s
Application and library allocation	210 MiB
Steady application and HYPRE allocation	371 MiB
Temporary peak during AMS setup	1,340 MiB
Context-inclusive sampled device peak	1829 MiB
Dense BEM matrix storage	none

The temporary peak in Table 7.2 occurs during construction of the HYPRE/AMS hierarchy. It is not the persistent boundary operator footprint. After releasing the hierarchy data, the steady combined allocation is 371 MiB. This state has no dense boundary matrix or general near-pair table.

The same conductor mesh and excitation were run in LS-DYNA on four CPU cores. Each electromagnetic interval represents 10^{-4} s, so the 20 completed intervals cover 2 ms.

The two tables do not give a hardware-normalized speedup because DynaMag ran on an NVIDIA GeForce RTX 5090, while LS-DYNA used four CPU cores. Their completed physical horizons also differ by a factor of ten. However, the preparation and storage comparison remains direct because the mesh and excitation match. Linear extrapolation of the LS-DYNA interval mean gives about 21,500 s for 200 intervals after preprocessing. The projected end-to-end time for 20 ms is approximately 22,500 s. This is an extrapolation, not a completed LS-DYNA run.

Table 7.3. Four-core LS-DYNA runtime and memory on the TEAM11 mesh. The 20 electromagnetic intervals cover 2 ms.

Metric	Measured value
Preprocessing and FEM–BEM assembly	939 s
Time for 20 electromagnetic intervals	2,152 s
Mean time per interval	107.6 s
Median time per interval	107 s
Minimum–maximum interval time	107–114 s
Population standard deviation	1.5 s
End-to-end time for 20 intervals	3,091 s
Private committed memory after first solve	6.24 GiB
Resident memory after first solve	3.55 GiB
Peak resident memory	3.57 GiB

7.4 Block preconditioning and solver selection

7.4.1 Block factorization and the operator that should be approximated

At one implicit time step, each coordinate can be written abstractly as

$$\mathcal{A} \begin{bmatrix} \mathbf{u} \\ \mathbf{q} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix}, \quad \mathcal{A} = \begin{bmatrix} A & B \\ C & D \end{bmatrix}. \quad (7.5)$$

$\mathbf{u}$ is the volume edge coordinate, and $\mathbf{q}$ is the boundary coordinate. A is a time-weighted mass-plus-curl operator. B and C couple the volume and boundary traces. D contains the boundary metric and exterior action.

With A invertible on the constrained volume space, the exact factorization is

$$\mathcal{A} = \begin{bmatrix} I & 0 \\ CA^{-1} & I \end{bmatrix} \begin{bmatrix} A & 0 \\ 0 & S \end{bmatrix} \begin{bmatrix} I & A^{-1}B \\ 0 & I \end{bmatrix}, \quad S = D - CA^{-1}B. \quad (7.6)$$

This identity leads directly to both the lower and Schur forms. A fixed lower preconditioner applies

$$\mathbf{z}_u = \tilde{A}^{-1} \mathbf{r}_u, \quad (7.7)$$

$$\mathbf{z}_q = \tilde{D}^{-1} (\mathbf{r}_q - C \mathbf{z}_u). \quad (7.8)$$

It retains the exact available sparse lower coupling but replaces the Schur complement by a boundary surrogate. An upper correction adds

$$\mathbf{z}_u \leftarrow \mathbf{z}_u - \tilde{A}^{-1} B \mathbf{z}_q. \quad (7.9)$$

Equations (7.7)–(7.9) do not form $CA^{-1}B$. The same matrix-free FEM–BEM operator evaluates $B\mathbf{z}_q$ and the physical residual.

7.4.2 Volume edge block

The volume block has the generic structure

$$A(\alpha_0, \Delta t) = \alpha_0 M_\mu + \Delta t C_\rho, \quad (7.10)$$

where M_μ is an $H(\text{curl})$ mass matrix and C_ρ is the resistive curl–curl term. Gradients lie in or near the curl–curl nullspace. Classical incomplete factorizations remain useful sparse-system tools (Scott and Tuma, 2023), but by themselves they do not encode this discrete gradient subspace. Scalar AMG on the edge matrix likewise does not represent the structure. The CUDA paths use one fixed auxiliary-space Maxwell cycle (AMS, preset f2) as $\tilde{A}^{-1}$. Its gradient and coordinate operators are prepared at a reference time weight. The physical operator still uses the current BDF coefficient.

One fixed cycle defines the outer-solver interpretation. AMS is a right preconditioner application, not an inner Krylov solve with its own tolerance. The outer method is flexible right-preconditioned GMRES, which measures the physical residual while allowing fixed multilevel and

mixed-precision actions.

7.4.3 Hollow-sphere evidence

The hollow conducting sphere has an analytical transient, while its FEM–BEM block retains the trace, nullspace, and time-mass structure of an application case. The following tests use A – k and direct H – k . Their matrices differ, so their absolute iteration counts are not directly comparable. Each test therefore measures a preconditioner within one coordinate.

A – k diagonal and Schur–Neumann comparison

Table 7.4 gives the controlled comparison on the 15,360-element cubed-sphere mesh. For $\mu_r = 1$, the diagonal path is already easy. Schur–Neumann adds online AK , KA , AMS , and boundary-AMG work, so it is slower. At $\mu_r = 4$, it reduces iterations by 28.8%. The small sphere still does not amortize the extra work.

Table 7.4. A – k hollow-sphere preconditioner ablation.

μ_r	Block form	Accepted	Rejected	Iterations	Solve time (s)
1	Diagonal	105	0	814	4.41
1	Depth-3 Schur–Neumann	106	0	825	8.94
4	Diagonal	105	0	770	4.20
4	Depth-3 Schur–Neumann	105	0	548	6.42

The same depth-three Schur–Neumann form helps the larger $\mu_r = 4$ long cylinder. It reduces the iteration count from 2,292 to 1,097 and solve time from 14.95 to 13.33 s. The comparison uses the best block-diagonal sparse surrogate as its baseline. DynaMag therefore selects the block form from the strength and topology of the coupling.

Direct H – k sphere

The direct H – k hollow sphere uses 100 fixed backward-Euler steps and the depth-three Schur preconditioner. AMS acts on the interior $H(\text{curl})$ block, and AMG acts on a sparse

boundary/surface Schur surrogate. The current run converges in 1,693 iterations. The maximum is 29, and the final step takes 15. The combined maximum/RMS field errors are 0.04022/0.00447 T. Neither the operator nor the preconditioner stores a dense BEM matrix. The final H- κ path uses 637 iterations over 107 adaptive steps, with a mean of 5.95 and a maximum of 16 per step. Table 6.6 gives the direct comparison and its time-integration caveat.

7.4.4 Selection rules for coupled dynamics

The experiments give four implementation rules.

1. Preserve the de Rham structure of the edge block with a fixed auxiliary-space action; scalar diagonal scaling is not an adequate replacement.
2. Keep the boundary surrogate sparse and fixed-cost. A preconditioner that performs a separately converged boundary solve has merely moved the original problem into a hidden inner iteration.
3. Use the lower block when volume–boundary coupling is moderate. Add a matrix-free upper or Schur–Neumann correction only when measured iteration reduction repays its extra BEM and AMS actions.
4. Rebuild coefficient- and geometry-dependent preconditioner data when their mathematical inputs change. The currently guarded moving-mesh H- κ path reflects incomplete lifecycle wiring of this rule, not a failure of the FEM–BEM formulation.

These rules preserve the airbox-free, matrix-free design. They also keep the cost of preconditioning visible in the outer iteration count.

8 Conclusions and Future Directions

This dissertation presents DynaMag, a finite-element micromagnetic solver for dynamic geometry, material properties, and electromagnetic feedback. During fixed-connectivity motion, the mesh topology and algebraic graphs remain fixed, while geometry, material data, and long-range operators are updated when their inputs change.

DynaMag uses hybrid superposition for magnetostatics and FEM–BEM for eddy current. Hybrid superposition evaluates the boundary potential and recovers the interior through a sparse Poisson solve, while B-AIM evaluates free-space and periodic Green actions with matched near correction. The $H-k$ formulation provides an intermediate direct-field derivation. The final $H-\kappa$ formulation removes the volumetric eddy-current airbox and supports periodic transients. Sphere, cylinder, and slab tests validate the electromagnetic components, and a conducting disk demonstrates coupled LLG–Maxwell dynamics.

8.1 Summary of results

Invariant data under dynamics. Fixed-connectivity motion preserves incidence, material identity, degree-of-freedom numbering, and adjacency graphs. However, it does not preserve geometric coefficients or coordinate-dependent long-range data. This division supports rigid motion and smooth morphing, but it does not claim to solve remeshing or topology change.

Accelerator efficiency. The numerical paths use SzCompute views on CPU, NVIDIA, AMD, and Metal. In the matched CUDA rod study, constant tabulated materials were statistically

unresolved from the static baseline. The explicit mesh/frame update was 3.9% of the moving runtime. The full translated trajectory was 17.8% slower because it also included conservative magnetostatic replanning and adaptive variation. Local state updates were cheap in this test, whereas global geometry-dependent plans cost more.

Magnetostatic memory. The hybrid superposition integral–Poisson method needs long-range superposition only on selected boundaries and obtains the volume field from a sparse interior solve (Duan et al., 2023). Uniqueness of the Dirichlet problem shows that this split is exact at the continuum level. Its memory advantage comes from reducing long-range targets and near data, not from changing the underlying magnetostatic boundary condition.

Open and periodic boundary operators. A prepared B-AIM sequence projects sources, performs the grid convolution, interpolates targets, and replaces local interactions. The periodic $\mathbf{r}/r^3$ kernel is differentiated analytically and includes the finite partial-period zero modes. Direct validation gives a relative L^2 error of 2.34×10^{-3} for the cylinder with 1D periodicity and 3.20×10^{-9} for the slab with 2D periodicity. The latter also reproduces the 7.5×10^5 A/m jump to within 2.39×10^{-3} A/m. In contrast, the ordinary periodic-Poisson $1/r$ inverse remains subject to zero-mean compatibility. The finite-part action used by H- κ fixes only the divergent additive constant and retains the finite nonzero line or sheet field in the gradient.

Dynamic LLG–Maxwell coupling. Validation proceeds from individual operators to electromagnetic transients, accepted-state history, and coupled observables. It includes analytical motion tests, periodic-gradient comparisons, and fixed-geometry LLG–Maxwell coupling. On the 145,152-element disk mesh, the first average-magnetization zero crossings are 0.20324 ns with eddy reaction and 0.32023 ns in the matched control.

8.2 Principal contributions

The main contribution of this work is broader than a single micromagnetic solver. DynaMag provides a numerical structure in which geometry, material state, field operators, and hardware execution are separated. A change in one part of the physical model does not require the entire solver to be rewritten. This structure makes it possible to study systems in which several physical fields evolve together.

The foundation combines SzCompute with a general finite-element layer. SzCompute provides the same vector, matrix, reduction, and memory operations on CPU, CUDA, ROCm, and Metal backends. The finite-element layer provides H^1 and Nédélec spaces, matrix-free element actions, sparse solvers, and block preconditioners. AMS, AMG, and the block-solver interface can therefore be reused when another field is added. The field may introduce new coefficients or unknowns without replacing the memory and solver foundation.

The on-the-fly evaluation pattern is equally important. Element operators use the current geometry and material properties instead of relying on a large set of fixed cached matrices. Long-range terms are also applied without dense storage. This keeps memory use controlled and avoids tying a simulation to one geometry. Elasticity, heat transfer, and fluid dynamics can build on the same mesh state and execution model. Their additional element spaces and preconditioners can be incorporated through the existing interfaces.

The magnetostatic and eddy-current methods remove the main long-range barriers to this design. Hybrid superposition limits the expensive magnetostatic interaction to selected boundaries and recovers the volume field with a sparse solve. FEM–BEM represents the unbounded electromagnetic exterior without an airbox, while B-AIM makes the repeated boundary action practical. Electromagnetic feedback can therefore follow moving or morphing magnetic bodies without requiring an exterior volume mesh.

Together, these contributions turn DynaMag into a foundation for exploring new coupled physics. Dynamic magnetization is no longer isolated from changes in shape, material state, or electromagnetic response. The same structure can support further physical models while preserving

accelerator performance and low memory use.

8.3 Future work

The first direction is to improve preconditioning efficiency. The coupled FEM–BEM systems still spend substantial work in Krylov iteration, especially as the mesh and coupling strength increase. Better sparse Schur approximations should reduce this cost. The multilevel hierarchy should also be reused across time steps and updated only when geometry or coefficients change. These improvements must preserve the matrix-free boundary action and avoid a hidden inner solve.

The second direction is distributed execution. Large device and material models will require several GPUs and eventually several compute nodes. The volume mesh, boundary sources, and B-AIM grid must be partitioned together. Communication should be overlapped with element evaluation and FFT work. Distributed AMS, AMG, and block preconditioning will then provide the solver foundation needed for large coupled systems.

The third direction is validation against experiments. Analytical solutions remain essential for individual operators, but they do not reproduce the uncertainty of a fabricated device or measured material. Future studies should use measured geometry, material data, and excitation. Switching behavior, magnetic loss, and frequency-dependent response can then be compared directly with experimental observations.

The fourth direction is broader multiphysics coupling. Thermal fields can change magnetic and electrical properties, while elastic deformation can change geometry and anisotropy. Fluid motion can also transport or load magnetic structures. These models should first be added through controlled one-way coupling and then extended to two-way transient feedback. Their state updates can use the same accepted-step, memory, and operator lifecycle already used by DynaMag.

8.4 Summary

DynaMag supports dynamic LLG with time-dependent materials and fixed-connectivity mesh motion. It also supports free-space and periodic $H-\kappa$ eddy current and fixed-geometry LLG–Maxwell coupling. The long-range operators use neither dense matrices nor a volumetric airbox. Their state remains on the selected provider during repeated applications.

The next stage is to improve coupled preconditioning, extend the framework to distributed systems, validate it against experiments, and add thermal, elastic, and fluid models. These developments can reuse the finite-element, FEM–BEM, B-AIM, and accelerator-portable structure established here.

REFERENCES

- Abert, C., Hrkac, G., Page, M., Praetorius, D., Ruggeri, M., Suess, D., 2014. Spin-polarized transport in ferromagnetic multilayers: An unconditionally convergent FEM integrator. *Computers & Mathematics with Applications* 68, 639–654.
- Aharoni, A., 2001. *Introduction to the Theory of Ferromagnetism*. Second ed., Oxford University Press.
- Aharoni, A., Shtrikman, S., 1958. Magnetization Curve of the Infinite Cylinder. *Physical Review* 109, 1522–1528.
- Ai, F., Duan, J., Lomakin, V., 2025a. Periodic micromagnetic finite element method. *Journal of Magnetism and Magnetic Materials* 615, 172776.
- Ai, F., Lin, Z., Duan, J., Lomakin, V., 2025b. Periodic Phase Diagrams in Micromagnetics With an Eigenvalue Solver. *IEEE Transactions on Magnetics* 61, 1–5.
- Ai, F., Lomakin, V., 2024. Fast Fourier Transform periodic interpolation method for superposition sums in a periodic unit cell. *Computer Physics Communications* 304, 109291.
- Allwood, D.A., Xiong, G., Faulkner, C.C., Atkinson, D., Petit, D., Cowburn, R.P., 2005. Magnetic Domain-Wall Logic. *Science* 309, 1688–1692.
- Alouges, F., 2008. A new finite element scheme for Landau-Lifchitz equations. *Discrete and Continuous Dynamical Systems - S* 1, 187–196.
- Alouges, F., Kritsikis, E., Steiner, J., Toussaint, J.C., 2014. A convergent and precise finite element scheme for Landau–Lifschitz–Gilbert equation. *Numerische Mathematik* 128, 407–430.
- Alouges, F., Soyeur, A., 1992. On global weak solutions for Landau-Lifshitz equations: Existence and nonuniqueness. *Nonlinear Analysis: Theory, Methods & Applications* 18, 1071–1084.
- Arlett, P., Bahrani, A., Zienkiewicz, O., 1968. Application of finite elements to the solution of Helmholtz’s equation. *Proceedings of the Institution of Electrical Engineers* 115, 1762–1766.
- Aurada, M., Feischl, M., Praetorius, D., 2012. Convergence of some adaptive FEM-BEM coupling for elliptic but possibly nonlinear interface problems. *ESAIM: Mathematical Modelling and Numerical Analysis* 46, 1147–1173.
- Baczewski, A.D., Dault, D.L., Shanker, B., 2012. Accelerated Cartesian Expansions for the Rapid

- Solution of Periodic Multiscale Problems. *IEEE Transactions on Antennas and Propagation* 60, 4281–4290.
- Bagge, J., Tornberg, A.K., 2023. Fast Ewald summation for Stokes flow with arbitrary periodicity. *Journal of Computational Physics* 493, 112473.
- Bagnères-Viallix, A., Baras, P., 1991. On a method to calculate the demagnetizing field in a micromagnetic structure. *Journal of Applied Physics* 69, 4599–4601.
- Barnett, A., Greengard, L., 2011. A new integral representation for quasi-periodic scattering problems in two dimensions. *BIT Numerical Mathematics* 51, 67–90.
- Bartels, S., 2015. Projection-free approximation of geometrically constrained partial differential equations. *Mathematics of Computation* 85, 1033–1049.
- Bartels, S., Ko, J., Prohl, A., 2007. Numerical analysis of an explicit approximation scheme for the Landau-Lifshitz-Gilbert equation. *Mathematics of Computation* 77, 773–788.
- Bartels, S., Prohl, A., 2006. Convergence of an Implicit Finite Element Method for the Landau-Lifshitz-Gilbert Equation. *SIAM Journal on Numerical Analysis* 44, 1405–1419.
- Bauer, G.E.W., Saitoh, E., van Wees, B.J., 2012. Spin caloritronics. *Nature Materials* 11, 391–399.
- Bañas, L., Bartels, S., Prohl, A., 2008. A Convergent Implicit Finite Element Discretization of the Maxwell-Landau-Lifshitz-Gilbert Equation. *SIAM Journal on Numerical Analysis* 46, 1399–1422.
- Berger, L., 1996. Emission of spin waves by a magnetic multilayer traversed by a current. *Physical Review B* 54, 9353–9358.
- Berkov, D.V., Gorn, N.L., 1998. Quasistatic remagnetization processes in two-dimensional systems with random on-site anisotropy and dipolar interaction: Numerical simulations. *Physical Review B* 57, 14332–14343.
- Berkov, D.V., Gorn, N.L., 2002. Thermally activated processes in magnetic systems consisting of rigid dipoles: equivalence of the Ito and Stratonovich stochastic calculus. *Journal of Physics: Condensed Matter* 14, L281–L287.
- Bleszynski, E., Bleszynski, M., Jaroszewicz, T., 1996. AIM: Adaptive integral method for solving large-scale electromagnetic scattering and radiation problems. *Radio Science* 31, 1225–1251.
- Boag, A., Michielssen, E., Brandt, A., 2002. Nonuniform polar grid algorithm for fast field evaluation. *IEEE Antennas and Wireless Propagation Letters* 1, 142–145.
- Bossavit, A., 1985. Two dual formulations of the 3-D eddy-currents problem. *COMPEL - The international journal for computation and mathematics in electrical and electronic engineering* 4, 103–116.
- Bossavit, A., 1998. *Computational Electromagnetism: Variational Formulations, Complementarity, Edge Elements*. Academic Press.

- Brown, Jr., W.F., 1963. Micromagnetics. Number 18 in Interscience Tracts on Physics and Astronomy, Interscience Publishers, New York.
- Bruckner, F., Ducevic, A., Heistracher, P., Abert, C., Suess, D., 2021. Strayfield calculation for micromagnetic simulations using true periodic boundary conditions. *Scientific Reports* 11, 9202.
- Buffa, A., Ciarlet, P., 2001a. On traces for functional spaces related to Maxwell's equations Part I: An integration by parts formula in Lipschitz polyhedra. *Mathematical Methods in the Applied Sciences* 24, 9–30.
- Buffa, A., Ciarlet, P., 2001b. On traces for functional spaces related to Maxwell's equations Part II: Hodge decompositions on the boundary of Lipschitz polyhedra and applications. *Mathematical Methods in the Applied Sciences* 24, 31–48.
- Capolino, F., Wilton, D., Johnson, W., 2007. Efficient computation of the 3D Green's function for the Helmholtz operator for a linear array of point sources using the Ewald method. *Journal of Computational Physics* 223, 250–261.
- Carbou, G., Fabrie, P., 1998. Time Average in Micromagnetism. *Journal of Differential Equations* 147, 383–409.
- Chang, R., Li, S., Lubarda, M.V., Livshitz, B., Lomakin, V., 2011. FastMag: Fast micromagnetic simulator for complex magnetic structures (invited). *Journal of Applied Physics* 109, 07D358.
- Chew, W.C., Jin, J.M., Lu, C.C., Michielssen, E., Song, J., 1997. Fast solution methods in electromagnetics. *IEEE Transactions on Antennas and Propagation* 45, 533–543.
- Chubykalo-Fesenko, O., Nowak, U., Chantrell, R.W., Garanin, D., 2006. Dynamic approach for micromagnetics close to the Curie temperature. *Physical Review B* 74, 094436.
- Cimrák, I., 2007. Existence, regularity and local uniqueness of the solutions to the Maxwell–Landau–Lifshitz system in three dimensions. *Journal of Mathematical Analysis and Applications* 329, 1080–1093.
- Cimrák, I., 2008. A Survey on the Numerics and Computations for the Landau-Lifshitz Equation of Micromagnetism. *Archives of Computational Methods in Engineering* 15, 277–309.
- Duan, J., Wang, X., Lomakin, V., 2023. Hybrid Superposition Integral—Poisson Solver Method for the Magnetostatic Field in Finite Element Micromagnetic Solvers. *IEEE Transactions on Magnetics* 59, 1–5.
- d'Aquino, M., Serpico, C., Miano, G., 2005. Geometrical integration of Landau–Lifshitz–Gilbert equation based on the mid-point rule. *Journal of Computational Physics* 209, 730–753.
- Evans, R.F.L., Hinzke, D., Atxitia, U., Nowak, U., Chantrell, R.W., Chubykalo-Fesenko, O., 2012. Stochastic form of the Landau-Lifshitz-Bloch equation. *Physical Review B* 85, 014433.
- Feischl, M., Tran, T., 2017. The Eddy Current–LLG Equations: FEM-BEM Coupling and A Priori Error Estimates. *SIAM Journal on Numerical Analysis* 55, 1786–1819.

- Fessler, J., Sutton, B., 2003. Nonuniform fast fourier transforms using min-max interpolation. *IEEE Transactions on Signal Processing* 51, 560–574.
- Forster, H., Schrefl, T., Dittrich, R., Scholz, W., Fidler, J., 2003. Fast boundary methods for magnetostatic interactions in micromagnetics. *IEEE Transactions on Magnetics* 39, 2513–2515.
- Fredkin, D., Koehler, T., 1990. Hybrid method for computing demagnetizing fields. *IEEE Transactions on Magnetics* 26, 415–417.
- Frei, E.H., Shtrikman, S., Treves, D., 1957. Critical Size and Nucleation Field of Ideal Ferromagnetic Particles. *Physical Review* 106, 446–455.
- Fu, S., Chang, R., Volvach, I., Kuteifan, M., Menarini, M., Lomakin, V., 2019. Block Inverse Preconditioner for Implicit Time Integration in Finite Element Micromagnetic Solvers. *IEEE Transactions on Magnetics* 55, 1–11.
- Garanin, D.A., 1997. Fokker-Planck and Landau-Lifshitz-Bloch equations for classical ferromagnets. *Physical Review B* 55, 3050–3057.
- Gilbert, T., 2004. Classics in Magnetism A Phenomenological Theory of Damping in Ferromagnetic Materials. *IEEE Transactions on Magnetics* 40, 3443–3449.
- Greengard, L., Lee, J.Y., 2004. Accelerating the Nonuniform Fast Fourier Transform. *SIAM Review* 46, 443–454.
- Greengard, L., Rokhlin, V., 1987. A fast algorithm for particle simulations. *Journal of Computational Physics* 73, 325–348.
- Grinstein, G., Koch, R.H., 2003. Coarse Graining in Micromagnetics. *Physical Review Letters* 90, 207201.
- Hackbusch, W., Khoromskij, B.N., 2000. A Sparse $\mathcal{H}$ -Matrix Arithmetic. Part II: Application to Multi-Dimensional Problems. *Computing* 64, 21–47.
- Henkelman, G., Jónsson, H., 2000. Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points. *The Journal of Chemical Physics* 113, 9978–9985.
- Hessel, A., Oliner, A.A., 1965. A New Theory of Wood’s Anomalies on Optical Gratings. *Applied Optics* 4, 1275.
- Hiptmair, R., 2002. Finite elements in computational electromagnetism. *Acta Numerica* 11, 237–339.
- Hu, Z., 2014. Infinite Boundary Terms of Ewald Sums and Pairwise Interactions for Electrostatics in Bulk and at Interfaces. *Journal of Chemical Theory and Computation* 10, 5254–5264.
- Johnson, C., Nédélec, J.C., 1980. On the coupling of boundary integral and finite element methods. *Mathematics of Computation* 35, 1063–1079.

- Kalinikos, B.A., Slavin, A.N., 1986. Theory of dipole-exchange spin wave spectrum for ferromagnetic films with mixed exchange boundary conditions. *Journal of Physics C: Solid State Physics* 19, 7013–7033.
- Kalkstein, D., Soven, P., 1971. A green's function theory of surface states. *Surface Science* 26, 85–99.
- af Klinteberg, L., Shamshirgar, D.S., Tornberg, A.K., 2017. Fast Ewald summation for free-space Stokes potentials. *Research in the Mathematical Sciences* 4, 1.
- Knoll, D., Keyes, D., 2004. Jacobian-free Newton–Krylov methods: a survey of approaches and applications. *Journal of Computational Physics* 193, 357–397.
- Kruzík, M., Prohl, A., 2006. Recent Developments in the Modeling, Analysis, and Numerics of Ferromagnetism. *SIAM Review* 48, 439–483.
- Le, K.N., Page, M., Praetorius, D., Tran, T., 2015. On a decoupled linear FEM integrator for eddy-current-LLG. *Applicable Analysis* 94, 1051–1067.
- Le, K.N., Tran, T., 2013. A convergent finite element approximation for the quasi-static Maxwell–Landau–Lifshitz–Gilbert equations. *Computers & Mathematics with Applications* 66, 1389–1402.
- Lebecki, K.M., Donahue, M.J., Gutowski, M.W., 2008. Periodic boundary conditions for demagnetization interactions in micromagnetic simulations. *Journal of Physics D: Applied Physics* 41, 175005.
- Li, S., Chang, R., Boag, A., Lomakin, V., 2012. Fast Electromagnetic Integral-Equation Solvers on Graphics Processing Units. *IEEE Antennas and Propagation Magazine* 54, 71–87.
- Li, S., Livshitz, B., Lomakin, V., 2010a. Fast evaluation of Helmholtz potential on graphics processing units (GPUs). *Journal of Computational Physics* 229, 8463–8483.
- Li, S., Van Orden, D.A., Lomakin, V., 2010b. Fast Periodic Interpolation Method for Periodic Unit Cell Problems. *IEEE Transactions on Antennas and Propagation* 58, 4005–4014.
- Liu, C., Aygün, K., Yılmaz, A.E., 2020. A parallel FFT-accelerated layered-medium integral-equation solver for electronic packages. *International Journal of Numerical Modelling: Electronic Networks, Devices and Fields* 33, e2684.
- Livshitz, B., Boag, A., Bertram, H.N., Lomakin, V., 2009. Nonuniform grid algorithm for fast calculation of magnetostatic interactions in micromagnetics. *Journal of Applied Physics* 105, 07D541.
- Logg, A., Mardal, K.A., Wells, G. (Eds.), 2012. *Automated Solution of Differential Equations by the Finite Element Method*. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Lomakin, V., Michielssen, E., 2005. Enhanced transmission through metallic plates perforated by arrays of subwavelength holes and sandwiched between dielectric slabs. *Physical Review B* 71,

235117.

- Lomakin, V., Michielssen, E., 2006. Transmission of Transient Plane Waves Through Perfect Electrically Conducting Plates Perforated by Periodic Arrays of Subwavelength Holes. *IEEE Transactions on Antennas and Propagation* 54, 970–984.
- Lomakin, V., Michielssen, E., 2007. Beam Transmission Through Periodic Subwavelength Hole Structures. *IEEE Transactions on Antennas and Propagation* 55, 1564–1581.
- Mansuripur, M., Giles, R., 1988. Demagnetizing field computation for dynamic simulation of the magnetization reversal process. *IEEE Transactions on Magnetics* 24, 2326–2328.
- Marshall, S.L., 2000. A periodic Green function for calculation of coulombic lattice potentials. *Journal of Physics: Condensed Matter* 12, 4575–4601.
- Meng, J., Boag, A., Lomakin, V., Michielssen, E., 2010. A multilevel Cartesian non-uniform grid time domain algorithm. *Journal of Computational Physics* 229, 8430–8444.
- Monk, P., 2003. *Finite Element Methods for Maxwell’s Equations*. Oxford University Press.
- Nedelec, J.C., 1980. Mixed finite elements in $\mathbb{R}^3$. *Numerische Mathematik* 35, 315–341.
- Newell, A.J., Williams, W., Dunlop, D.J., 1993. A generalization of the demagnetizing tensor for nonuniform magnetization. *Journal of Geophysical Research: Solid Earth* 98, 9551–9555.
- Pei, R., Askham, T., Greengard, L., Jiang, S., 2023. A fast method for imposing periodic boundary conditions on arbitrarily-shaped lattices in two dimensions. *Journal of Computational Physics* 474, 111792.
- Peng, S., Tamir, T., Bertoni, H., 1975. Theory of Periodic Dielect Waveguides. *IEEE Transactions on Microwave Theory and Techniques* 23, 123–133.
- Phillips, J., White, J., 1997. A precorrected-FFT method for electrostatic analysis of complicated 3-D structures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 16, 1059–1072.
- Prohl, A., 2001. *Computational Micromagnetism*. Vieweg+Teubner Verlag, Wiesbaden.
- Saad, Y., Schultz, M.H., 1986. GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems. *SIAM Journal on Scientific and Statistical Computing* 7, 856–869.
- Scholz, W., Fidler, J., Schrefl, T., Suess, D., Dittrich, R., Forster, H., Tsiantos, V., 2003. Scalable parallel micromagnetic solvers for magnetic nanostructures. *Computational Materials Science* 28, 366–383.
- Schryer, N.L., Walker, L.R., 1974. The motion of 180° domain walls in uniform dc magnetic fields. *Journal of Applied Physics* 45, 5406–5421.
- Scott, J., Tũma, M., 2023. Incomplete Factorizations, in: *Algorithms for Sparse Linear Systems*.

- Birkhäuser, Cham. Nečas Center Series, pp. 185–203.
- Shamshirgar, D.S., Bagge, J., Tornberg, A.K., 2021. Fast Ewald summation for electrostatic potentials with arbitrary periodicity. *The Journal of Chemical Physics* 154, 164109.
- Slonczewski, J., 1996. Current-driven excitation of magnetic multilayers. *Journal of Magnetism and Magnetic Materials* 159, L1–L7.
- Stoner, E.C., Wohlfarth, E.P., 1948. A mechanism of magnetic hysteresis in heterogeneous alloys. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 240, 599–642.
- Suess, D., 2007. Micromagnetics of exchange spring media: Optimization and limits. *Journal of Magnetism and Magnetic Materials* 308, 183–197.
- Suess, D., Tsiantos, V., Schrefl, T., Fidler, J., Scholz, W., Forster, H., Dittrich, R., Miles, J., 2002. Time resolved micromagnetics using a preconditioned time integration method. *Journal of Magnetism and Magnetic Materials* 248, 298–311.
- Takano, K., 2005. Micromagnetic-FEM models of a perpendicular writer and reader. *IEEE Transactions on Magnetics* 41, 696–701.
- Tarjan, R.E., 1975. Efficiency of a Good But Not Linear Set Union Algorithm. *Journal of the ACM* 22, 215–225.
- Thiaville, A., Nakatani, Y., Miltat, J., Suzuki, Y., 2005. Micromagnetic understanding of current-driven domain wall motion in patterned nanowires. *Europhysics Letters (EPL)* 69, 990–996.
- Van Orden, D., Lomakin, V., 2009. Rapidly Convergent Representations for 2D and 3D Green's Functions for a Linear Periodic Array of Dipole Sources. *IEEE Transactions on Antennas and Propagation* 57, 1973–1984.
- Van Orden, D., Lomakin, V., 2012. Rapidly Convergent Representations for Periodic Green's Functions of a Linear Array in Layered Media. *IEEE Transactions on Antennas and Propagation* 60, 870–879.
- Vansteenkiste, A., Leliaert, J., Dvornik, M., Helsen, M., Garcia-Sanchez, F., Van Waeyenberge, B., 2014. The design and verification of MuMax3. *AIP Advances* 4, 107133.
- Visintin, A., 1985. On Landau-Lifshitz' equations for ferromagnetism. *Japan Journal of Applied Mathematics* 2, 69–84.
- Wilton, D., Rao, S., Glisson, A., Schaubert, D., Al-Bundak, O., Butler, C., 1984. Potential integrals for uniform and linear source distributions on polygonal and polyhedral domains. *IEEE Transactions on Antennas and Propagation* 32, 276–281.
- Wysocki, A.L., Antropov, V.P., 2017. Micromagnetic simulations with periodic boundary conditions: Hard-soft nanocomposites. *Journal of Magnetism and Magnetic Materials* 428,

274–286.

- Yao, Z., Jambunathan, R., Zeng, Y., Nonaka, A., 2022. A massively parallel time-domain coupled electrodynamics–micromagnetics solver. *The International Journal of High Performance Computing Applications* 36, 167–181.
- Yee, K., 1966. Numerical solution of initial boundary value problems involving maxwell’s equations in isotropic media. *IEEE Transactions on Antennas and Propagation* 14, 302–307.
- Zhang, S., Li, Z., 2004. Roles of Nonequilibrium Conduction Electrons on the Magnetization Dynamics of Ferromagnets. *Physical Review Letters* 93, 127204.
- Śmigaj, W., Betscke, T., Arridge, S., Phillips, J., Schweiger, M., 2015. Solving Boundary Integral Problems with BEM++. *ACM Transactions on Mathematical Software* 41, 1–40.
- Žutić, I., Fabian, J., Das Sarma, S., 2004. Spintronics: Fundamentals and applications. *Reviews of Modern Physics* 76, 323–410.