UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering and the Environment

Computational Engineering and Design Group

Computational Methods in Micromagnetics

by

Dmitri Chernyshenko

Thesis for the degree of Doctor of Philosophy

June 2016

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Doctor of Philosophy

COMPUTATIONAL METHODS IN MICROMAGNETICS

by Dmitri Chernyshenko

With the continued growth of computational power, computational modelling has become an increasingly important part of modern science. The field of micromagnetism has benefited from the increase of computational power, leading in recent decades to the development of many important micromagnetic methods. This thesis aims to address some computational challenges relevant to the field of micromagnetism today.

The computation of the demagnetising field is often the most time-consuming part of a micromagnetic simulation. In the finite difference method, this computation is usually done using the Fourier transform method, in which the demagnetising field is computed as the convolution of the magnetisation field with the demagnetising tensor. An analytical formula for the demagnetising tensor is available, however due to numerical cancellation errors it can only be applied for close distances between the interacting cells. For far distances between the interacting cells other approaches, such as asymptotic expansion, have to be used. In this thesis, we present a new method to compute the demagnetising tensor by means of numerical integration. The method offers improved accuracy over existing methods for the intermediate range of distances.

In the finite element method, the computation of the demagnetising field is commonly done using the hybrid FEM/BEM method. The fast multipole method offers potential theoretical advantages over the hybrid FEM/BEM method particularly for the current and future generations of computing hardware. In micromagnetics, it has been applied to compute the demagnetising field in the finite difference setting and to compute the magnetostatic interaction between nanoparticles, however no implementation of the fast multipole method in finite elements is yet available. As one of the steps towards it, in this thesis we develop a new formula for the energy of the magnetostatic interaction between linearly magnetized polyhedrons. This formula can be used to compute the direct interaction between finite element cells in the fast multipole method.

Ferromagnetic resonance is a popular experimental technique for probing the dynamical properties of magnetic systems. We extend the eigenvalue method for the computation of resonance modes to the computation of the FMR spectrum, and apply it to compute ferromagnetic resonance for a proposed FMR standard reference problem.

Contents

List of Figures

List of Tables

Declaration of Authorship

I, Dmitri Chernyshenko, declare that this thesis titled “Computational Methods in Micromagnetics” and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:
   (a) Chapter 3 is based on work published as “D. Chernyshenko and H. Fangohr, *Computing the demagnetizing tensor for finite dif- ference micromagnetic simulations via numerical integration*, Journal of Magnetism and Magnetic Materials, vol. 381, pp. 440–445, 2015”
   (b) Chapter 6 is based on work published as “H. Fangohr, D. S. Chernyshenko, M. Franchin, T. Fischbacher, and G. Meier, *Joule heating in nanowires*, Phys. Rev. B, vol. 84, p. 054437, 2011”

Signed: .....

Date: .....

Acknowledgements

First of all, I would like to thank my supervisor Prof. Hans Fangohr, for without his continued support even during difficult times I would not have been able to complete this work. His advice and assistance has been invaluable throughout the duration of my PhD.

I would like to thank my colleagues and collaborators who in many ways contributed to the research presented here: Maximilian Albert, Alexander Baker, Gabriel Balaban, Marijan Beg, Marc-Antonio Bisotti, Beckie Carey, David Cortes, Matteo Franchin, Anders Johansen, Ondrej Hovorka, Andreas Knittel, Guido Meier, Neil O'Brien, Mark Vousden, and Weiwei Wang. Special thanks go to my fellow PhD students, Andreas, Marc-Antonio, Marijan, Matteo, Max, and Weiwei — some of you have already finished your PhDs but your companionship and friendship was invaluable through the years.

I would like to thank Dr. Thomas Fischbacher for his valuable advice and explanations regarding micromagnetic simulation during the start of my PhD.

I would like to thank the late Prof. Peter de Groot for his teaching and advice about solid state physics and the physical side of micromagnetism. To a person with a mathematical and computational background the perspective of a physicist was immensely relevant.

I would like to thank the co-directors of the Doctoral Training Centre for Complex System Simulation, Prof. Seth Bullock and Prof. Jonathan Essex, and my Doctoral Training Centre tutor Dr. Jason Noble for their support throughout my PhD.

The funding for my PhD research was provided by the Faculty of Engineering and the Environment at the University of Southampton and by the EPSRC Centre for Doctoral Training grant EP/G03690X/1.

Finally, I would like to thank my family for their continued support. Without their help and encouragement I would have never been able to embark on, or complete, the work presented here.

List of publications

Below is the complete list of publications utilising some of the work described in this thesis:

1. H. Fangohr, D. S. Chernyshenko, M. Franchin, T. Fischbacher, and G. Meier, "Joule heating in nanowires," *Physical Review B*, vol. 84, p. 054437 (2011)

2. D. Chernyshenko and H. Fangohr, "Computing the demagnetizing tensor for finite difference micromagnetic simulations via numerical integration," *Journal of Magnetism and Magnetic Materials*, vol. 381, pp. 440–445 (2015)

3. D. Chernyshenko and H. Fangohr, "Computation of the magnetostatic interaction between linearly magnetized polyhedrons.,' in preparation for submission (2015)

4. A. Baker, S. Zhang, W. Wang, G. Ashton, M. Beg, M. Albert, D. Chernyshenko, M.-A. Bisotti, M. Franchin, C. L. Hu, R. Stamps, T. Hesjedal, and H. Fangohr, "Proposal of a micromagnetic standard problem for ferromagnetic resonance simulations," in preparation for submission (2015)

5. W. Wang, M. Albert, M. Beg, M.-A. Bisotti, D. Chernyshenko, D. Cortes, I. Hawke, and H. Fangohr, "Magnon driven domain wall motion with the Dzyaloshinskii-Moriya interaction," *Physical Review Letters* 114, 087203 (2015)

6. W. Wang, M. Dvornik, M.-A. Bisotti, D. Chernyshenko, M. Beg, M. Albert, A. Vansteenkiste, B.V. Waeyenberge, A.N. Kuchko, V.V. Kruglyak, and H. Fangohr, "Phenomenological description of the nonlocal magnetization relaxation in magnonics, spintronics, and domain-wall dynamics," *Physical Review B* 92, 054430 (2015)

7. M. Beg, D. Chernyshenko, M.-A. Bisotti, W. Wang, M. Albert, R.L. Stamps, and H. Fangohr "Finite size effects, stability, hysteretic behaviour, and reversal mechanism of skyrmionic textures in nanostructures," in preparation for submission (2015)

Chapter 1

Introduction

Technological progress in recent decades has been driven in part by the growth of computing power. The first electronic general-purpose computer ENIAC, dating back to 1946, weighted 30 tons and consumed $150\,\mathrm{kW}$ of power to perform 5000 operations per second. The invention of the transistor in 1947 and the integrated circuit in 1958 led to a rapid increase in computational capabilities in the following years. Transistor Computer, the first transistorised computer in the world, was developed at the University of Manchester in 1953; it still relied on vacuum tubes to power its $125\,\mathrm{kHz}$ clock and was notoriously unreliable. However, within only a few years transistor-based computers fully took over — this was the start of the semiconductor industry.

In 1965 Gordon Moore, the co-founder of Intel, made the prediction that the density of transistors in an integrated circuit would double every year; later in 1975, he revised this to a doubling every two years. This famous prediction led the semiconductor industry on an impressive multi-decade journey through multiple technology and process shifts. Currently, the world's fastest supercomputer is Tianhe-2, containing 3 million CPU cores and capable of theoretical peak performance of $55\,\mathrm{petaFLOP/s}$. The next supercomputing milestone, an exascale system capable of performing 10^{18} operations per seconds, is predicted to be reached by 2018.

The demand for ever faster supercomputers is driven by the increasing demands of computational modelling. Computational modelling is a vital part of modern science, as it can support theory by providing a cheap way to experiment, and it can support experiments by explaining observed phenomena when theory alone cannot do it.

One computational model that has benefited from the growth of available computing power is the model of micromagnetism. The micromagnetic model describes the nanoscale dynamics of magnetisation in ferromagnetic materials; it was developed in the 1950s by Brown [1] but was limited in applications until the late 1980s, when computers became powerful enough to model micromagnetic systems.

The broad theme of this thesis is the development and application of new computational methods in micromagnetism. This was motivated in part by the following broad challenges facing micromagnetism today: the development of efficient algorithms aligned with the capabilities of modern computers, engineering better and more reliable software to deal with the increasing complexity of models, and the integration of new types of physics to allow multi-physics, multi-scale modelling.

It is not possible to address all of these challenges in a single work, in fact nor by a single person or group. But hopefully, some of the methods and software developed as part of this research will be a contribution towards that.

1.1 Thesis outline

The thesis is structured in the following way: Chapter 2 provides an overview of the micromagnetic model and some of the current state-of-the-art methods in micromagnetism. The continuous micromagnetic model is described in Section 2.1. Section 2.2 contains an overview of the main approaches to the discretisation in the model: the finite difference and finite element methods. Section 2.3 covers the computation of the demagnetising field — the most time-consuming part of the simulation both for finite difference and for finite element methods. Section 2.4 covers issues related to the time integration of the dynamical equation. Section 2.5 describes the computation of ferromagnetic resonance in micromagnetic systems via the solution of the eigenvalue problem for the linearized dynamical equation.

Chapter 3 describes a novel method to compute the demagnetising tensor in finite difference micromagnetics via numerical integration. This work was published in [2].

Chapter 4 describes a novel semi-analytical method for the computation of the energy of the magnetostatic interaction between linearly magnetised polyhedrons. This work is currently being prepared for submission to a journal.

Chapter 5 describes a proposal for a standard problem for modelling ferromagnetic resonance in micromagnetic systems. It is based on the collaborative paper [3], currently in preparation for submission to a journal. The author's contribution is the implementation and application of the eigenvalue method for computing the resonance modes, frequencies, and spectrum.

Chapter 6 studies the effect of Joule heating in ferromagnetic nanowires used in typical spin torque transfer experiments. This work was published in [4]. The author's contributions are the simulations done using ANSYS and the analytical estimates (i.e. Section 6.3.3 and onwards).

In addition to the novel research in Chapters 3-6, the use of the linear least-squares method to compute the FMR spectrum in Section 2.5.1.4 is also new (to the best of the author's knowledge).

1.2 Software

To support the research described in this thesis, two software packages have been developed:

- the finite difference package *simlib*, used for the simulations in Chapters 3, 4 and for the eigenvalue method in Chapter 5. The author is the sole developer of *simlib*.
- the finite element package *finmag*, used for the finite element simulation in Chapter 5. The author participated in the initial implementation for some parts of the package (in particular, those shared between the FD and FE approaches such as time integration, as well as parts of the implementation of the FEM/BEM method), however *finmag* is a collaborative effort with many developers (see Acknowledgements).

Chapter 2

Background

2.1 Micromagnetics

The micromagnetic model describes the dynamics of magnetisation in a solid ferromagnetic material at length scales typically measured in tens of nanometres to micrometres, and time scales typically measured in nanoseconds. It is a phenomenological model which ignores the atomic structure and describes magnetisation as a continuous field. In the model, interactions between the magnetic moments in the material give rise to an an effective magnetic field which in turn produces the torque that drives the magnetisation. The theory was started with the work of Landau and Lifshitz on the structure of domain walls [5] and was significantly extended by Brown [1]. The time-dependent equation was later augmented by Gilbert [6] to more accurately describe the behaviour of magnetisation in the presence of large damping.

In the micromagnetic model, the magnetisation $\mathbf{M}$ is a continuous field of constant magnitude $M = M_s$. Intuitively it can be understood as the local density of the magnetic moments of atoms in the magnetic material. The direction of magnetisation evolves according to the Landau-Lifshitz-Gilbert equation

$$\frac{d\mathbf{M}}{dt} = -\frac{\gamma}{1+\alpha^2}\left[\mathbf{M} \times \mathbf{H}_{\text{eff}} + \frac{\alpha}{M_s}\mathbf{M} \times (\mathbf{M} \times \mathbf{H}_{\text{eff}})\right] \tag{2.1}$$

In this equation, $\mathbf{H}_{\text{eff}}$ is the effective magnetic field (described below), γ is the gyromagnetic ratio and α is the damping coefficient that drives the magnetisation towards the static equilibrium value.

On its own, the Landau-Lifshitz-Gilbert equation is not sufficient to describe the dynamics of magnetisation — the complete model has to also specify the effective field $\mathbf{H}_{\text{eff}}$. The effective field represents the total contribution of physical interactions between the magnetic moments in the material as well as the external environment and contains several terms:

$$\mathbf{H}_{\text{eff}} = \mathbf{H}_{\text{applied}} + \mathbf{H}_{\text{exchange}} + \mathbf{H}_{\text{demag}} + \mathbf{H}_{\text{anisotropy}} + \dots \tag{2.2}$$

where

- $\mathbf{H}_{\text{applied}}$ is the externally applied (Zeeman) field. The magnetisation will attempt to align along it to the extent permitted by the other components of the effective field.

- $\mathbf{H}_{\text{exchange}}$ is the phenomenological field that describes the exchange interaction between neighbouring spins. The exchange interaction attempts to align neighbouring spins; in a typical material, this interaction is the strongest on short scales and neighbouring spins are typically aligned close to each other. This provides some justification for treating magnetisation as a continuous field.

- $\mathbf{H}_{\text{demag}}$ is the demagnetising field due to the dipole-dipole interaction between the individual spins. It is the combined magnetic field produced by all magnetic moments in the material.

- $\mathbf{H}_{\text{anisotropy}}$ is the phenomenological anisotropy field that attempts to align the magnetisation in the directions preferred by the crystal lattice.

In micromagnetics, the exchange and anisotropy fields arise as a consequence of corresponding energy contributions. Similarly to (2.2), the total energy E of the ferromagnetic body is

$$E = E_{\text{applied}} + E_{\text{exchange}} + E_{\text{demag}} + E_{\text{anisotropy}} + \dots \tag{2.3}$$

where

$$E_{\text{applied}} = -\mu_0 \int_\Omega \mathbf{M} \cdot \mathbf{H}_{\text{applied}} \, dV \tag{2.4}$$

$$E_{\text{exchange}} = \frac{A}{M_s^2} \int_\Omega (\boldsymbol{\nabla} M_x)^2 + (\boldsymbol{\nabla} M_y)^2 + (\boldsymbol{\nabla} M_z)^2 \, dV \tag{2.5}$$

$$E_{\text{demag}} = -\frac{\mu_0}{2} \int_\Omega \mathbf{M} \cdot \mathbf{H}_{\text{demag}} \, dV \tag{2.6}$$

$$E_{\text{anisotropy}} = \int_\Omega \varepsilon_{\text{anisotropy}}\left(\frac{\mathbf{M}}{M_s}\right) dV \tag{2.7}$$

Here, $\varepsilon(\vec{s})$ is the anisotropy energy density; it is assumed to be a function of the magnetisation direction $\mathbf{M}/M_s$ and not an arbitrary scalar field.

2.1.1 Brown's equations

In an equilibrium magnetic configuration, the internal energy is at a local minimum. To derive the condition on $\mathbf{M}$ necessary for equilibrium, the internal energy E has to be minimised under the constraint $M = M_s$. If the internal energy is described by a volume integral $E = \int \epsilon \, dV$, the minimisation results in what is known as Brown's equations

$$\mathbf{M} \times \boldsymbol{\delta}\epsilon = 0 \tag{2.8}$$

where $\boldsymbol{\delta}$ is the Euler-Lagrange operator (2.14). A detailed derivation of Brown's equation is reproduced below in section 2.1.2.

Brown's equations state that the torque produced by the field $\mathbf{H}_{\text{eff}} = -\frac{1}{\mu_0}\boldsymbol{\delta}\epsilon$ must vanish at equilibrium. Micromagnetic theory generalises the above condition and postulates that even outside equilibrium, magnetisation evolves in time according to the Landau-Lifshitz-Gilbert equation (2.1) with the effective field $\mathbf{H}_{\text{eff}}$.

For completeness, we provide the following expressions for the exchange and anisotropy fields which are easily derived from (2.14):

$$\mathbf{H}_{\text{exchange}} = -\frac{1}{\mu_0}\boldsymbol{\delta} E_{\text{exchange}} = \frac{2A}{\mu_0 M_s^2}\nabla^2\mathbf{M} \tag{2.9}$$

$$\mathbf{H}_{\text{anisotropy}} = -\frac{1}{\mu_0}\boldsymbol{\delta} E_{\text{anisotropy}} = -\frac{1}{\mu_0 M_s}\boldsymbol{\nabla}_{\mathbf{s}}\,\varepsilon_{\text{anisotropy}} \tag{2.10}$$

2.1.2 Derivation of Brown's equations

Let us assume that the total energy E of a ferromagnetic body Ω is described in terms of a volumetric energy density ϵ, so that $E = \int_\Omega \epsilon \, dV$. In the general case, the total energy may also contain surface energy terms, but the following derivation does not take them into account.

Let us denote by M_i ($i = 1, 2, 3$) the components of the magnetisation vector $\mathbf{M}$, and by M_{i,x_j} their partial derivatives $M_{i,x_j} = \frac{\partial M_i}{\partial x_j}$. In the general case, the energy density ϵ depends on the magnetisation M_i itself (in the applied, demagnetisation, and anisotropy energy terms) and on the partial derivatives M_{i,x_j} (in the exchange energy term). To find the equilibrium condition, we have to look for the stationary points of the functional

$$E = \int_\Omega \epsilon(M_i, M_{i,x_j})\, dV \tag{2.11}$$

subject to the constraint $|\mathbf{M}| = M_s$. We can eliminate the constraint by introducing a Lagrangian multiplier $\lambda = \lambda(x_i)$ and seeking the stationary point of the modified functional

$$E' = \int_\Omega \epsilon(M_i, M_{i,x_j}) + \lambda(|\mathbf{M}|^2 - M_s^2)\, dV = \int_\Omega \epsilon'(M_i, M_{i,x_j}, \lambda)\, dV \tag{2.12}$$

The stationary points can be found by solving the 3 Euler-Lagrange equations

$$\Big(\frac{\partial}{\partial M_i} - \sum_j \frac{d}{dx_j}\frac{\partial}{\partial M_{i,x_j}}\Big)\epsilon' = 0 \tag{2.13}$$

(the corresponding equation for λ produces the constraint $M^2 - M_s^2 = 0$). To simplify the notation, we introduce the Euler-Lagrange operator $\boldsymbol{\delta}$ with the components

$$\delta_i = \frac{\partial}{\partial M_i} - \sum_j \frac{d}{dx_j}\frac{\partial}{\partial M_{i,x_j}} \tag{2.14}$$

and rewrite the Euler-Lagrange equations as

$$\boldsymbol{\delta}\epsilon' = 0 \tag{2.15}$$

$$\epsilon' = \epsilon + \lambda(|\mathbf{M}|^2 - M_s^2) \tag{2.16}$$

Since $\boldsymbol{\delta}\lambda(|\mathbf{M}|^2 - M_s^2) = 2\lambda\mathbf{M}$, we can eliminate λ by multiplying the equation (2.15) by $\mathbf{M}\times$, which produces the desired equilibrium equation

$$\mathbf{M} \times \boldsymbol{\delta}\epsilon = 0 \tag{2.17}$$

2.2 Micromagnetic simulation

The micromagnetic equations cannot be solved analytically except in very simple cases [7], therefore it is necessary to solve them numerically in a computer simulation. To do so, the continuous equation (2.1) needs to be replaced with a discretised version, with each of the variables (and operators) entering the equation discretised appropriately.

Perhaps the most common approach is to perform the spatial discretisation first and approximate the continuous field $\mathbf{M}$ using a fixed number of degrees of freedom $a_1, \ldots, a_n$:

$$\mathbf{M}(\mathbf{r}) = \mathbf{M}_{\text{approx}}(a_1, \ldots, a_n)(\mathbf{r}) \tag{2.18}$$

Here, $\mathbf{M}_{\text{approx}}(a_1, \ldots, a_n)(\mathbf{r})$ is the continuous field that corresponds to the specific values of $a_1, \ldots, a_n$. The exact meaning of $a_1, \ldots, a_n$ and $\mathbf{M}_{\text{approx}}$ depends on the spatial discretisation method. The two main approaches to spatial discretization are Finite Differences (FD) and Finite Elements (FE), which are described in sections 2.2.1 and 2.2.2.

After spatial discretisation, the continuous partial differential equation (2.1) is replaced with a system of first-order ordinary differential equations with n unknowns $a_1(t), \ldots, a_n(t)$ that can subsequently be solved using a variety of standard time integration algorithms (Runge-Kutta, multistep, BDF).

2.2.1 Finite Differences

In the finite difference (FD) method, the simulated ferromagnetic system is divided into an regular array (mesh) of $n_1 \times n_2 \times n_3$ cuboids $\sigma_{i_1 i_2 i_3}$ with $i_1 = 1 \ldots n_1$, $i_2 = 1 \ldots n_2$, $i_3 = 1 \ldots n_3$. Inside each cuboid, magnetisation is assumed to have a constant value $\mathbf{M}_{i_1 i_2 i_3}$ resulting in $3n_1 n_2 n_3$ total degrees of freedom.

If the magnetic region has an irregular shape T, magnetisation outside of that region is pinned to zero: $\mathbf{M}_{i_1 i_2 i_3} = 0$ if $\sigma_{i_1 i_2 i_3}$ is outside T.

In the finite difference method a common way to discretise the effective field and the Landau-Lifshitz-Gilbert equation is via the so-called energy-based approach [8].

In this method, the total energy (2.3) is written as a function of the degrees of freedom $\mathbf{M}_{i_1 i_2 i_3}$

$$E(\mathbf{M}_{i_1 i_2 i_3}) = E_{\text{applied}}(\mathbf{M}_{i_1 i_2 i_3}) + E_{\text{exchange}}(\mathbf{M}_{i_1 i_2 i_3}) + \ldots \tag{2.19}$$

Similarly to the procedure in section 2.1.2, this function is then minimised with respect to the degrees of freedom $\mathbf{M}_{i_1 i_2 i_3}$ (subject to the constraint $|\mathbf{M}| = 1$). This produces a discrete effective field

$$\mathbf{H}_{\text{eff},i_1 i_2 i_3} = -\frac{1}{\mu_0 |\sigma|} \frac{\partial E}{\partial \mathbf{M}_{i_1 i_2 i_3}} \tag{2.20}$$

where $|\sigma|$ is the volume of the cuboidal cells in the mesh.

The time derivatives of the magnetisation $d\mathbf{M}_{i_1 i_2 i_3}/dt$ can then be computed pointwise:

$$\frac{d\mathbf{M}_{i_1 i_2 i_3}}{dt} = -\frac{\gamma}{1+\alpha^2}\left[\mathbf{M}_{i_1 i_2 i_3} \times \mathbf{H}_{\text{eff},i_1 i_2 i_3} + \frac{\alpha}{M_s}\mathbf{M}_{i_1 i_2 i_3} \times (\mathbf{M}_{i_1 i_2 i_3} \times \mathbf{H}_{\text{eff},i_1 i_2 i_3})\right] \tag{2.21}$$

This energy-based approach to the discretisation of the Landau-Lifshitz-Gilbert equation guarantees that up to the tolerance of the ODE solver, the discrete total energy is minimised at an equilibrium point, and in general, in any simulation the total energy will be a non-increasing function of time.

2.2.2 Finite Elements

In the finite element method, the simulated magnetic region of space is divided into a polyhedral mesh [9]. The continous magnetisation field $\mathbf{M}(r)$ is approximated by a linear combination of basis functions ϕ_i associated with that mesh:

$$\mathbf{M}_{\mathrm{approx}}(\mathrm{r}) = \sum \mathbf{M}_i \phi_i(\mathbf{r}) \tag{2.22}$$

where $\mathbf{M}_i$ are the corresponding degrees of freedom.

In the most common case in micromagnetic simulations, the functions $M_{\mathrm{approx}}(\mathbf{r})$ are piecewise polynomials, typically of the first order (i.e. linear). The micromagnetic package *finmag* used for the simulations in Chapter 5 employs tetrahedral meshes and Lagrange elements where the approximation function space $\mathbf{M}_{\mathrm{approx}}(\mathbf{r})$ is the set of all continuous (vector-valued) functions on the mesh, linear within each tetrahedron. In this case, each node i in the mesh has a corresponding degree of freedom $\mathbf{M}_i$ and the basis functions $\phi_i(\mathbf{r})$ are piecewise linear "tent" functions that have value 1 at node i and value 0 at all other nodes.

2.2.3 Effective field and the box method for finite elements

To discretise the effective field in the finite element approach, similarly to the finite difference case we again write the total energy functional

$$E(\mathbf{M}_i) = E_{\mathrm{applied}}(\mathbf{M}_i) + E_{\mathrm{exchange}}(\mathbf{M}_i) + E_{\mathrm{demag}}(\mathbf{M}_i) + E_{\mathrm{anisotropy}}(\mathbf{M}_i) + \dots \tag{2.23}$$

Similar to the finite difference case we can minimise the total energy with respect to $\mathbf{M}_i$ to derive the equilibrium equation and the effective field. The FENICS finite element library used in *finmag* can perform this derivation automatically for the applied, exchange, and anisotropy energy terms where the energy functional can be easily represented as a volume integral over the mesh.

A slight disadvantage of this approach is that the computation of the applied, exchange, and anisotropy fields involves solving a sparse linear system which requires computational effort. Additionally, for most meshes the value of the field at each node will potentially be influenced by magnetisation at all other nodes, making the discretised applied, exchange, and anisotropy interactions nonlocal.

To avoid this problem, the effective field can instead be calculated using the so-called "box method" [9].

Recall that in the continuous case, the effective field is the variational derivative of the energy functional with respect to $\mathbf{M}(\mathbf{r})$:

$$\mathbf{H}_{\mathrm{eff}} = -\frac{1}{\mu_0}\frac{\delta U}{\delta \mathbf{M}(\mathbf{r})} \tag{2.24}$$

Our previous approach was to write the total energy as a function of the degrees of freedom $\mathbf{M}_i$, and solve the minimisation problem afterwards. Alternatively, we can directly define the discretisation of the operator $\boldsymbol{\delta}$ to be

$$\frac{\delta_{\text{discrete}} U}{\delta \mathbf{M}_i} := \frac{1}{V_i} \frac{\partial U}{\partial \mathbf{M}_i} \tag{2.25}$$

where

$$V_i := \int \phi_i(r) dr \tag{2.26}$$

is the volume of the portion of the mesh associated with i-th node. The effective field can thus be computed as

$$\mathbf{H}_i = -\frac{1}{\mu_0 V_i} \frac{\partial U}{\partial \mathbf{M}_i} \tag{2.27}$$

Since a change of magnetisation at the node $\mathbf{M}_i$ only influences the energy density in the neighbouring elements, changing $\mathbf{M}_i$ will only alter the applied, exchange, and anisotropy fields corresponding to node i and directly neighbouring nodes, giving us the desired locality property.

However, since the effective field $\mathbf{H}_i$ is no longer obtained by energy minimisation, the total energy (2.23) may increase during simulation, and will generaly be not at a minimum at an equilibrium configuration. To avoid this, the discrete energy function can be modified so that the energy is computed from the discrete effective field (2.27). During a simulation, this modified energy function will then be non-increasing (again, subject to the accuracy of the time integrator).

2.3 Demagnetising field

Both for finite difference (FE) and finite element (FD) micromagnetic simulations, the most time-consuming part is the computation of the demagnetising field. The time derivative of magnetisation (the Landau-Lifshitz-Gilbert equation) is usually computed pointwise; the exchange and anisotropy fields are local — their value only depends on the value of the magnetisation in close proximity. These terms can be computed efficiently and the computation can be made parallel to utilise multiple CPU cores.

On the other hand, the demagnetising field is global — the magnetostatic potential and field at each point is the sum of the contributions from all discretisation cells. For a system of n cells, the naive approach to computing the demagnetising field would require the evaluation of $O(n^2)$ interactions, which is infeasible for a system of any practical size. Therefore, both for FD and FE methods the computation has to utilise specialized, more efficient algorithms.

2.3.1 Demagnetising field: overview

The demagnetising field $\mathbf{H}_{\text{demag}}$ is the classical magnetostatic field produced by all the magnetic moments in the system. For a continuous system it is described by Maxwell's equations, however in micromagnetics the standard equations are usually transformed into a form more suitable for computation.

In micromagnetics, the demagnetising field $\mathbf{H}_{\text{demag}}$ is considered separately from the magnetic field originating outside of the ferromagnet, described by the applied (Zeeman) field $\mathbf{H}_{\text{applied}}$. For

the rest of Section 2.3, $\mathbf{H}$ will denote the demagnetising field only, and E the demagnetising energy (i.e. $\mathbf{H} = \mathbf{H}_{\text{demag}}$, $E = E_{\text{demag}}$).

For the demagnetising field,

$$B = \mu_0(\mathbf{H} + \mathbf{M}) \qquad \text{constitutive relation} \tag{2.28}$$

$$\nabla \cdot \mathbf{B} = 0 \qquad \text{Gauss's law} \tag{2.29}$$

$$\nabla \times \mathbf{H} = 0 \qquad \text{Ampere's law} \tag{2.30}$$

From the last equation, there exists a scalar potential function $\phi(\mathbf{r})$ defined on the whole space $\mathbb{R}^3$ such that

$$\mathbf{H} = -\nabla\phi \tag{2.31}$$

Then, from the first two equations

$$\nabla^2\phi = \begin{cases} \nabla \cdot \mathbf{M} & \text{inside the material} \\ 0 & \text{outside the material} \end{cases} \tag{2.32}$$

At the boundary of the material, the normal component of $\mathbf{B}$ is continuous and so is the tangential component of H. Therefore, the gradient of ϕ is continuous in the tangential direction and has a discontinuity of $-\mathbf{n} \cdot \mathbf{M}$ in the normal direction, where $\mathbf{n}$ is the normal vector to the boundary:

$$\nabla\phi_{\text{outside}} - \nabla\phi_{\text{inside}} = -(\mathbf{n} \cdot \mathbf{M})\,\mathbf{n} \qquad \text{at the boundary} \tag{2.33}$$

With an additional regularity condition, these equations allow a unique solution [7, 1]. Therefore the computation of the demagnetising field can be reduced to the problem of finding the magnetostatic potential ϕ that satisfies the equations (2.32), (2.33). Once ϕ is known, the demagnetising field can be computed using equation (2.31), and the demagnetising energy can be found using the relation

$$E = -\frac{\mu_0}{2}\int_V \mathbf{M} \cdot \mathbf{H}\, dr \tag{2.34}$$

where the integration is over the ferromagnetic region V.

The equations (2.32), (2.33) permit the formal solution

$$\phi(\mathbf{r}) = -\frac{1}{4\pi}\int_V \nabla' \cdot \mathbf{M}(\mathbf{r}')\frac{1}{|\mathbf{r} - \mathbf{r}'|}\, dr' + \frac{1}{4\pi}\int_{\partial V} \mathbf{n} \cdot \mathbf{M}(\mathbf{r}')\frac{1}{|\mathbf{r} - \mathbf{r}'|}\, ds' \tag{2.35}$$

$$\mathbf{H}(\mathbf{r}) = -\frac{1}{4\pi}\int_V \nabla' \cdot \mathbf{M}(\mathbf{r}')\,\nabla'\frac{1}{|\mathbf{r} - \mathbf{r}'|}\, dr' + \frac{1}{4\pi}\int_{\partial V} \mathbf{n} \cdot \mathbf{M}(\mathbf{r}')\,\nabla'\frac{1}{|\mathbf{r} - \mathbf{r}'|}\, ds' \tag{2.36}$$

or alternatively, integrating by parts,

$$\phi(\mathbf{r}) = \frac{1}{4\pi}\int_V \mathbf{M}(\mathbf{r}') \cdot \nabla'\frac{1}{|\mathbf{r} - \mathbf{r}'|}\, dr' \tag{2.37}$$

$$\mathbf{H}(\mathbf{r}) = -\frac{1}{4\pi}\int_V \mathbf{M}(\mathbf{r}') \cdot \nabla\nabla'\frac{1}{|\mathbf{r} - \mathbf{r}'|}\, dr' \tag{2.38}$$

The first pair of equation can be interpreted as the potential $\phi(\mathbf{r})$ and field $\mathbf{H}(\mathbf{r})$ arising due to volume charge density $-\nabla \cdot \mathbf{M}$ inside the material plus surface charge density $\mathbf{n} \cdot \mathbf{M}$ on the boundary of the material. The second pair of equations can be interpreted as the potential $\phi(\mathbf{r})$

and field $\mathbf{H}(\mathbf{r})$ being the superposition of potentials and fields of primitive magnetic dipoles with dipole moment $\mathbf{M}(\mathbf{r})$ throughout the magnetic region.

Computational procedures usually either rely on the differential equation formulation (2.32), (2.33) to compute the potential $\phi(\mathbf{r})$ — as done in the FEM/BEM method described in Section 2.3.3, or use the integral formulation (2.38) to compute the energy and field directly — as done in the Fourier transform finite difference method described in Section 2.3.2 as well as the analytical calculation in Chapter 4.

2.3.2 Demagnetising field: finite difference method

In the finite difference method, space is discretized into a regular array of cuboids. Suppose that the magnetic region V is divided into an $n = n_1 \times n_2 \times n_3$ grid of cuboidal cells $\sigma_{i_1 i_2 i_3}$ centered at $\mathbf{r}_{i_1 i_2 i_3}$ each with dimensions $c_x \times c_y \times c_z$, volume $|\sigma| = c_x c_y c_z$, and constant magnetisation $\mathbf{M}_{i_1 i_2 i_3}$. From equations (2.38) and (2.34) the total demagnetising energy E of the system is

$$\begin{aligned} E &= \frac{\mu_0}{8\pi} \int_V \int_V \mathbf{M}(\mathbf{r}) \cdot \nabla\nabla' \frac{1}{|\mathbf{r} - \mathbf{r}'|} \cdot \mathbf{M}(\mathbf{r}')\, dr\, dr' \\ &= \frac{1}{2} \sum_{i_1, i_2, i_3} \sum_{i'_1, i'_2, i'_3} E_{i_1 i_2 i_3 \leftrightarrow i'_1 i'_2 i'_3} \end{aligned} \tag{2.39}$$

where

$$E_{i_1 i_2 i_3 \leftrightarrow i'_1 i'_2 i'_3} = \frac{\mu_0}{4\pi} \int_{\sigma_{i_1 i_2 i_3}} \int_{\sigma_{i'_1 i'_2 i'_3}} \mathbf{M}(\mathbf{r}) \cdot \nabla\nabla' \frac{1}{|\mathbf{r} - \mathbf{r}'|} \cdot \mathbf{M}(\mathbf{r}')\, dr\, dr' \tag{2.40}$$

is the energy of the pairwise interaction between the cells $\sigma_{i_1 i_2 i_3}$ and $\sigma_{i'_1 i'_2 i'_3}$.

In the finite difference method, magnetisation is assumed to be constant within each cell. Therefore, the factors $\mathbf{M}(\mathbf{r})$ and $\mathbf{M}(\mathbf{r}')$ can be taken outside of the double integral. The remaining double integral only depends on the separation between the cells $\sigma_{i_1 i_2 i_3}$ and $\sigma_{i'_1 i'_2 i'_3}$, leading to the following formula

$$E_{i_1 i_2 i_3 \leftrightarrow i'_1 i'_2 i'_3} = \mu_0 |\sigma|\, \mathbf{M}_{i_1 i_2 i_3} \cdot \mathbf{N}(\mathbf{r}_{i_1 i_2 i_3} - \mathbf{r}_{i'_1 i'_2 i'_3}) \cdot \mathbf{M}_{i'_1 i'_2 i'_3} \tag{2.41}$$

$$\mathbf{N}(\mathbf{r}) := \frac{1}{4\pi|\sigma|} \int_\sigma \int_\sigma \nabla_{r_1} \nabla_{r_2} \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|}\, dr_1\, dr_2 \tag{2.42}$$

Combining the equations (2.41), (2.39), and (2.20) also gives us the formula for the discrete demagnetising field (in the energy-based approach):

$$\mathbf{H}_{i_1 i_2 i_3} = -\sum_{i'_1 i'_2 i'_3} \mathbf{M}_{i'_1 i'_2 i'_3} \cdot \mathbf{N}(\mathbf{r}_{i_1 i_2 i_3} - \mathbf{r}_{i'_1 i'_2 i'_3}) \tag{2.43}$$

In the previous equations, we have introduced the so-called demagnetising tensor $\mathbf{N}(\mathbf{r})$. It is a dimensionless symmetric rank 2 tensor which describes the magnetostatic interaction between cells in the mesh. It does not depend on the magnetisation of the cells or their exact position, only on the separation vector $\mathbf{r} = \mathbf{r}_{i_1 i_2 i_3} - \mathbf{r}_{i'_1 i'_2 i'_3}$ and the dimensions of the cell c_x, c_y, c_z. An analytical formula for $\mathbf{N}(\mathbf{r})$ is known [10, 11, 12, 13]; more detail about the demagnetising tensor and its computation is provided in Chapter 3.

In principle, the analytical formula allows us to compute the demagnetising tensor $\mathbf{N}(\mathbf{r}_{i_1 i_2 i_3} - \mathbf{r}_{i'_1 i'_2 i'_3})$ and energy $E_{i_1 i_2 i_3 \leftrightarrow i'_1 i'_2 i'_3}$ for each pair of cells in the mesh, and add them together to compute the demagnetising energy E and field $\mathbf{H}_{i_1 i_2 i_3}$. However, this method requires $O(n^2)$ operations and for any practical mesh size n this computation is not feasible.

To arrive at a more efficient algorithm, notice that since the mesh is regular, the separation vector $\mathbf{r}_{i_1 i_2 i_3} - \mathbf{r}_{i'_1 i'_2 i'_3} = (c_x(i_1 - i'_1), c_y(i_2 - i'_2), c_z(i_3 - i'_3))$ only depends on the difference between the indices of the separated cells $i_1 - i'_1$, $i_2 - i'_2$, $i_3 - i'_3$. Therefore, we can introduce the integer-indexed demagnetising tensor

$$\mathbf{N}_{j_1, j_2, j_3} := \mathbf{N}(c_x j_1, c_y j_2, c_z j_3) \tag{2.44}$$

and rewrite (2.43) as

$$\mathbf{H}_{i_1 i_2 i_3} = -\sum_{i'_1 i'_2 i'_3} \mathbf{M}_{i'_1 i'_2 i'_3} \cdot \mathbf{N}_{i_1 - i'_1, i_2 - i'_2, i_3 - i'_3} \tag{2.45}$$

Although this is still a double sum involving a total of $O(n^2)$ operations, we can recognize a discrete convolution, which can be computed efficiently using the discrete fast Fourier transform (DFFT) algorithm. The only potential difficulty is that the discrete Fourier transform operates on circular sums and convolutions, while for the standard (non-periodic) boundary conditions the above sum is non-periodic. To address this issue, we can pad the array $\mathbf{M}_{i_1 i_2 i_3}$ with zeroes to size $2n_x \times 2n_y \times 2n_z$ and extend the sum in (2.45) to the new doubled indices in all dimensions. This produces a circular sum, to which the discrete Fourier transform can then be applied.

The full algorithm for the computation of the finite difference demagnetising field is therefore:

1. Compute the demagnetising tensor entries $\mathbf{N}_{i_1 - i'_1, i_2 - i'_2, i_3 - i'_3}$ and store them in an $2n_x \times 2n_y \times 2n_z \times 6$ array (with a minus sign for equation (2.45), as well as any dimensional factors as required).

2. Compute the 3d forward Fourier transform of this array, producing a complex-valued $2n_x \times 2n_y \times 2n_z \times 6$ array $\mathbf{N}^*$ in the frequency space.

3. Given a discretized magnetisation array $\mathbf{M}_{i_1 i_2 i_3}$, pad it to size $2n_x \times 2n_y \times 2n_z \times 3$ with zeroes.

4. Compute the forward Fourier transform of the magnetisation, producing a $2n_x \times 2n_y \times 2n_z \times 3$ array $\mathbf{M}^*$ in the frequency space.

5. Compute the elementwise dot product of $\mathbf{M}^*$ and $\mathbf{N}^*$, resulting in the array $\mathbf{H}^*$ in the frequency space.

6. Compute the backward Fourier transform of $\mathbf{H}^*$, producing the demagnetising field $\mathbf{H}_{i_1 i_2 i_3}$ in real space.

In step 1, the computation of the demagnetising tensor $N_{i_1 - i'_1, i_2 - i'_2, i_3 - i'_3}$ can be done using the analytical formula when the separation between the cells is small. When the separating distance is

large, the evaluation of the analytical formula becomes inaccurate due to numerical cancellation errors [14]. Due to this, the micromagnetic package OOMMF [15] uses an asymptotic formula to compute the demagnetising tensor for large separating distances. In Chapter 3, we develop a novel approach to address the numerical cancellation issue based on numerical integration of the equation (2.42).

The method described in this section was used for the finite difference simulation code *simlib* used in Chapters 3, 4, and for the eigenvalue method in 5.

2.3.3 Demagnetising field: finite element method

In the finite element method, the mesh is irregular and the Fourier transform method described in the previous chapter cannot be applied. Instead, a solution to the open boundary Poisson equation (2.32) for the magnetostatic potential can be sought. A direct approach would require the discretization of the space outside of the material in order to satisfy the regularity condition at infinity. This may substantially increase the size of the mesh as well as make the meshing procedure more complicated. A review of specialized methods for the open boundary problem can be found in [16] and [17].

An alternative method that only requires the discretisation of space inside the material is the hybrid FEM/BEM method [18]. In this method, the magnetostatic potential ϕ is split into two parts

$$\phi = \phi_1 + \phi_2 \tag{2.46}$$

Inside the material, the potential ϕ_1 is required to satisfy the Poisson equation with a Neumann boundary condition, outside, it is 0:

$$\begin{aligned} \nabla^2\phi_1 &= \nabla\cdot\mathbf{M} && \text{inside the material} \\ \phi_1 &= 0 && \text{outside the material} \\ \frac{\partial\phi_1}{\partial\mathbf{n}} &= \mathbf{n}\cdot\mathbf{M} && \text{at the boundary} \end{aligned} \tag{2.47}$$

For the equations (2.32) and (2.33) to hold, the second part of the potential ϕ_2 must therefore satisfy:

$$\nabla^2\phi_2 = 0 \qquad \text{everywhere} \tag{2.48}$$

$$\nabla\phi_{2,\text{outside}} - \nabla\phi_{2,\text{inside}} = 0 \qquad \text{at the boundary} \tag{2.49}$$

$$\phi_{2,\text{outside}} - \phi_{2,\text{inside}} = -\phi_{1,\text{inside}} \qquad \text{at the boundary} \tag{2.50}$$

The equation for ϕ_1 is a standard Poisson equation and can be solved using the ordinary finite element method. The equations for ϕ_2 describe the scalar potential of dipolar density $\phi_{1,\text{inside}}$ at the boundary; the solution is

$$\phi_2(\mathbf{r}) = \frac{1}{4\pi}\int_{\partial V}\phi_1(\mathbf{r}')\frac{\mathbf{n}\cdot(\mathbf{r}-\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|^3}\,ds' \qquad \text{inside the material} \tag{2.51}$$

$$\phi_2(\mathbf{r}) = \frac{1}{4\pi}\int_{\partial V}\phi_1(\mathbf{r}')\frac{\mathbf{n}\cdot(\mathbf{r}-\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|^3}\,ds' + \left(\frac{\Omega_{\partial V}(\mathbf{r})}{4\pi} - 1\right)\phi_1(\mathbf{r}) \qquad \text{at the boundary} \tag{2.52}$$

where $\Omega_{\partial V}(\mathbf{r})$ is the solid angle covered by the boundary ∂V as seen from the point $\mathbf{r}$.

To avoid solving for ϕ_2 on the whole mesh, instead the equation (2.52) can be used to calculate the values of ϕ_2 at the boundary, and the values of ϕ_2 inside can be found by solving the Laplace equation using FEM inside the material.

From equation (2.52), ϕ_2 is a linear function of the boundary charge ϕ_1. After finite element discretization, the degree-of-freedom vector $\hat{\phi}_{2,\text{boundary}}$ is the product of a boundary element matrix $\mathbf{B}$ with the degree-of-freedom vector $\hat{\phi}_{2,\text{boundary}}$

$$\hat{\phi}_{2,\text{boundary}} = \mathbf{B} \cdot \hat{\phi}_{1,\text{boundary}} \tag{2.53}$$

The boundary element matrix $\mathbf{B}$ is independent of the magnetization M; it can be computed using the analytical Lindholm formulas for the double layer potential [19, 20]. The complete procedure for computing the demagnetising field is therefore

1. Compute the boundary element matrix $\mathbf{B}$ (this step can be done once at the start of the simulation).
2. Solve the Poisson equation (2.47) for the potential ϕ_1 using FEM.
3. Compute the potential ϕ_2 at the boundary using equation (2.52).
4. Solve the Laplace equation for ϕ_2 inside the material based on the known values of ϕ_2 at the boundary using FEM.
5. Compute the demagnetising field $\mathbf{H} = -\nabla(\phi_1 + \phi_2)$.

An alternative hybrid FEM/BEM method has been developed which imposes slightly different boundary conditions [21]:

$$\phi = \phi_1' + \phi_2' \tag{2.54}$$

$$\nabla^2 \phi_1' = \nabla \cdot \mathbf{M} \qquad \text{inside the material} \tag{2.55}$$

$$\phi_1' = 0 \qquad \text{outside } \textit{and} \text{ at the boundary} \tag{2.56}$$

$$\nabla^2 \phi_2' = 0 \qquad \text{everywhere} \tag{2.57}$$

$$\frac{\partial \phi_{2,\text{outside}}'}{\partial \mathbf{n}} - \frac{\partial \phi_{2,\text{inside}}'}{\partial \mathbf{n}} = -\mathbf{n} \cdot \mathbf{M} + \frac{\partial \phi_1'}{\partial \mathbf{n}} \qquad \text{at the boundary} \tag{2.58}$$

$$\phi_{2,\text{outside}}' - \phi_{2,\text{inside}}' = 0 \qquad \text{at the boundary} \tag{2.59}$$

Again, the Poisson equation for ϕ_1' can be solved using FEM, while the potential ϕ_2' can be found using the single layer integral

$$\phi_2'(\mathbf{r}) = \frac{1}{4\pi} \int_{\partial V} \Big(-\mathbf{n} \cdot \mathbf{M}(\mathbf{r}') + \frac{\partial \phi_1'(\mathbf{r}')}{\partial \mathbf{n}} \Big) \frac{1}{|\mathbf{r} - \mathbf{r}'|} \, ds' \qquad \text{inside the material} \tag{2.60}$$

In both methods, the boundary element matrix $\mathbf{B}$ is dense — in a mesh with m surface nodes it has the dimensions $3m \times 3m$. Due to this, when the mesh geometry has a large number of surface nodes (for instance, when modelling a thin film material), the memory requirements to store the

matrix can be quite high. Approaches based on hierarchical matrix compression have been applied to deal with this issue.

The micromagnetic simulation package *finmag* supports both the [18] and the [21] variants of the hybrid FEM/BEM to compute the demagnetising field.

2.3.4 Fast multipole method

The fast multipole method (FMM) is a general algorithm to compute pairwise interactions in a arbitrary system of n bodies. Direct computation requires the evaluation of n^2 interactions, while the fast multipole method reduces this complexity to $O(n \log n)$ by approximating the long-range interactions between the bodies.

In micromagnetics, the fast multipole method has been successfully applied to compute the demagnetising field in the finite difference setting [22] as well as to compute interactions between nanoparticles [23, 24]. However, the application of FMM to the finite element setting is particularly desireable since the fast multipole method potentially offers two advantages against the commonly used hybrid FEM/BEM (described in Section 2.3.3): the first advantage is straightforward handling of thin geometries (as no boundary matrix is used), and the second is easy parallelisation compared to the FEM/BEM method which requires an iterative solution of the linear systems for ϕ_1 and then ϕ_2.

This section outlines the prerequisites for the computation of the demagnetising field using FMM for the finite element method. Later in the thesis, Chapter 4 addresses one of the principal prerequisites — the computation of the direct part of the fast multipole method.

The fast multipole method can be applied if

- Each body a_i in the system produces a spatial field $\Phi_i(r)$; these fields add to produce the total field $\Phi(\mathbf{r}) = \sum_i \Phi_i(\mathbf{r})$.
- The interaction/motion between the bodies is due to the interaction between the body and the total field: if the total field $\Phi(\mathbf{r})$ is known, then so is the interaction of each body with it.
- Far away from body a_i, its field $\Phi_i(\mathbf{r})$ is smooth, permitting a far-field approximation (the multipole expansion).

The success of the original fast multipole algorithm [25, 26, 27] has lead to a large number of variations and improvements, however any fast multipole algorithm has the following informal stucture:

- Iterating through the bodies, build a hierarchical structure that represents the multipole expansion of the "far away field" for *all* the bodies in the system.
- For each body, compute the interaction with bodies "close" to it directly.
- For each body, compute their interaction with the "far away field" computed in step 1. The hierarchical structure in step 1 is designed in such a way so that for each pair of interacting

bodies, their pairwise interaction is only ever considered once - either as a "close" direct interaction, or as a far away interaction with the multipole field.

This general structure of the algorithm is flexible in terms of: the kind of multipole expansion used in the far-field approximation, the concept of the body itself, the field $\Phi(\mathbf{r})$, the criteria for when bodies are "close", the hierarchical structure used in step 1, and so on.

The original FMM algorithm and early modifications used the spherical expansion for the far field. However, recently algorithms based on the Cartesian expansion [28] became popular. The Cartesian expansion algorithms perform particularly well when the multipole expansion is of relatively low order [29]. For the calculation of the demagnetising field, a high level of accurace is not required, since other parts of the simulation such as the spatial discretization introduce errors of their own. Therefore, the Cartesian expansion is a natural choice.

For the choice of the interacting objects, the two main candidates are the interactions between the polyhedrons (tetrahedrons) of the mesh (i.e., between the cells of the mesh), or between the finite element basis (tent) functions directly (i.e., between the nodes of the mesh). Since a typical mesh usually has fewer nodes than cells, using basis functions would result in fewer bodies and fewer pairs of interactions. On the other hand, it is more difficult to determine the criterion for closeness for the basis functions — a less efficient criterion would lead to a larger fraction of interactions that have to be computed directly, resulting in a loss of efficiency.

The next choice is to determine the potential/field Φ as well the base function of the multipole expansion. The natural option is to use the scalar potential $1/|\mathbf{r}|$ as the base for the multipole expansion. For a magnetic body, the first moment of the expansion (the magnetic charge) would be zero, while the second moment (the magnetic dipolar moment) is the dominant term at long range.

A concrete implementation of the fast multipole algorithm would require several equations that need to be established:

- The criteria for when the bodies are close and the interaction is evaluated directly versus via the far-field approximation.
- The multipole expansion of a body in terms of the scalar potential $1/|\mathbf{r}|$. As can be seen seen from equations (2.37) and (2.38), this would encode both the magnetostatic potential $\phi(\mathbf{r})$ and the magnetostatic field $\mathbf{H}(\mathbf{r}) = -\nabla\phi$ of the body.
- The interaction of a body with the global field: given the coefficients of a multipole expansion, how to compute degrees of freedom for the demagnetising field.
- The evaluation of a direct interaction.

For the energy-based formulation (Section 2.2.1), the items 2 and 3 are very straightforward to calculate (using equations (2.37) and (2.38)), so the main difficulty is to determine the cutoff criteria for the direct interaction, as well as to efficiently compute the direct interaction itself. In the energy-based formulation, this involves computing the magnetostatic energy of two interacting, linearly magnetised polyhedrons. A novel semi-analytical method to do this is developed in Chapter 4.

2.4 Time integration of the Landau-Lifshitz-Gilbert equation

After spatial discretization with either the finite difference or the finite element method, the dynamics of the system are described by the semi-discretized Landau-Lifshitz-Gilbert equation for n magnetisation variables $\mathbf{m}_i$:

$$\frac{d\mathbf{m}_i}{dt} = -\frac{\gamma}{1+\alpha^2}\left[\mathbf{m}_i \times \mathbf{H}_{\text{eff},i} + \alpha\, \mathbf{m}_i \times (\mathbf{m}_i \times \mathbf{H}_{\text{eff},i})\right] \qquad i = 1..n \tag{2.61}$$

This is an ordinary differential equation with $3n$ degrees of freedom, written using normalised magnetisation $\mathbf{m} = \mathbf{M}/M_s$, which is often easier to work with in code.

This equation can be integrated using standard methods for solving ODEs, however this equation is stiff and explicit methods such as Runge-Kutta or explicit multistep methods are often inefficient — implicit methods such as backwards differentiation formula (BDF) usually perform better. The stiffness of the equation is due to the considerably higher strength of the exchange interaction compared to other interactions present in the system. In principle, to resolve the high strength of the exchange, small timesteps are required during time integration. However, in a typical simulation neighbouring spins will seek to minimise the exchange energy and will remain nearly parallel throughout the whole simulation — the neighbouring spins will rotate in lockstep to ensure the exchange energy stays low, and the motion of the spins is on the timescale of the other interactions. Intuitively, explicit integration methods might potentially get "stuck" probing high-energy short-timescale motions where neighbouring spins move out of alignment, resulting in a very short timestep required to resolve those motions (or numerical instability if a longer timestep is forced).

The stiffness of equation (2.61) is a particular concern for finite element simulations. As simulation cells get smaller, the strength of the exchange field increases relative to the strength of the demagnetising field — in the FE case due to the $1/V_i$ factor in equation (2.27), and in the FD case due to the $1/|\sigma|$ factor in equation (2.20). The non-regular meshes used in the finite element method will nearly always have some smaller cells for which the stiffness problem is more pronounced.

To efficiently resolve the stiff dynamics of the system, implicit methods require the knowledge of the Jacobean of the differential equation (or rather, the inverse of it). The Jacobean is the derivative of the right-hand side of the differential equation with respect to the degree of freedom variables — in the case of the LLG, with respect to magnetisation. The derivation of the Jacobean for the Landau-Lifshitz-Gilbert is shown in Section 2.4.1, however, before this can be done, one extra challenge for the computation of the dynamical equation needs to be addressed.

This extra challenge is the requirement for the equation (2.61) to conserve the magnitude of the magnetisation, $|\mathbf{m}| = 1$. Mathematically, any solution to (2.61) conserves it, however due to the accumulation of numerical errors in a simulation the magnitude of the magnetisation eventually will diverge from 1.

To deal with this, so-called geometrical integrators have been developed that observe the constraint exactly [30], however such integrators are specific to the Landau-Lifshitz-Gilbert equation

and cannot be implemented using existing ODE solvers such as SUNDIALS [31]. Instead, to maintain the constraint we add a longitudinal relaxation term to the equation (2.61) that brings the magnetisation back to its standard magnitude:

$$\frac{d\mathbf{m}_i}{dt} = -\frac{\gamma}{1+\alpha^2}\left[\mathbf{m}_i \times \mathbf{H}_{\text{eff},i} + \alpha\, \mathbf{m}_i \times (\mathbf{m}_i \times \mathbf{H}_{\text{eff},i})\right] + \frac{1}{t_{\text{relax}}}(1 - |\mathbf{m}_i|^2)\, \mathbf{m}_i \tag{2.62}$$

Here, t_{relax} is the time constant for the longitudinal relaxation which needs to be chosen based on other parameters of the simulation. If t_{relax} is too large, magnetisation may still diverge from the magnitude of 1. If it is too small, the equation will become more stiff and the ODE solver will be forced to choose a smaller time step. A heuristic rule used in the finite difference simulation package *simlib* sets t_{relax} to 1/10th of the characteristic time of the exchange interaction:

$$t_{\text{relax}} = 0.1 \cdot t_{\text{exchange}} \tag{2.63}$$

$$t_{\text{exchange}} = 1 \Big/ \frac{\gamma\, H_{\text{exchange}}}{1+\alpha^2} \tag{2.64}$$

$$H_{\text{exchange}} = \frac{2A}{\mu_0 M_s \sigma^{2/3}} \tag{2.65}$$

where A is the exchange constant, M_s the saturation magnetisation, and σ the volume of the cells in the mesh.

2.4.1 The Jacobean of the Landau-Lifshitz-Gilbert equation

If the discrete Landau-Lifshitz-Gilbert equation (2.62) is written in the form

$$\frac{d\mathbf{m}}{dt} = \mathcal{L}(\mathbf{m}, \mathbf{H}_{\text{eff}}(\mathbf{m})) \tag{2.66}$$

then the Jacobean of the equation is the derivative

$$\mathbf{J} = \nabla_{\mathbf{m}} \mathcal{L}(\mathbf{m}, \mathbf{H}_{\text{eff}}(\mathbf{m})) \tag{2.67}$$

For a system with n degrees of freedom (in magnetisation) the Jacobean $\mathbf{J}$ is a $3n \times 3n$ matrix. It is clearly infeasible to store this matrix in memory. Instead, the implicit solver is usually provided with a routine to compute the product $\mathbf{Jv}$ of the Jacobean with an arbitrary vector $\mathbf{v}$. Using this routine, the ODE solver will then use an iterative algorithm to find the inverse products $\mathbf{J}^{-1}\mathbf{w}$ required for the implicit integration formula.

Therefore, the goal is to derive a formula for the Jacobean-times-vector product $\mathbf{Jv}$. The starting point is the identity

$$\mathbf{Jv} = \frac{d}{d\epsilon}\mathcal{L}(\mathbf{m} + \epsilon\mathbf{v}, \mathbf{H}_{\text{eff}}(\mathbf{m} + \epsilon\mathbf{v}))\Big|_{\epsilon=0} \tag{2.68}$$

Let $\mathbf{H}'$ be the derivative of the effective field with respect to magnetisation

$$\mathbf{H}' = \nabla_{\mathbf{m}} \mathbf{H}_{\text{eff}}(\mathbf{m}) \tag{2.69}$$

Similarly to the Jacobean, $\mathbf{H}'$ is a $3n \times 3n$ matrix, while $\mathbf{H}'\mathbf{v}$ is a vector of length $3n$. Using the identity $\mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}}) = \mathbf{m}(\mathbf{m} \cdot \mathbf{H}_{\text{eff}}) - \mathbf{H}_{\text{eff}}(\mathbf{m} \cdot \mathbf{m})$ and the Landau-Lifshitz-Gilbert

equation with the longitudinal relaxation term (2.62) gives the following formula for the product

$$\begin{aligned}\mathbf{J}\mathbf{v} = &-\frac{\gamma}{1+\alpha^2}\big[\mathbf{v}\times\mathbf{H}_{\text{eff}} + \mathbf{m}\times\mathbf{H}'\big] \\ &-\frac{\gamma\alpha}{1+\alpha^2}\big[\mathbf{v}(\mathbf{m}\cdot\mathbf{H}_{\text{eff}}) + \mathbf{m}(\mathbf{v}\cdot\mathbf{H}_{\text{eff}} + \mathbf{m}\cdot\mathbf{H}') - \mathbf{H}'(\mathbf{m}\cdot\mathbf{m}) - 2\mathbf{H}_{\text{eff}}(\mathbf{m}\cdot\mathbf{v})\big] \\ &+\frac{1}{t_{\text{relax}}}\big[\mathbf{v}(1-|\mathbf{m}|^2) - 2\mathbf{m}(\mathbf{v}\cdot\mathbf{m})\big]\end{aligned} \tag{2.70}$$

In practice, only the exchange and anisotropy are included in the Jacobean since both for finite difference and for finite element methods (with the box method) the matrix $\mathbf{H}'$ is sparse and the product $\mathbf{H}'\mathbf{v}$ can be computed using a single sparse matrix-vector multiplication. Since the characteristic timescales for the demagnetising interaction are usually much smaller, even with demagnetisation excluded from the Jacobean the implicit algorithm remains efficient.

The time integration methods described in this section were implemented in the finite element package *finmag* used in Chapter 5 as well as the finite difference package *simlib* used in Chapters 3, 4, and for the eigenvalue method in 5.

2.5 Ferromagnetic resonance

Ferromagnetic resonance (FMR) is a common experimental technique to study the dynamics of magnetisation using microwave fields. The absorption of the microwave field is measured as a function of frequency; absorption is maximised when the frequency of the excitation matches one of the resonant frequencies of the magnetic system. The absorption spectrum can then be used to infer the properties of the system [32, 33, 34, 35].

Several methods are available to model ferromagnetic resonance in computation

- Direct method: application of a time-dependent periodic magnetic field and measuring the amplitude of the precession of magnetisation induced by the field.
- Ringdown method: the system is perturbed from an equilibrium state by a small excitation and subsequent dynamics are recorded. The Fourier transform of recorded data will reveal the resonance frequencies and corresponding modes [36].
- Eigenvalue method: instead of computing the time evolution, the problem is reformulated as an eigenvalue problem; the eigenvectors correspond to normal modes and eigenvalues to frequencies and damping times of the modes [37].

Modelling of ferromagnetic resonance is investigated in more detail in Chapter 5. The next section provides a background for the eigenvalue method (method 3); it is based on an appendix from the draft paper [3] that Chapter 5 is based on. Section 2.5.1.4 introduces a novel method to compute the FMR spectrum of the damped system using the linear least-squares approximation.

2.5.1 Eigenvalue approach

In this section, we provide a brief summary of the eigenvalue method described in [37], with modifications required to compute the Gilbert damping and excitation dependent FMR spectrum of

the system along with the resonance frequencies and corresponding normal modes.

The dynamics of the micromagnetic system are governed by the Landau-Lifshitz-Gilbert (LLG) equation:

$$\dot{\mathbf{m}} = -\frac{\gamma}{1+\alpha^2}\left[\mathbf{m}\times\mathbf{H}_{\text{eff}} + \alpha\,\mathbf{m}\times(\mathbf{m}\times\mathbf{H}_{\text{eff}})\right] = \mathcal{L}(\mathbf{m}), \tag{2.71}$$

where $\mathbf{m}$ is the normalized magnetization: $\mathbf{m} = \mathbf{M}/M_{\text{s}}$, with $|\mathbf{M}| = M_{\text{s}}$ being the saturation magnetisation. If the system is in its equilibrium state $\mathbf{m}_0$, then $\mathcal{L}(\mathbf{m}_0) = 0$. Small perturbations from the equilibrium can be described as $\mathbf{m} = \mathbf{m}_0 + \varepsilon\mathbf{v}$, with $\mathbf{v} \perp \mathbf{m}_0$ since the $|\mathbf{m}| = 1$ condition is imposed. For a small ε, terms of order ε^2 and higher can be neglected, which results in the linearized equation (for the general case):

$$\dot{\mathbf{v}} = \left.\frac{\partial\mathcal{L}}{\partial\mathbf{m}}\right|_{\mathbf{m}=\mathbf{m}_0}\mathbf{v}. \tag{2.72}$$

If the fixed linear operator $\hat{L} = \left.\frac{\partial\mathcal{L}}{\partial\mathbf{m}}\right|_{\mathbf{m}=\mathbf{m}_0}$ is defined, the linearized equation can be written as $\dot{\mathbf{v}} = \hat{L}\mathbf{v}$. This is an ordinary differential equation, which can be solved by an ansatz of the form $\mathbf{v} = \mathrm{Re}(\tilde{\mathbf{v}}e^{i2\pi ft})$. The normal modes (eigenvectors) $\tilde{\mathbf{v}}$ and oscillation frequencies (eigenvalues) f can be found from the following eigenvalue problem

$$i2\pi f\tilde{\mathbf{v}} = \hat{L}\tilde{\mathbf{v}} \tag{2.73}$$

2.5.1.1 Linearized equation without damping

First, we consider the case when the damping term in the LLG equation is neglected ($\alpha = 0$). Without damping, the magnetic moments precess indefinitely, and the LLG equation preserves energy. In this simplest case the calculation of the linearized operator $\hat{L}$ is fairly straightforward and results in the following linearized equation of motion [37]

$$\dot{\mathbf{v}} = \gamma\,\mathbf{m}_0\times\hat{A}\mathbf{v}, \qquad \hat{A} = |\mathbf{H}_{\text{eff}}(\mathbf{m}_0)|\,\mathrm{Id} - \frac{\partial\mathbf{H}_{\text{eff}}}{\partial\mathbf{v}} \tag{2.74}$$

where $\hat{A}$ is a positive definite Hermitian operator. The normal modes $\tilde{\mathbf{v}}$ and frequencies f of the linearized equation can be found from the eigenvalue problem

$$-i2\pi f\mathbf{m}_0\times\tilde{\mathbf{v}} = \gamma\,\hat{A}\tilde{\mathbf{v}} \tag{2.75}$$

The left-hand side of this eigenvalue problem contains the operator $\hat{B}\mathbf{v} = -i\,\mathbf{m}_0\times\tilde{\mathbf{v}}$ describing uniform precession. This operator is Hermitian but not positive definite (its eigenvalues are ±1).

Because of energy conservation, the oscillation frequencies f_k that satisfy this eigenvalue problem will be real and the normal modes $\tilde{\mathbf{v}}_k$ corresponding to different frequencies will be orthogonal. These properties enable the efficient numerical solution of (2.75); the eigenvalues f are the resonant frequencies and the complex magnitudes of the eigenvectors $\tilde{\mathbf{v}}$ are the normal mode amplitude plots (the complex phase of $\tilde{\mathbf{v}}$ corresponds to the phase of the oscillations at the corresponding sites).

However, in order to compute the FMR spectrum via the eigenvalue approach, we have to consider the more complicated case of non-zero damping.

2.5.1.2 Linearized equation with damping: perturbative analysis

For the case of sufficiently small non-zero damping α, a perturbative analysis can be performed to determine the corrections to the eigenvalues. In this case, eigenvalues have the form $\lambda = i2\pi f - 1/\tau$, where τ is the characteristic time for the mode to decay to $1/e$ of its starting amplitude value. It turns out that in the first order, the resonance frequencies are unchanged, and the damping times can be found using a relatively simple analytic calculation without solving the perturbed eigenvalue equation numerically [37]. Additionally, the coupling between the perturbed normal modes is small if their frequencies are sufficiently separated — this property will be useful for the calculation of the FMR spectrum.

2.5.1.3 Linearized equation with damping: numerical solution

To compute the actual FMR spectrum, we have to perform some extra work beyond the perturbation analysis. More precisely, we have to derive the linearized equation in the presence of damping, and solve the corresponding eigenvalue problem. The derivation of the linearized equation with damping is straightforward but slightly tedious. We skip this derivation and instead compute the linearized equation using the following numerical differentiation trick.

For the linearized equation, we have to compute the directional derivative

$$\hat{L}v = \left(\frac{\partial \mathcal{L}}{\partial \mathbf{m}}\bigg|_{\mathbf{m}=\mathbf{m}_0} \right)[v] = \frac{d}{d\epsilon}\mathcal{L}(m_0 + \epsilon v)|_{\epsilon=0} \tag{2.76}$$

For the test problem, the components of the effective field (demagnetization, exchange, bias) are all either constant, or linear functions of $\mathbf{m}$; therefore as a function of ϵ, $\mathcal{L}(m_0 + \epsilon v)$ is a degree 3 polynomial (the highest degree coming from the damping term $\mathbf{m} \times \mathbf{m} \times H_{\text{eff}}$). This means that a numerical differentiation rule of order 3 or higher will compute the derivative $\frac{d}{d\epsilon}\mathcal{L}(m_0 + \epsilon v)|_{\epsilon=0}$ exactly.

2.5.1.4 Linearized equation with damping: spectrum computation

The previous sections outlined the method used to compute the frequencies and normal mode shapes for the linearized equation, with or without damping. In this section we describe the subsequent computation of the FMR spectrum, which also depends on the initial state of the system. To determine the contributions of each normal mode to the total spectrum, we have to compute the coupling between the initial state and the normal modes in the presence of damping. More precisely, let n be the total number of the degrees of freedom (for a mesh with N nodes, $n = 3N$). Let $\tilde{v}_i$, $i = 1, 2, \ldots, n$ be the set of eigenvalues (normal modes) without damping, $\tilde{v}_i^{(p)}$, $i = 1, 2, \ldots, n$ the set of perturbed eigenvalues in the presence of damping, and f_i and τ_i the corresponding mode frequencies and damping times. Let $v_{\text{initial}} = m_{\text{initial}} - m_0$ be the initial state of the system. Due to the orthogonality property described in Section 2.5.1.1, we can assume that the non-perturbed eigenvalues form an orthonormal basis with the respect to the Hermitian inner product defined by the operator $\hat{A}$ ($(x, y) := x \cdot \hat{A}y^*$): $(v_i, v_j) = \delta_{ij}$.

To solve the linearized equation of motion Eq. (2.72), we need to expand the initial state v_{initial} in the perturbed $\tilde{v}_i^{(p)}$ basis:

$$v_{\text{initial}} = \sum_{i=1}^{n} C_i \tilde{v}_i^{(p)} \tag{2.77}$$

Once the coefficients C_i are known, the full solution of the linearized equation (2.72) is

$$m(t) = m_0 + \sum_{i=1}^{n} C_i e^{(2\pi i \omega_i - 1/\tau_i)t} \tag{2.78}$$

Given this full analytic solution, we can then calculate the spectrum using either of the methods described in Section 5.2.3. Unfortunately, this expansion requires the knowledge of the complete set of eigenvectors $\tilde{v}_i^{(p)}$, which is numerically unfeasible to compute. Instead, we will attempt to reconstruct the spectrum based on the first k perturbed modes with the lowest frequencies (we used $k = 40$). We would like to do this by finding the "best" approximation

$$v_{\text{initial}} = \sum_{i=1}^{k} c_i \tilde{v}_i^{(p)} + R \tag{2.79}$$

We will look for this approximation in the subspace spanned by the first k' non-perturbed normal modes $\tilde{v}_i$, with $k' > k$ (we used $k' = 60$). Due to the frequency separation property mentioned earlier, we can expect that this restriction will not affect the residue (the high-frequency modes will not measurably contribute to the low-frequency spectrum). When restricted to this subspace, we arrive at a system of k' equations with k unknowns c_i

$$(v_{\text{initial}}, v_j) = \sum_{i=1}^{k} c_i (\tilde{v}_i^{(p)}, v_j), \quad i = 1 \dots k' \tag{2.80}$$

This linear system is overspecified but any residue will only contain the high-frequency modes with frequencies above f'_k, which do not contribute to the spectrum for frequencies below f_k that we are trying to compute. It can be solved with a standard linear least-squares (linear regression) method, allowing the determination of the coupling coefficients c_i and thus the full solution (2.78).

Chapter 3

Demagnetizing tensor in finite difference micromagnetics

In the finite difference method, the demagnetizing field is usually computed as a convolution of the magnetization vector field with the demagnetizing tensor (see Section 2.3.2). An analytical expression for the demagnetizing tensor is available, however at distances far from the cuboidal cell, the numerical evaluation of the analytical expression can be very inaccurate.

Due to this large-distance inaccuracy numerical packages such as OOMMF compute the demagnetizing tensor using the explicit formula at distances close to the originating cell, but at distances far from the originating cell a formula based on an asymptotic expansion has to be used. In this chapter, we describe a method to calculate the demagnetizing field by numerical evaluation of the multidimensional integral describing the demagnetizing tensor. This method improves the accuracy of computation at intermediate distances from the origin.

The work described in this chapter was published in the Journal of Magnetism and Magnetic Materials [2].

3.1 Introduction

Micromagnetic simulation of ferromagnetic nanostructures is a widespread tool to support research and device design in a variety of fields, including magnetic data storage and sensing. The micromagnetic theory is based on partial differential equations proposed in [38] combined with an equation of motion that can be solved to determine the time-development of the magnetization vector function.

3.1.1 Micromagnetics

Numerical simulations in micromagnetics commonly solve the Landau-Lifshitz-Gilbert and the associated partial differential equations using a finite difference discretization of space (including OOMMF, LLG Micromagnetics Simulator, Micromagus, Mumax, Micromagnum [15, 39, 40, 41, 42]). In the finite difference method, space is discretized using a regular grid with cuboidal cells, and the magnetization and other scalar and vector fields involved in the computation are assumed

to be constant within each of these cubiodal cells. The magnetization vector field $\mathbf{M}$ is the primary degree of freedom. As a function of $\mathbf{M}$, which is represented by a constant within every cell, various fields such as the exchange, anisotropy, Zeeman, and demagnetizing fields are computed. These are added together, and enter the equation of motion for $\mathbf{M}$ as the effective field.

The most demanding part of the calculation is to determine the demagnetizing field: to compute the demagnetizing field for one of the discretization cells, an integral over the whole magnetic domain has to be carried out, which translates into a (triple) sum (in a 3d system) over all cells in the finite difference discretization.

For a finite difference discretization of a three-dimensional sample with n_1 discretization points in the x-direction, n_2 points in the y-direction and n_3 points in the z-direction, there are $n = n_1 n_2 n_3$ cuboidal cells in total. To compute the demagnetizing field in each one of these cells, we need to consider the total contribution of all n cells. Thus, to work out the demagnetizing field for all n cells requires $\mathcal{O}(n^2)$ operations using a naive approach. For realistic mesh sizes this is infeasible. Instead, usually the demagnetizing field $\mathbf{H}$ is expressed as a discrete convolution of the so-called demagnetizing tensor $\mathbf{N}$ and the representation of the magnetization $\mathbf{M}$

$$\mathbf{H}_{i_1 i_2 i_3} = -\sum_{j_1=1}^{n_1} \sum_{j_2=1}^{n_2} \sum_{j_3=1}^{n_3} \mathbf{N}(\mathbf{r}_{i_1 i_2 i_3} - \mathbf{r}_{j_1 j_2 j_3}) \cdot \mathbf{M}_{j_1 j_2 j_3}, \tag{3.1}$$

The triple indices $i_1 i_2 i_3$ are used to index cells in the 3d-discretization, i.e. i_1 counting cells in the x-direction, and correspondingly i_2 and i_3 in y- and z- direction. The vector $\mathbf{r}_{i_1 i_2 i_3}$ points to the centre of the cell $i_1 i_2 i_3$ and $\mathbf{M}_{i_1 i_2 i_3}$ is the magnetization in that cell, and $\mathbf{H}_{i_1 i_2 i_3}$ is the demagnetizing field in that cell.

The discrete convolution (3.1) is typically carried out as a product in Fourier space as the regular spacing of the finite difference cells allows straightforward use of the Fast Fourier Transform and its inverse. The demagnetizing tensor $\mathbf{N}$ needs to be computed once for given geometry and discretization, normally at the setup stage of the simulation. In this work we propose a new procedure for the accurate computation of entries in the demagnetizing tensor.

The tensor $\mathbf{N}(\mathbf{r})$ describes the energy of the demagnetizing interaction between two uniformly magnetized cuboids σ_1 and σ_2 of volume $|\sigma|$ separated by the translation vector $\mathbf{r}$. It is a symmetric tensor of rank 2, which we write in the dimensionless form following the convention from [11]

$$\mathbf{N}(\mathbf{r}) = -\frac{1}{4\pi|\sigma|} \int_{\sigma_1} d\mathbf{r}_1 \int_{\sigma_2(\mathbf{r})} \nabla_{\mathbf{r}_1} \nabla_{\mathbf{r}_2} \frac{1}{|\mathbf{r}_1 - \mathbf{r}_2|} \, d\mathbf{r}_2 \tag{3.2}$$

$$E_{\sigma_1 \leftrightarrow \sigma_2} = \mu_0 |\sigma| \, \mathbf{M}_1 \cdot \mathbf{N}(\mathbf{r}) \cdot \mathbf{M}_2 \tag{3.3}$$

The computation of the demagnetizing field using the formula (3.1) follows the commonly used energy-based approach to the discretization of the Landau-Lifshitz-Gilbert equation [8]. In this approach (used in OOMMF [15] and in our work), the components of the effective field in each cuboidal cell are obtained via a minimization procedure applied to the discretized total energy $E_{\text{total}} = \sum_{\sigma_1, \sigma_2} E_{\sigma_1 \leftrightarrow \sigma_2}$. In contrast, in the field-based approach the discretized fields are obtained from the corresponding continuous fields by e.g. sampling at the cell center [43, 44, 45, 8], and the demagnetizing tensor (3.2) is not used.

3.1.2 Numerical accuracy of $\mathbf{N}(\mathbf{r})$ evaluation

The components of the demagnetizing tensor can be computed via an analytical formula [10, 11, 12, 13]. However, when $\mathbf{r} = \mathbf{r}_a - \mathbf{r}_b$ is large compared to the size of the mesh cells (i.e. when the interacting mesh cells at $\mathbf{r}_a$ and $\mathbf{r}_b$ are far apart on the grid), evaluation of this expression on a computer can result in a loss of significant digits, to a point where the computed answer contains no significant digits at all [14].

The loss of accuracy is caused by catastrophic cancellation: the terms of the analytical expression correspond to indefinite integrals with r^3 order of growth, while the demagnetizing tensor itself (the definite integral) is of the order $1/r^3$. On modern CPUs, double-precision floating numbers contain approximately 15 significant digits, therefore the relative error in the computation of the demagnetizing tensor using the analytical formula is of the order $10^{-15}r^6$. One can therefore expect that for cell separations greater than $10^{15/6} \approx 300$ the analytical computation will contain no significant digits at all (i.e., the relative error will be greater than 1).

Indeed, the above estimate is confirmed if the result of the computation using the analytical formula is compared to the exact (up to machine precision) value of the demagnetizing tensor. The exact value is computed using specialized high-precision libraries (see Sec. 3.2.2 and 3.3.7). As seen in Fig. 3.1, the relative error η of the analytical computation grows as r^6 and crosses the $\eta = 1$ threshold at a separation of about 300 cells.

The micromagnetic simulation package OOMMF counteracts this inaccuracy problem by utilizing an asymptotic expansion of the demagnetizing tensor [14] in terms of powers of $1/r$ up to 6th order. In this chapter we investigate an alternative approach to deal with the catastrophic cancellation problem, which is to compute the integral (3.2) directly using numerical integration.

3.2 Method

As outlined in section 3.1, the 6-d integral described in (3.2) can be computed using

- an analytical formula [11],
- an asymptotic expansion [14],
- numerical integration.

For computing integrals in one dimension, multiple highly-accurate methods are available. However, the computation of multidimensional integrals is hindered by the so-called curse of dimensionality, where the number of integration points increases exponentially with the increase in dimension. Since the demagnetizing tensor $\mathbf{N}(\mathbf{r})$ has to be computed for all possible grid offsets $\mathbf{r}$, the integration method for this six-dimensional integral needs to be both accurate and fast.

The tradeoff between accuracy and computational effort can be achieved using the sparse grid family of methods, also known as Smolyak quadrature [46] (for a brief review of other alternatives see [47]).

The sparse grid method is used to extend one-dimensional integration rules to integration formulas in multiple dimensions. Below we summarize the key ideas of the method [47]. Starting with a one-dimension family of integration formulas I_k for computing the 1d integral

$$\int_0^1 f(x)dx \approx I_k[f] = \sum_{i=1}^{n_k} a_{ki} f(x_{ki}) \tag{3.4}$$

we can write the formal identity

$$\begin{aligned}\int_0^1 f(x)dx &= I_0[f] + (I_1[f] - I_0[f]) + \dots \\ &= \Delta_0[f] + \Delta_1[f] + \dots\end{aligned}$$

where $\Delta_k = I_k - I_{k-1}$ and $\Delta_0 = I_0$.

An example for a family of integration formulas I_k are Gaussian integration formulae with $k+1$ points.

Now, we apply the above formal identity d times to obtain a d-dimensional integration rule

$$\int_{[0,1]^d} f(\mathbf{x})d\mathbf{x} = \Big(\prod_{j=1}^{d}(\Delta_0^{(j)} + \Delta_1^{(j)} + \dots)\Big)[f] \tag{3.5}$$

Here the $\Delta_k^{(j)}$ operator applies the approximation (3.4) to the j-th argument of the function f, and the $\prod$ symbol represents the operator product, i.e. repeated application of $\sum_k \Delta_k^{(j)}$ in all d dimensions.

The formal expansion (3.5) is infinite; to obtain a practical integration formula we need to truncate it. Smolyak quadrature achieves this by expanding the product (3.5) and grouping together terms $\Delta_{i_1}^{(1)} \cdots \Delta_{i_d}^{(d)}$ with the same term order $k = i_1 + \dots + i_d$:

$$\int_{[0,1]^d} f(\mathbf{x})d\mathbf{x} = Q_0[f] + Q_1[f] + \dots \tag{3.6}$$

$$\approx Q_0[f] + Q_1[f] + \dots + Q_k[f] \tag{3.7}$$

where

$$Q_k[f] = \sum_{i_1+\dots+i_d=k} \Delta_{i_1}^{(1)} \Delta_{i_2}^{(2)} \cdots \Delta_{i_d}^{(d)}[f] \tag{3.8}$$

Compared to the evaluation of the d-dimensional integral using a naive d-fold product (which required n_k^d integration points for rule I_k), the Smolyak formula (3.7) greatly reduces the number of integration points required to achieve a certain level of accuracy, as long as the integrand is reasonably smooth.

Different one-dimensional families (3.4) will result in different multidimensional formulas (3.7); in our testing for the demagnetizing integrand (3.2) we obtained the best results when using the "delayed Kronrod-Patterson sequence" developed by Petras [47]. The delayed sequence is based on the Kronrod-Patterson family of 1d integration formulas I_k^{KP} [48], however some of the formulas are repeated to lower the rank of approximation (and the required number of integration points), determined by the "delay sequence" k_i:

$$I_i^{\mathrm{delayed}} = I_{k_i}^{\mathrm{KP}} \tag{3.9}$$

For the maximum delay sequence [47] the formulas are repeated so that the 1-d rule I_i^{delayed} is accurate for polynomials up to rank i:

$$k_i^{\text{full}} = 0, 1, 1, 2, 2, 2, 3, 3, 3, 3, 3, 3, 4, \ldots \tag{3.10}$$

3.2.1 Integrand

3.2.1.1 6d method

A straightforward way to compute the demagnetizing tensor numerically would be to apply the sparse grid formulas to the 6-d integral (3.2).

3.2.1.2 4d method

Due to the high dimensionality of the 6d integral, the number of required integration points will be quite high. To reduce the dimensionality, we can transform the 6-d volume integral (3.2) to a 4-d surface integral using a variant of Gauss's theorem. The procedure is described in [11] and results in the following formulas for the components of the demagnetizing tensor where we have used the notation of [11]

$$N_{xx}(X,Y,Z) = \frac{1}{4\pi\Delta x\,\Delta y\,\Delta z}\big[2F(X,Y,Z) - F(X+\Delta x,Y,Z) - F(X-\Delta x,Y,Z)\big]$$

$$N_{xy}(X,Y,Z) = \frac{1}{4\pi\Delta x\,\Delta y\,\Delta z}\big[G(X,Y,Z) - G(X-\Delta x,Y,Z) - G(X,Y+\Delta y,Z) + G(X-\Delta x,Y+\Delta y,Z)\big]$$

with

$$F(X,Y,Z) = \int_0^{\Delta z}\int_0^{\Delta y}\int_0^{\Delta z}\int_0^{\Delta y} \frac{dz\,dy\,dz'\,dy'}{\sqrt{X^2 + (y+Y-y')^2 + (z+Z-z')^2}}$$

and

$$G(X,Y,Z) = \int_{Y-\Delta y}^{Y}\int_{Z-\Delta z}^{Z}\int_{z}^{z+\Delta z}\int_{X}^{X+\Delta x} \frac{dy\,dz\,dz'\,dx'}{\sqrt{x'^2 + y^2 + z'^2}}$$

where $(X,Y,Z) = \mathbf{r}$ is the vector between the two interacting cells, and Δx, Δy, Δz are the edge lengths of the cuboidal cells.

To simplify the application of numerical integration formulas, we transform the above expressions so that integration is performed over the unit cube $[0,1]^4$.

$$N_{xx}(X,Y,Z) = \int_0^1\int_0^1\int_0^1\int_0^1 \big[2\mathcal{F}(X,Y,Z) - \mathcal{F}(X+\Delta x,Y,Z) - \mathcal{F}(X-\Delta x,Y,Z)\big]\,dz\,dy\,dz'\,dy' \tag{3.11}$$

$$N_{xy}(X,Y,Z) = \int_0^1\int_0^1\int_0^1\int_0^1 dy\,dz\,dz'\,dx' \times$$
$$\times\big[\mathcal{G}(X,Y,Z) - \mathcal{G}(X-\Delta x,Y,Z) - \mathcal{G}(X,Y+\Delta y,Z) + \mathcal{G}(X-\Delta x,Y+\Delta y,Z)\big] \tag{3.12}$$

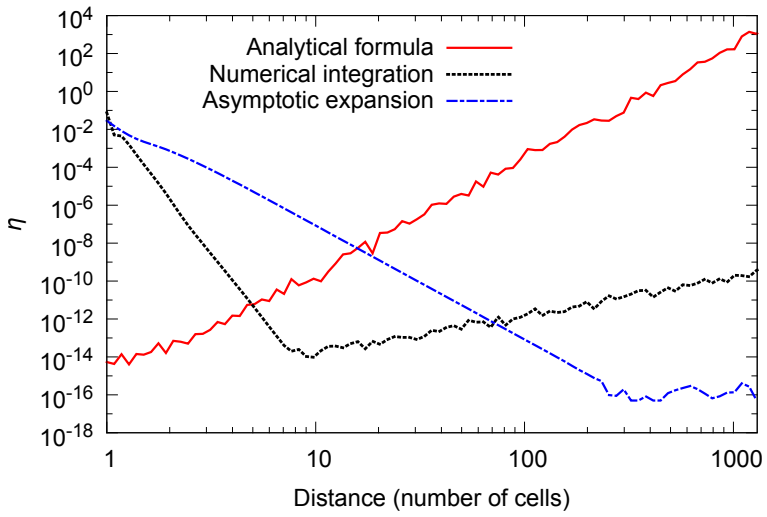

Figure 3.1: The relative error η in the computation of the demagnetizing tensor as a function of the distance between the interacting cuboids, $1\times1\times1$ cell size, Kronrod-Patterson sparse grid integration with full delay, order $k = 7$ [47], double precision arithmetic (precision $\sim 10^{-16}$).

with

$$\mathcal{F}(X,Y,Z) = \frac{\Delta y\,\Delta z}{4\pi\,\Delta x}\frac{1}{\sqrt{X^2+(\Delta y\,y+Y-\Delta y\,y')^2+(\Delta z\,z+Z-\Delta z\,z')^2}}$$

$$\mathcal{G}(X,Y,Z) = \frac{\Delta z}{4\pi}\frac{1}{\sqrt{(X+\Delta x\,x')^2+(Y+\Delta y(y-1))^2+(Z+\Delta z(z+z'-1))^2}}$$

The remaining components of the tensor can be obtained by variable substitution (for example, $N_{yz}(X,Y,Z) = N_{xy}(Y,Z,X)$).

In this approach, we are essentially computing two dimensions of the 6-d integral analytically, and the remaining four numerically. As we shall see, the two analytical steps introduce a small amount of cancellation, but the required number of integration points is significantly reduced.

3.2.2 Error estimation

To determine the accuracy of the computed demagnetizing tensor $\mathbf{N}$, we compute its exact (to machine precision) value $\mathbf{N}_{\text{exact}}$ via the analytical formula using the GNU MPFR high-precision floating point library [49], and evaluate the relative error η:

$$\eta = \frac{||\mathbf{N}-\mathbf{N}_{\text{exact}}||}{||\mathbf{N}_{\text{exact}}||} \tag{3.13}$$

where the matrix norm is defined as $||\mathbf{N}|| = \sqrt{\sum_{i,j} N_{ij}^2}$.

3.3 Results

3.3.1 Overview

Figure 3.1 shows the comparison of the accuracy of the analytical formula, the asymptotic expansion, and the numerical sparse grid integration methods as a function of the distance between cells

using (8 byte) double floating point. For the numerical sparse grid integration, the 4d method (Sec. 3.2.1.2) has been used.

While the aspect ratio of the cuboidal cell with edge lengths $\Delta x, \Delta y$ and Δz affects the results somewhat, they are independent of the absolute size of the cuboidal cell. We have chosen $\Delta x, \Delta y$ and $\Delta z = (1, 1, 1)$ so that the distance $|\mathbf{r}|$ between interacting cuboids is expressed in the number of cells between the interacting cuboids. For example, for a micromagnetic simulation with a $5\,\mathrm{nm} \times 5\,\mathrm{nm} \times 5\,\mathrm{nm}$ cell size, the distance of 10 on the plot in Figure 3.1 corresponds to a $50\,\mathrm{nm}$ distance on the mesh.

We start by discussing the analytical formula shown as a red solid line in Fig. 3.1. Its relative error η for very short distances is 10^{-14}. We cannot expect an error below 10^{-16} as this is the precision of the double floating point numbers used. As the distance between interacting cells increases, the analytical formula becomes less accurate. At a separation of 100 cells, it is about 10^{-4}, meaning that only the first 4 digits are correct. In fact, beyond a distance of about 300 cells the relative error becomes greater than 1, indicating that no digits of the double float can be expected to be correct and that not considering the demagnetizing tensor beyond that point would be more accurate than computing it analytically.

The asymptotic expansion (blue dash-dotted line in Fig. 3.1) starts with a large relative error $\eta \sim 10^{-2}$ for short distances, which decreases as the distance increases. At about 200 cells distance, the relative error is $\sim 10^{-15}$ and remains of that magnitude for larger distances. The smooth reduction of the error with distance reflects the way the high-order moments of the cuboid interaction decay with increasing distance, thus making the asymptotic approximation increasingly more accurate. The asymptotic expansion is more accurate than the exact analytical expression for distances greater than about 11 cells.

The precise number of cells for which this crossover occurs depends on the aspect ratio of the discretization cell as well as the direction of the separation vector between the two interacting cuboids. In practical implementations of finite difference micromagnetic codes a crossover point needs to be identified. In OOMMF, by default the asympotic formula is used for distances above 32 cells, which for this example corresponds to an error of 10^{-6}.

The numerical sparse grid integration error (black dotted line in Fig. 3.1) also starts around 10^{-2} for short distances and decays to 10^{-14} for a cell distance of about 7. For very short distances, the integrand (3.2) varies quickly (it diverges for $|\mathbf{r}| \to 0$) and numerical integration is inaccurate. For cell distances between 7 and 70, numerical integration is more accurate than the analytical expression, and more accurate than the asymptotic expansion. Beyond radius 70, the asymptotic expansion is more accurate than numerical integration. The slight inrease in the relative error of the numerical formula with increasing distance is caused by the cancellation introduced by Gauss's theorem (see 3.2.1.2).

In summary, the numerical evaluation of the analytical formula is most accurate for short distances, and the asymptotic expansion is most accurate for long distances. The new sparse grid integration method introduced here is most accurate for intermediate distances.

3.3.2 Sparse grid integration parameters and execution performance

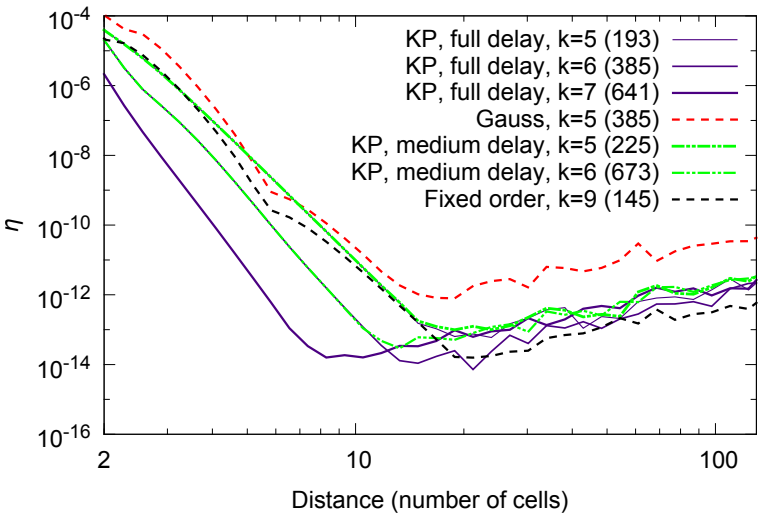

Figure 3.2: Comparison of numerical integration schemes for computing the 4d integral (3.11)-(3.12); KP: sparse grid integration based on the delayed Kronrod-Patterson 1d rule [47]; Gauss: sparse grid integration based on the delayed Gauss 1d rule; fixed order: 9th order rule with 145 integration points [50, 51]; k: the rule's order of approximation; $(\cdot)$: the number of integration points.

In the sparse grid method, the order of approximation is a parameter that can be adjusted to reach the desired level of accuracy or performance. For lower orders of approximation, one can also use fixed-order integration formulas that usually provide the same accuracy with fewer integration points. An extensive library of such formulas is available [50, 52, 53, 54].

Since the computation of the demagnetizing tensor is a one-off cost, higher order, more accurate integration formulas would be preferable. In Figure 3.2 we show the results for a number of sparse grid formulas as well as for a 145-point fixed-order formula from [50, 51].

The most accurate method considered here is the 641-point 7th order sparse grid method based on the 1d Kronron-Patterson sequence with full delay [47]. The delayed Kronrod-Patterson methods with orders 6 and 5 have fewer points and decreased accuracy. Formulas based on Gauss or Kronrod-Patterson sequences with medium delay require extra integration points to achieve the same level of accuracy and are thus suboptimal compared to the Kronrod-Patterson family with full delay.

The 145-point fixed order formula is more accurate than the 193-point 5th order Kronrod-Patterson formula, but slightly less accurate than the 385-point 6th order formula.

Based on these results, we recommend using either the 145-point fixed order formula or the Kronrod-Patterson fully delayed formula with order 6 or 7, depending on required accuracy and performance.

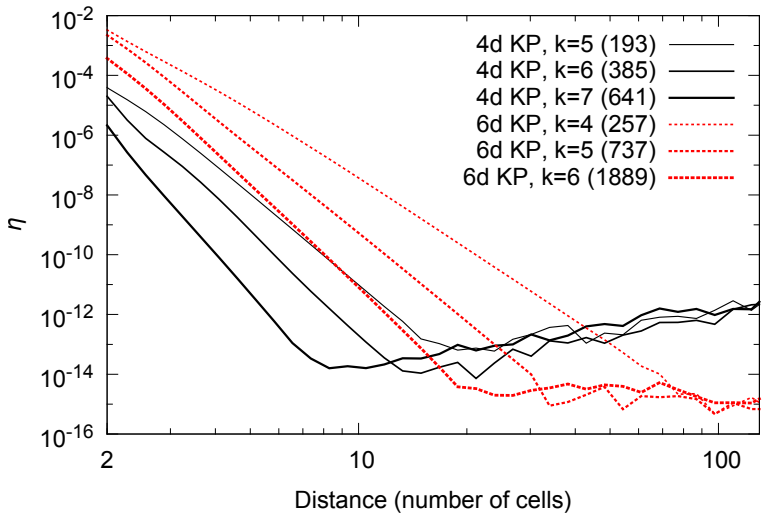

Figure 3.3: Comparison of numerical integration of the 4d integral (3.11)-(3.12) versus the 6d integral (3.2), Kronrod-Patterson 1d rule with full delay, $1 \times 1 \times 1$ cells.

3.3.3 Comparison of 4d and 6d integration

As an alternative to computing the 4d integral (method 3.2.1.2), we can also compute all of the 6 integrations numerically (method 3.2.1.1). The results for the 6d integration method are shown in Figure 3.3. All red dashed curves show results for the 6d method, and all black solid lines show results for the 4d method.

Both for the three 4d data sets and the three 6d data sets, we can see that the decrease in the error with increasing distance is faster for higher order methods. The 4d lines show minimal error between 8 and 11 cells distance, and for larger distances the error increases a little — this reflects the numerical cancellation from subtracting large terms that increase with distance and originate from the analytical integration that has been carried out over 2 dimensions. On the contrary, the 6d data sets — where the whole 6d integral has been computed numerically — does not suffer from this and the error remains small ($< 10^{-14}$) for larger distances. However, the number of evaluation points is higher compared to the 4d method.

As seen in Figure 3.2, even the lowest order quadrature formula gives a sufficient accuracy gain in the intermediate range to be better than the analytic and the asymptotic expression. However, in time-dependent micromagnetic simulations the computation of the demagnetizing tensor is a one-time setup cost, and the setup cost is only slightly influenced by the order of the integration formula — one might as well choose the higher order, more accurate formula.

Comparing the 4d and 6d method, we suggest the 4d method as it is faster to compute. The 4d method is less accurate than the 6d for the largest distances but this is of little practical concern — for those distances, the asymptotic expansion can be used instead.

3.3.4 Performance

In table 3.1, we show performance measurements for computing the entries in the demagnetizing tensor. For the fixed order (k=9) 145-point integration rule, the computational cost is comparable

to the analytical formula, while the KP full delay (k=7) 641-point rule is approximately 5 times slower. The total cost of evaluating the demagnetizing tensor on a $400 \times 40 \times 1$ mesh was $95\,\mathrm{ms}$ (Table 3.2) for the 145-point rule and $351\,\mathrm{ms}$ for the 641-point rule. Since a typical dynamical micromagnetic simulation usually requires 10,000 or more time steps, the one-time cost of setting up the demagnetizing tensor is minor even for the 641-point rule. For the sample system studied here, evaluating the effective field and dm/dt once takes $4\,\mathrm{ms}$. We also note that as the mesh grows larger, more of the entries can be computed using the extremely fast asymptotic formula, resulting in sublinear scaling of the total cost with mesh size n.

Method	Time per cell, ns
Analytical	8.5
Integration, 145 points	9.8
Integration, 641 points	43.2
Asymptotic	0.2

Table 3.1: Per-cell cost of computing the entries of the demagnetizing tensor, workstation: dual CPU Intel E5506 2.13 GHz (8 threads), compiler: GCC 4.7.2.

Method	Total time, ms
Combined, 145 points	95
Combined, 641 points	351
LLG dm/dt evaluation	4

Table 3.2: Total cost of computing the demagnetizing tensor for a $2000 \times 200 \times 20\,\mathrm{nm}^3$ mesh, $5 \times 5 \times 20\,\mathrm{nm}$ cells.

3.3.5 Single point floating precision

The recent rise of General Purpose computing on Graphical Processing Units (GPGPU) has re-invigorated single-precision floating point operations: on these architectures single precision floating point operations are generally much faster, and on cheaper cards the only type of floating point operations provided. Additionally, the available RAM on the GPU card is limited, providing another incentive to use single rather than double precision floating point numbers.

We repeat the study presented in figure 3.1 but use single precision numbers for all methods and show the results in figure 3.4. The qualitative findings are the same as for double precision numbers: the most accurate methods are as a function of increasing distance: (i) the analytical formula, (ii) the sparse grid numerical integration technique and (iii) the asymptotic expression.

However, the relevant cross-over points have moved to shorter distances. The analytical expression becomes less accurate than numerical integration for more than 2 cells distance, and the asymptotic expression is more accurate than numerical integration for spacings greater than 8 cells.

As mentioned in 3.3.1 we find that the analytical expression for double precision (Figure 3.1) provides only 4 significant digits (i.e. a relative error of 10^{-4}) for a distance of 100 cells. For single

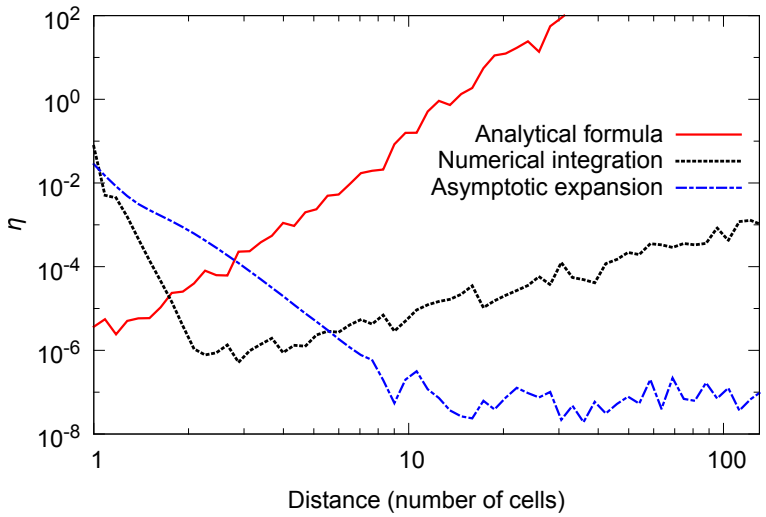

Figure 3.4: The relative error η of computing the demagnetizing tensor using single precision arithmetic (precision $\sim 10^{-7}$), $1\times1\times1$ cell size, Kronrod-Patterson sparse grid integration with full delay, order $k = 7$.

precision (Figure 3.4) we find that the analytical expression provides the same level of accuracy (i.e. 4 significant digits) only for distances up to 3 cells. Correspondingly, the distance for which the relative error exceeds 1, moves from over 300 cells with double precision to 11 cells with single precision.

The accurate calculation of the demagnetizing tensor entries using single precision floating point numbers only is challenging – using the best methods currently known and combining the three methods shown in figure 3.4, the relative error can be kept around or below 10^{-4}.

For practical use of GPGPU single precision calculations for micromagnetic simulation, we recommend to compute the demagnetizing tensor using double-precision – either on the GPU with some reduction in speed if the GPU hardware supports this, or on the CPU with a more significant time penalty. As the computation of the demagnetizing tensor (3.2) only needs to be done once and subsequent computations of the demagnetizing field as required for energy minimisation or time stepping only require to carry out the convolution (3.1), it should be acceptable to increase the accuracy of the demagnetizing tensor for a one-off time penalty in the setup phase.

3.3.6 Forward and backward Fast Fourier Transform

Our tests showed that the forward and inverse fast Fourier transforms required to compute the convolution (3.1) did not introduce any significant numerical error in the calculation of the demagnetizing field (either single or double precision).

3.3.7 Other high accuracy methods

We note for completeness that there are other options to compute the demagnetizing tensor more accurately than any of the methods outlined above if high accuracy is of utmost importance.

There are high precision arithmetic libraries available which provide software implementations of floating point operations: the library user can choose the number of significant digits used in the calculations (where 8 would approximately correspond to single precision accuracy and 16 to double precision floating point numbers). The larger the number of significant digits to be used, the slower is the execution of these operations. We have used this technique to obtain the reference data $\mathbf{N}_{\text{exact}}$ required to compute the error η (3.13). Using such libraries requires significant changes to source code, and execution is extremely slow. It is impractical to use such libraries for micromagnetic simulation.

3.4 Summary

We have compared the accuracy of computing the demagnetizing tensor using the analytical formula, the asymptotic expansion, and numerical integration. We obtain and provide quantitative data on the relative error of demagnetizing tensor entries computed using all three methods.

We propose a new method using numerical integration to compute the entries of the demagnetizing tensor which allows to increase the accuracy from an error of 10^{-8} to an error of only 10^{-12} for intermediate distances of between 4 and 80 simulation cells for the commonly used double precision floating point numbers. In the context of micromagnetic simulations, we find that the 7th order 641-point Kronrod-Patterson sparse grid formula with full delay [47] and the fixed-order 145-point rule [50, 51] provide a reasonable tradeoff between accuracy and performance.

In the context of recent GPGPU use in micromagnetic simulations, we also obtain accuracy data for the three methods using single precision floating point numbers.

Chapter 4

Computation of the magnetostatic interaction between linearly magnetized polyhedrons

In this Chapter we present a method to accurately compute the energy of the magnetostatic interaction between linearly (or uniformly, as a special case) magnetized polyhedrons. The method has applications in finite element micromagnetics, or more generally in computing the magnetostatic interaction when the magnetization is represented using the finite element method (FEM).

The magnetostatic energy is described by a six-fold integral that is singular when the interaction regions overlap, making direct numerical evaluation problematic. To resolve the singularity, we evaluate four of the six iterated integrals analytically resulting in a 2d integral over the surface of a polyhedron, which is nonsingular and can be integrated numerically. This provides a more accurate and efficient way of computing the magnetostatic energy integral compared to existing approaches.

The method was developed to facilitate the evaluation of the demagnetizing interaction between neighouring elements in finite-element micromagnetics and provides a possibility to compute the demagnetizing field using efficient fast multipole or tree code algorithms.

The work described in this chapter is currently being prepared for submission to a journal.

4.1 Introduction

In the continuum form of the Landau-Lifshitz-Gilbert equation, the effective field $\mathbf{H}_{\text{eff}}(\mathbf{r})$ is the functional derivative of the total energy functional $E(\mathbf{M})$ with respect to the magnetization $\mathbf{M}(\mathbf{r})$ [1]:

$$\mathbf{H}(\mathbf{r}) = -\frac{1}{\mu_0}\frac{\delta E}{\delta \mathbf{M}(\mathbf{r})} \tag{4.1}$$

$$E = E_{\text{Zeeman}} + E_{\text{demag}} + E_{\text{anisotropy}} + E_{\text{exchange}} + \dots \tag{4.2}$$

In numerical micromagnetics, the dynamics of magnetization are described by the semi-discretized Landau-Lifshitz-Gilbert equation, in which the motion of magnetization is computed from the discretized effective field. It is desireable to preserve the relation (4.1) between effective

field and total energy in the semi-discretized formulation [8] — if (4.1) holds for the discrete system, then the total energy will decrease in the simulation, simplifying the use of energy-based criteria for the control of the simulation or the search for an equilibrium.

In order to preserve (4.1), the effective field has to be computed from the discretized total energy function using the corresponding discrete counterpart to the functional derivative $\delta E/\delta \mathbf{M}(\mathbf{r})$. For the exchange, anisotropy, and Zeeman terms this is usually straightforward, however for the demagnetizing field it is more difficult. In finite difference (FD) micromagnetics, it can be achieved by computing the total demagnetizing energy of the system using the analytical expression [10, 11, 12, 13] for the demagnetizing tensor, and then differentiating with respect to the degrees of freedom [15, 8]. However, in finite element (FE) micromagnetics, the demagnetizing field is usually computed using the FEM/BEM method [18, 21] where the field is derived from the magnetostatic potential, and the energy is not computed exactly.

It is therefore desireable to be able to accurately and efficiently compute the total magnetostatic energy of a system represented by a set of polyhedral elements, with magnetization linear inside each polyhedron (as in the FE method with linear Lagrange elements). The total energy of the system is the sum of pairwise interactions between the polyhedrons and in this chapter we describe how to compute this pairwise interaction.

Given two interacting magnetized polyhedrons τ and τ' with arbitrary magnetizations $\mathbf{M}(\mathbf{r})$ and $\mathbf{M}'(\mathbf{r})$, the energy $E_{\tau\leftrightarrow\tau'}$ of the magnetostatic interaction between them is:

$$E_{\tau\leftrightarrow\tau'} = -\mu_0 \int_\tau \mathbf{M}(\mathbf{r}) \cdot \mathbf{H}'_{\mathrm{demag}}(\mathbf{r})\, d\mathbf{r} = \frac{\mu_0}{4\pi} \int_\tau \int_{\tau'} \mathbf{M}(\mathbf{r}) \cdot \big(\nabla_{\mathbf{r}} \nabla_{\mathbf{r}'} \frac{1}{|\mathbf{r}-\mathbf{r}'|}\big) \cdot \mathbf{M}'(\mathbf{r}')\, d\mathbf{r}\, d\mathbf{r}' \tag{4.3}$$

where $H'_{\mathrm{demag}}(\mathbf{r}) = -\frac{1}{4\pi}\nabla_{\mathbf{r}} \int_{\tau'} \mathbf{M}'(\mathbf{r}') \cdot \nabla_{\mathbf{r}'}(1/|\mathbf{r}-\mathbf{r}'|)\, d\mathbf{r}'$ is the demagnetizing (stray) field produced by the polyhedron τ'.

The straightforward approach of numerically computing the integral (4.3) is problematic because it requires explicit integration over a 6-dimensional region of space; additionally, when the polyhedrons overlap or coincide, the integrand is singular and regular integration methods cannot be applied. Many analytical results are available for similar 3-fold integrals arising during the calculation of the field [55, 56, 57, 58, 59]. However, none of the formulas can be easily adapted to this more complex 6d case. In [60] a method is developed for removing the singularity in (4.3) that can be applied to the uniform magnetization case but does not generalize easily to the linear magnetization case. A Fourier-transform method has been devised for the computation of magnetostatic energy [61, 62], however for the case investigated in this chapter, the required 3d numerical integration in the Fourier space is still somewhat impractical.

The rest of the chapter describes the proposed method for the computation of this integral. The main approach is to analytically perform 4 out of 6 iterated integrals resulting in a 2d surface integral that is nonsingular and can be evaluated numerically using standard methods. This semi-analytical approach is similar to [63]; the use of notation and vector analysis in the analytical derivation is similar to the techniques in [58, 59].

4.2 Formulation of the problem

For the purposes of computation, an arbitrary linear vector-valued function in space $\mathbf{M}(\mathbf{r})$ can be represented by a 3×4 matrix $||M_{ij}||$: $\mathbf{M}(\mathbf{r}) = ||M_{ij}|| \cdot (1, r_x, r_y, r_z)^T$. However, performing analytical calculations for this general case is quite inconvenient; instead we only consider vector-valued linear functions of the form $A(\mathbf{r})\mathbf{M}$ where $A(\mathbf{r})$ is a scalar linear function and $\mathbf{M}$ is a constant vector. For the common case of a tetrahedral element, an arbitrary linear vector-valued function $\mathbf{M}(\mathbf{r})$ can be reconstructed from the vertex values $\mathbf{M}_i$, $i = 1 \dots 4$: $\mathbf{M}(\mathbf{r}) = \sum_{i=1}^{4} A_i(r)\mathbf{M}_i$, where $A_i(r)$ are the shape functions of the tetrahedron.

We perform the computations for a pair of interacting linearly magnetized polyhedrons τ and τ' with magnetizations $A(\mathbf{r})\mathbf{M}$ and $B(\mathbf{r})\mathbf{M}'$, where $\mathbf{M}$ and $\mathbf{M}'$ are constant magnetization vectors and $A(\mathbf{r})$, $B(\mathbf{r})$ are dimensionless linear functions in space. From (4.3), the energy $E_{\tau\leftrightarrow\tau'}$ of the magnetostatic interaction between the polyhedrons is $E_{\tau\leftrightarrow\tau'} = \frac{\mu_0}{4\pi}\,\mathbf{M}\cdot\mathbf{N}\cdot\mathbf{M}'$, where $\mathbf{N}$ is the symmetric 3×3 "demagnetizing tensor"

$$\mathbf{N} = \int_{\tau}\int_{\tau'} A(\mathbf{r})B(\mathbf{r}')\nabla_{\mathbf{r}}\nabla_{\mathbf{r}'}\frac{1}{|\mathbf{r}-\mathbf{r}'|}\,d\mathbf{r}\,d\mathbf{r}' \tag{4.4}$$

The goal of this chapter is to compute this sixfold integral given the coordinates of the vertices of τ and τ' and the coefficients of the linear functions $A(\mathbf{r})$ and $B(\mathbf{r})$. When the polyhedrons τ and τ' are separated, the integral can be computed numerically, however when the polyhedrons overlap or coincide, the integrand is singular and standard numerical integration is inaccurate.

To deal with this issue, we analytically reduce the double volume integral (4.4) to a double surface integral, then evaluate the surface integral over $\mathbf{r}$ analytically, and the second surface integral integral over $\mathbf{r}'$ numerically. This procedure is similar to the one employed in [63] — the four analytical steps result in a surface integral with a bounded integrand that can be integrated numerically with reasonable accuracy and efficiency.

A short notice on units: the demagnetizing tensor commonly used in finite difference micromagnetics [11] is dimensionless, however the tensor $\mathbf{N}$ computed in this chapter (4.4) has units of volume.

4.3 Method

The analytical derivation proceeds in the following three steps:

- transform the double volume integral (4.4) to a double surface integral using Gauss's theorem, removing linear factors via integration by parts (Section 4.4),
- express the integrand for the outer surface integral over $\mathbf{r}'$ as a linear combination of primitive

terms (Section 4.5),

$$I_0(\tau;\mathbf{r}') = \int_{\partial\tau-\mathbf{r}'} \frac{1}{|\mathbf{R}|} ds \tag{4.5}$$

$$\mathbf{I}_1(\tau;\mathbf{r}') = \int_{\partial\tau-\mathbf{r}'} \frac{\mathbf{R}}{|\mathbf{R}|} ds \tag{4.6}$$

$$\mathbf{I}_2(\tau;\mathbf{r}') = \int_{\partial\tau-\mathbf{r}'} \frac{\mathbf{R}^{\otimes 2}}{|\mathbf{R}|} ds \tag{4.7}$$

where $\partial\tau$ is the polygonal surface of the polyhedron τ, $\partial\tau - \mathbf{r}'$ is the same surface shifted by $\mathbf{r}'$ according to the the substitution $\mathbf{R} = \mathbf{r} - \mathbf{r}'$, and $\otimes$ denotes tensor multiplication (i.e. $\mathbf{R}^{\otimes 2}$ is a symmetric tensor of rank 2).

- analytically integrate the primitive terms $\mathbf{I}_k$ over each polygonal surface, again by applying integration by parts, Stokes' theorem (for integration over a surface), and gradient theorem (for integration over a line), Section 4.5.1.

The equation (4.17) derived in step 2 (see Section 4.5) together with the corresponding analytical formulas for the primitive terms $\mathbf{I}_k$ form the main result of the chapter.

4.3.1 Auxiliary functions

The derivation of the analytical formulas for $\mathbf{I}_k$ requires the computation of a number of auxiliary functions:

$$J_0(\tau;\mathbf{r}') = \int_{\partial\tau-\mathbf{r}'} |\mathbf{R}| ds \tag{4.8}$$

$$\eta_0(\mathbf{R}_1,\mathbf{R}_2) = \int_{\mathbf{R}_1}^{\mathbf{R}_2} \frac{1}{|\mathbf{R}|} dl \tag{4.9}$$

$$\boldsymbol{\eta}_1(\mathbf{R}_1,\mathbf{R}_2) = \int_{\mathbf{R}_1}^{\mathbf{R}_2} \frac{\mathbf{R}}{|\mathbf{R}|} dl \tag{4.10}$$

$$\lambda_0(\mathbf{R}_1,\mathbf{R}_2) = \int_{\mathbf{R}_1}^{\mathbf{R}_2} |\mathbf{R}| dl \tag{4.11}$$

$$\boldsymbol{\lambda}_1(\mathbf{R}_1,\mathbf{R}_2) = \int_{\mathbf{R}_1}^{\mathbf{R}_2} \mathbf{R}|\mathbf{R}| dl \tag{4.12}$$

$$\tag{4.13}$$

The relation between the formulas for these functions is shown below (an arrow indicates that the formula at the source depends on the formula at the target):

$$\begin{array}{ccccccc} \eta_0 & \longleftarrow & I_0 & & & & \\ \uparrow & & \uparrow & & & & \\ \lambda_0 & \longleftarrow & \mathbf{I}_1 & & & & \\ \uparrow & & \uparrow & & & & \\ J_0 & \longleftarrow & \mathbf{I}_2 & \longrightarrow & \boldsymbol{\lambda}_1 & \longrightarrow & \boldsymbol{\eta}_1 \end{array} \tag{4.14}$$

4.4 Analytical derivation — reduction to a surface integral

The first step is the conversion of (4.4) to a double surface integral. Due to the high complexity of intermediate expressions the derivation was performed using a computer algebra system. Table 4.1 shows the identities used in the derivation, they are applied repeatedly in a straightforward manner to integrate the terms and simplify the result. We chose to express the vector/tensor identities using tensor notation which is compact and more suitable for computer algebra than using explicit indices (with Einstein summation).

After repeatedly applying the identities in Table 4.1, we arrive at the following formula for the demagnetizing tensor

$$\begin{aligned}\mathbf{N} = \int_{\partial\tau'}\int_{\partial\tau} \Big[&\frac{1}{2}A(\mathbf{r})(\mathbf{n}_{\mathbf{r}'}\cdot(\mathbf{r}-\mathbf{r}'))\,\mathbf{n}_{\mathbf{r}}\otimes\nabla B - \frac{1}{2}B(\mathbf{r}')(\mathbf{n}_{\mathbf{r}}\cdot(\mathbf{r}-\mathbf{r}'))\,\mathbf{n}_{\mathbf{r}'}\otimes\nabla A \\ &+\frac{1}{6}(\mathbf{n}_{\mathbf{r}}\cdot(\mathbf{r}-\mathbf{r}'))(\mathbf{n}_{\mathbf{r}'}\cdot(\mathbf{r}-\mathbf{r}'))\,\nabla A\otimes\nabla B + A(\mathbf{r})B(\mathbf{r}')\,\mathbf{n}_{\mathbf{r}}\otimes\mathbf{n}_{\mathbf{r}'}\Big]\frac{1}{|\mathbf{r}-\mathbf{r}'|}ds\,ds'\end{aligned} \tag{4.15}$$

where $\mathbf{n}_{\mathbf{r}}$, $\mathbf{n}_{\mathbf{r}'}$ are the normal vectors to the corresponding surfaces and ∇A, ∇B are the (constant) gradient vectors for the linear functions $A(\mathbf{r})$ and $B(\mathbf{r})$. As expected, the formula is symmetrical under replacement $A \leftrightarrow B$, $\mathbf{r} \leftrightarrow \mathbf{r}'$ and reduces to Gauss's theorem when $A(\mathbf{r})$ and $B(\mathbf{r})$ are constant.

Side note: in principle, all derivations could be performed with scalars instead of tensors by computing the scalar counterpart to the integral (4.4)

$$N(\mathbf{m},\mathbf{m}') = \mathbf{m}\cdot\mathbf{N}\cdot\mathbf{m}' = \int_{\tau}\int_{\tau'} A(\mathbf{r})B(\mathbf{r}')(\mathbf{m}\cdot\nabla_{\mathbf{r}})(\mathbf{m}'\cdot\nabla_{\mathbf{r}'})\frac{1}{|\mathbf{r}-\mathbf{r}'|}\,d\mathbf{r}\,d\mathbf{r}' \tag{4.16}$$

With this method the auxiliary vectors $\mathbf{m}$ and $\mathbf{m}'$ would have to be included in all intermediate derivations; this would remove the need to keep track of tensor indices at the cost of slightly expanded notation.

Gauss's theorem	$\displaystyle\int_V \nabla_{\mathbf{r}}\,\mathbf{F}\,dr = \int_{\partial V}\mathbf{n}\otimes\mathbf{F}\,dr$
gradient of a product	$a\,\nabla_{\mathbf{r}}\mathbf{F} = \nabla_{\mathbf{r}}[a\,\mathbf{F}] - \nabla a\otimes\mathbf{F}$
integration of $\frac{(\mathbf{r}-\mathbf{r}')^{\otimes k}}{\lvert\mathbf{r}-\mathbf{r}'\rvert}$	$\displaystyle\frac{(\mathbf{r}-\mathbf{r}')^{\otimes k}}{\lvert\mathbf{r}-\mathbf{r}'\rvert} = \frac{1}{k+2}\nabla_{\mathbf{r}}\cdot\frac{(\mathbf{r}-\mathbf{r}')^{\otimes k+1}}{\lvert\mathbf{r}-\mathbf{r}'\rvert}$

Table 4.1: Identities used in the conversion of the double volume integral (4.4) to a double surface integral. Here $\mathbf{F}(\mathbf{r})$ is a tensor of any rank and $a(\mathbf{r})$ is a scalar. When the gradient operator ∇ is applied to a vector or tensor, we assume that the new tensor index is added at the front (i.e. $\nabla f_{jk} = \partial_i f_{jk}$). Similarly, divergence applies to the first index of a tensor (i.e. $\nabla\cdot f_{jk} = \partial_j f_{jk}$).

4.5 Analytical result — integrand and the primitive terms $\mathbf{I}_k$

The next step is to express the inner integral over $\mathbf{r}$ in (4.15) in terms of the primitive integrals $\mathbf{I}_k$ (4.5)–(4.7). By using the identity $A(\mathbf{r}) = A(\mathbf{r}') - (\mathbf{r} - \mathbf{r}') \cdot \nabla A$, we get

$$\mathbf{N} = \int_{\partial\tau'} \Big[\frac{1}{2}A(\mathbf{r}')(\mathbf{n}' \cdot \mathbf{I}_1)(\mathbf{n} \otimes \nabla B) + B(\mathbf{r}')(\nabla A \cdot \mathbf{I}_1)(\mathbf{n} \otimes \mathbf{n}') - \frac{1}{2}B(\mathbf{r}')(\mathbf{n} \cdot \mathbf{I}_1)(\nabla A \otimes \mathbf{n}') + A(\mathbf{r}')B(\mathbf{r}')I_0\,\mathbf{n} \otimes \mathbf{n}' - \frac{1}{6}I_2(\mathbf{n}, \mathbf{n}')(\nabla A \otimes \nabla B) + \frac{1}{2}I_2(\nabla A, \mathbf{n}')(\mathbf{n} \otimes \nabla B)\Big] ds' \tag{4.17}$$

Note that the primitive integral $\mathbf{I}_k$ is a tensor of rank k and is a function of $\mathbf{r}'$.

4.5.1 Evaluation of $\mathbf{I}_k$

The integral I_0 has been computed in [58] (eq. (17) for $W_f(\mathbf{r})$ in [58])

$$I_0 = \sum_{F \in \partial\tau - \mathbf{r}'} \left[- \sum_{(\mathbf{R}_1, \mathbf{R}_2) \in \partial F} (\mathbf{n} \times \mathbf{u} \cdot \mathbf{R}_1)\, \eta_0(\mathbf{R}_1, \mathbf{R}_2) - (\mathbf{n} \cdot \mathbf{R}_f)\, \Omega(F) \right] \tag{4.18}$$

where the outer sum is over the facets F of the polyhedron surface $\partial\tau - \mathbf{r}'$, the inner sum is over the line edges $(\mathbf{R}_1, \mathbf{R}_2)$ of the facet, $\mathbf{n}$ is the facet normal, $\mathbf{u} = (\mathbf{R}_2 - \mathbf{R}_1)\,/|\mathbf{R}_2 - \mathbf{R}_1|$ is the unit vector along the edge, $\mathbf{R}_f$ is any point on the facet, and $\Omega(F))$ is the solid angle of the facet F from the origin.

The integral $\mathbf{I}_1$ has been computed in [59] (eq. (27) in [59])

$$\mathbf{I}_1 = \sum_{F \in \partial\tau - \mathbf{r}'} \left[- \sum_{(\mathbf{R}_1, \mathbf{R}_2) \in \partial F} (\mathbf{n} \times \mathbf{u})\, \lambda_0(\mathbf{R}_1, \mathbf{R}_2) + (\mathbf{n} \cdot \mathbf{R}_f)\, \mathbf{n}\, I_0(F) \right] \tag{4.19}$$

Note that in the expression for $\mathbf{I}_1$, each term of the sum over the facets references the integral $I_0(F)$ applied only to that facet, not the whole surface.

The derivation of the formula for the integral $\mathbf{I}_2$ is shown in 4.7.1

$$\mathbf{I}_2 = \sum_{F \in \partial\tau - \mathbf{r}'} \left[- \sum_{(\mathbf{R}_1, \mathbf{R}_2) \in \partial F} (\mathbf{n} \times \mathbf{u}) \otimes \lambda_1(\mathbf{R}_1, \mathbf{R}_2) + (\mathbf{n}^2 - \mathbf{Id})\, J_0(F) + (\mathbf{n} \cdot \mathbf{R}_f)\, \mathbf{n} \otimes \mathbf{I}_1(F) \right] \tag{4.20}$$

Again, each term of the sum over the facets references the integrals $J_0(F)$ and $\mathbf{I}_1(F)$ applied only to that facet.

The auxiliary function η_0 (used in the equations for I_0 and λ_1) has been computed in [58] (eq. (18) for $w_e(\mathbf{r})$)

$$\eta_0(\mathbf{R}_1, \mathbf{R}_2) = \ln \frac{|\mathbf{R}_1| + |\mathbf{R}_2| + |\mathbf{R}_2 - \mathbf{R}_1|}{|\mathbf{R}_2| + |\mathbf{R}_1| - |\mathbf{R}_2 - \mathbf{R}_1|}, \tag{4.21}$$

The auxiliary function λ_0 (used in the equations for $\mathbf{I}_1$ and J_0) has been computed in [59] (eq. (22) for $\lambda_e(\mathbf{r})$)

$$\lambda_0(\mathbf{R}_1, \mathbf{R}_2) = \frac{1}{2}\mathbf{u}(\mathbf{R}_2|\mathbf{R}_2| - \mathbf{R}_1|\mathbf{R}_1|) + \frac{1}{2}|\mathbf{R}_1 \times \mathbf{u}|^2\, \eta_0(\mathbf{R}_1, \mathbf{R}_2) \tag{4.22}$$

The auxiliary function J_0 (used in the equation for $\mathbf{I}_2$) has been computed in [58] (eq. (22) for $\Lambda_f(\mathbf{r})$)

$$J_0 = \sum_{F\in\partial\tau-\mathbf{r}'}\left[\frac{1}{3}\sum_{(\mathbf{R}_1,\mathbf{R}_2)\in\partial F}\mathbf{n}\times R_1\cdot\mathbf{u}\,\lambda_0(\mathbf{R}_1,\mathbf{R}_2)+\frac{1}{3}(\mathbf{R}_f\cdot\mathbf{n})^2 I_0\right] \tag{4.23}$$

The derivation of the formula for $\boldsymbol{\eta}_1$ (used in the equation for $\boldsymbol{\lambda}_1$) is shown in 4.7.2

$$\boldsymbol{\eta}_1(\mathbf{R}_1,\mathbf{R}_2) = \mathbf{u}(|\mathbf{R}_2|-|\mathbf{R}_1|)-\mathbf{u}\times(\mathbf{u}\times\mathbf{R}_1)\,\eta_0(\mathbf{R}_1,\mathbf{R}_2) \tag{4.24}$$

The derivation of the formula for $\boldsymbol{\lambda}_1$ (used in the equation for $\mathbf{I}_2$) is shown in 4.7.3

$$\boldsymbol{\lambda}_1(\mathbf{R}_1,\mathbf{R}_2) = \left(\frac{1}{2}\mathbf{Id}+\frac{1}{6}\mathbf{u}\,(\mathbf{u}\cdot)\right)\left[|\mathbf{R}|(\mathbf{u}\cdot\mathbf{R})(\mathbf{R}-\frac{1}{2}\mathbf{u}\,(\mathbf{u}\cdot\mathbf{R}))\Big|_{\mathbf{R}_1}^{\mathbf{R}_2}+|\mathbf{u}\times\mathbf{R}_1|^2(\mathbf{Id}-\frac{1}{2}\mathbf{u}\,(\mathbf{u}\cdot))\,\boldsymbol{\eta}_1(\mathbf{R}_1,\mathbf{R}_2)\right] \tag{4.25}$$

with the notation $f(\mathbf{R})\big|_{\mathbf{R}_1}^{\mathbf{R}_2} \equiv f(\mathbf{R}_2)-f(\mathbf{R}_1)$.

4.6 Numerical results

In order to numerically verify the analytical results in Section 4.5 we need a way to compute the energy integral (4.3) exactly or with sufficient precision. For a cuboid, we could do this if the magnetization was constant by using the analytical expression for the demagnetizing tensor [10, 11, 12, 13]. However, for a constant magnetization the gradient terms ∇A and ∇B in our analytical expressions would be zero, and for a more comprehensive test we have to cover the case of (nontrivial) linear magnetization. The authors are not aware of an analytical result that could be used as a reference in this case, instead, as a reference we used a series of finite difference micromagnetic simulations with progressively increasing mesh size until convergence was reached.

4.6.1 Test problem formulation

The test system is a magnetized cuboid with dimensions $1.7l \times 1.3l \times l$, where l is a arbitrary length parameter. A linear magnetization function $\mathbf{M}(\mathbf{r})$ of the cuboid can be written in the general form

$$\mathbf{M}(\mathbf{r}) = \mathbf{L}\cdot\begin{pmatrix} l \\ r_x \\ r_y \\ r_z \end{pmatrix}$$

where $\mathbf{L}$ is a 3×4 matrix. In this equation we have used the length parameter l to make the units of $\mathbf{L}$ uniform.

The 12 entries of the matrix $\mathbf{L}$ can also be written as a column vector $\mathbf{L}_{\text{vec}}$; we assume $\mathbf{L}_{\text{vec}}$ is formed by stacking the columns of $\mathbf{L}$ (i.e, the first 3 entries of $\mathbf{L}_{\text{vec}}$ come from the first column of $\mathbf{L}$, and so on).

The demagnetizing energy $E = E(\mathbf{L})$ of the cuboid is a quadratic function of $\mathbf{L}$ which we write in the form

$$E(L) = -\frac{1}{2}\mu_0 l^5\,\mathbf{L}_{\text{vec}}\cdot\mathcal{E}\cdot\mathbf{L}_{\text{vec}}$$

Here, $\mathcal{E}$ is a symmetric 12×12 matrix; the factor l^5 is introduced to make $\mathcal{E}$ dimensionless. The entries in $\mathcal{E}$ do not depend on l but only on the aspect ratio of the cuboid ($1.7 \times 1.3 \times 1$ in our case).

To verify our analytical formula, we compute the entries of the matrix $\mathcal{E}$ in two ways:

- Via finite difference micromagnetic simulations with a progressively finer mesh until convergence, producing the reference matrix $\mathcal{E}_{\text{ref}}$.
- By subdividing the cuboid into tetrahedrons and calculating pairwise interactions between the tetrahedrons using our new analytical formula (4.17)–(4.25), producing the test matrix $\mathcal{E}_{\text{test}}$.

To estimate the error of computing $\mathcal{E}_{\text{test}}$, we compute the relative error

$$\eta = \frac{||\mathcal{E}_{\text{test}} - \mathcal{E}_{\text{ref}}||}{||\mathcal{E}_{\text{ref}}||}$$

using the sum-of-squares matrix norm

$$||\mathcal{E}|| = \sum_{i,j} \varepsilon_{ij}^2$$

4.6.2 Test results

The reference matrix $\mathcal{E}_{\text{ref}}$ was computed by performing finite diference simulations with mesh sizes $8 \times 8 \times 8$, $16 \times 16 \times 16$, ..., up to $128 \times 128 \times 128$, and computing the Richardson's extrapolation estimate using the last two steps. The estimated relative error of computing $\mathcal{E}_{\text{ref}}$ (compared to the unknown exact value) was 10^{-9}.

For the computation of the test matrix $\mathcal{E}_{\text{test}}$ via the analytical formula (4.17)–(4.25), we tested several numerical integration rules for the triangle: two fixed-order rules from [50] with orders 3 and 10 (4 and 25 points respectively), two families of symmetric rules [64, 65] with varying number of points, and also as a baseline the repeated 1d Gauss rule (i.e. by applying the 1d Gauss rule to each of the 2 dimensions of the triangle).

The results are shown in Table 4.2, ordered by decreasing relative error η. In general, for a given number of points all rules displayed approximately the same order of accuracy, for example for each of the 4 rules with 25 points the relative error η was $\sim 5 \cdot 10^{-5}$, with the symmetric rule [64] showing slightly better accuracy (especially for the 175 point rule with error $4.0 \cdot 10^{-6}$ vs $1.1 \cdot 10^{-5}$ for the symmetric rule [65]). The most accurate rule considered was the repeated 1d Gauss rule with 6400 points; the number of points is clearly too high to use it in practice, but it does show excellent agreement with the reference result $\mathcal{E}_{\text{ref}}$ obtained from finite difference simulations.

Method	No. of points n	Rel. error η	Ref
Fixed order k=3	4	$3.3 \cdot 10^{-2}$	[50]
Symmetric family 2, k=5	7	$4.9 \cdot 10^{-3}$	[65]
Symmetric family 1, k=5	7	$4.9 \cdot 10^{-3}$	[64]
Fixed order k=10	25	$5.8 \cdot 10^{-4}$	[50]
1d Gauss, m=5	25	$5.5 \cdot 10^{-4}$	
Symmetric family 2, k=10	25	$5.1 \cdot 10^{-4}$	[65]
Symmetric family 1, k=10	25	$4.7 \cdot 10^{-4}$	[64]
Symmetric family 2, k=20	79	$4.2 \cdot 10^{-5}$	[65]
1d Gauss, m=10	100	$4.1 \cdot 10^{-5}$	
Symmetric family 1, k=20	85	$3.4 \cdot 10^{-5}$	[64]
Symmetric family 2, k=30	171	$1.1 \cdot 10^{-5}$	[65]
Symmetric family 1, k=30	175	$4.0 \cdot 10^{-6}$	[64]
Symmetric family 2, k=40	295	$1.7 \cdot 10^{-6}$	[65]
1d Gauss, m=80	6400	$1.2 \cdot 10^{-8}$	

Table 4.2: Numerical integration error using the analytical formula (4.17)–(4.25) with various triangle integration rules. For the rules from [50, 64, 65], k is the order of approximation; for the repeated 1d Gauss rule, m is the number of points of the 1d rule.

4.7 Derivations

4.7.1 Derivation of $\mathbf{I}_2$

It is easy to verify that $\nabla\mathbf{R}|\mathbf{R}| = \mathbf{Id}|\mathbf{R}| + \mathbf{R}^{\otimes 2}/|\mathbf{R}|$; using eq. (4.38) for the first index of the rank-2 tensor $\nabla\mathbf{R}|\mathbf{R}|$ we get

$$\begin{aligned}
\mathbf{I}_2(\partial\tau;\mathbf{r}') &= \int_{\partial\tau-\mathbf{r}'} \left[\nabla\mathbf{R}|\mathbf{R}| - \mathbf{Id}|\mathbf{R}|\right] ds \\
&= \int_{\partial\tau-\mathbf{r}'} \left[-\mathbf{n}\times(\mathbf{n}\times\nabla\mathbf{R}|\mathbf{R}|) + \mathbf{n}\otimes(\mathbf{n}\cdot\nabla\mathbf{R}|\mathbf{R}|) - \mathbf{Id}|\mathbf{R}|\right] ds \\
&= \int_{\partial\tau-\mathbf{r}'} \left[-\mathbf{n}\times(\mathbf{n}\times\nabla\mathbf{R}|\mathbf{R}|) + (\mathbf{n}^{\otimes 2} - \mathbf{Id})|\mathbf{R}| + (\mathbf{n}\cdot\mathbf{R})\,\mathbf{n}\otimes\frac{\mathbf{R}}{|\mathbf{R}|}\right] ds
\end{aligned}$$

In the expression $\mathbf{n}\times(\mathbf{n}\times\nabla\mathbf{R}|\mathbf{R}|)$ the cross product acts on the first index of the tensor $\nabla\mathbf{R}|\mathbf{R}|$.

After applying the Stokes' theorem (4.39) and noticing that $\mathbf{n}\cdot\mathbf{R}$ is constant over the facets of a polyhedron, we get the desired formula

$$\mathbf{I}_2 = \sum_{F\in\partial\tau-\mathbf{r}'} \left[-\sum_{(\mathbf{R}_1,\mathbf{R}_2)\in\partial F} (\mathbf{n}\times\mathbf{u})\otimes\lambda_1(\mathbf{R}_1,\mathbf{R}_2) + (\mathbf{n}^2 - \mathbf{Id})\,J_0(F) + (\mathbf{n}\cdot\mathbf{R}_f)\,\mathbf{n}\otimes\mathbf{I}_1(F)\right] \tag{4.26}$$

4.7.2 Derivation of $\boldsymbol{\eta}_1$

To compute $\eta_1(\mathbf{R}_1, \mathbf{R}_2)$, we decompose $\mathbf{R}/|\mathbf{R}|$ into components parallel and orthogonal to $\mathbf{u}$ (eq. (4.38))

$$\frac{\mathbf{R}}{|\mathbf{R}|} = \nabla|\mathbf{R}| = \mathbf{u}(\mathbf{u}\cdot\nabla)|\mathbf{R}| - \mathbf{u}\times(\mathbf{u}\times\frac{\mathbf{R}}{|\mathbf{R}|}) \tag{4.27}$$

Integrating over the line $(\mathbf{R}_1, \mathbf{R}_2)$ and using the fact that for points on the line $\mathbf{u}\times\mathbf{R} = \mathbf{u}\times\mathbf{R}_1$

$$\eta_1(\mathbf{R}_1, \mathbf{R}_2) = \mathbf{u}\int_{\mathbf{R}_1}^{\mathbf{R}_2}(\mathbf{u}\cdot\nabla)|\mathbf{R}|\,dl - \mathbf{u}\times(\mathbf{u}\times\mathbf{R}_1)\int_{\mathbf{R}_1}^{\mathbf{R}_2}\frac{1}{|\mathbf{R}|}\,dl \tag{4.28}$$

After applying the gradient theorem (eq. (4.37)) we get the desired equation

$$\eta_1(\mathbf{R}_1, \mathbf{R}_2) = \mathbf{u}|\mathbf{R}|\Big|_{\mathbf{R}_1}^{\mathbf{R}_2} - \mathbf{u}\times(\mathbf{u}\times\mathbf{R}_1)\,\eta_0(\mathbf{R}_1, \mathbf{R}_2) \tag{4.29}$$

4.7.3 Derivation of $\boldsymbol{\lambda}_1$

We begin by writing the following identities that can be verified by direct differentiation

$$(\mathbf{u}\cdot\nabla)[(\mathbf{u}\cdot\mathbf{R})\mathbf{R}|\mathbf{R}|] = \mathbf{R}|\mathbf{R}| + (\mathbf{u}\cdot\mathbf{R})\mathbf{u}|\mathbf{R}| + (\mathbf{u}\cdot\mathbf{R})^2\frac{\mathbf{R}}{|\mathbf{R}|} \tag{4.30}$$

$$(\mathbf{u}\cdot\nabla)[(\mathbf{u}\cdot\mathbf{R})^2|\mathbf{R}|] = 2(\mathbf{u}\cdot\mathbf{R})|\mathbf{R}| + (\mathbf{u}\cdot\mathbf{R})^3\frac{1}{|\mathbf{R}|} \tag{4.31}$$

Multiplying the second equation by $\mathbf{u}/2$ and subtracting from the first to eliminate $(\mathbf{u}\cdot\mathbf{R})|\mathbf{R}|$ on the right-hand side, we get

$$(\mathbf{u}\cdot\nabla)[(\mathbf{u}\cdot\mathbf{R})\mathbf{R}|\mathbf{R}|] - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot\nabla)[(\mathbf{u}\cdot\mathbf{R})^2|\mathbf{R}|] = \mathbf{R}|\mathbf{R}| + (\mathbf{u}\cdot\mathbf{R})^2\frac{\mathbf{R}}{|\mathbf{R}|} - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot\mathbf{R})^3\frac{1}{|\mathbf{R}|} \tag{4.32}$$

Solving for $\mathbf{R}|\mathbf{R}|$ and integrating

$$\boldsymbol{\lambda}_1(\mathbf{R}_1, \mathbf{R}_2) = \int_{\mathbf{R}_1}^{\mathbf{R}_2}(\mathbf{u}\cdot\nabla)\big[(\mathbf{u}\cdot\mathbf{R})(\mathbf{R} - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot\mathbf{R}))|\mathbf{R}|\big]\,dl - \int_{\mathbf{R}_1}^{\mathbf{R}_2}(\mathbf{u}\cdot\mathbf{R})^2[\mathbf{Id} - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot)]\frac{\mathbf{R}}{|\mathbf{R}|}\,dl \tag{4.33}$$

The first integral can be evaluated using the gradient theorem; to evaluate the second, we note that

$$(\mathbf{u}\cdot\mathbf{R})^2 = |\mathbf{R}|^2 - |\mathbf{u}\times\mathbf{R}|^2 \tag{4.34}$$

Again, the quantity $\mathbf{u}\times\mathbf{R}$ is constant and can be moved outside of the integral, resulting in the following equation for $\boldsymbol{\lambda}_1(\mathbf{R}_1, \mathbf{R}_2)$:

$$\boldsymbol{\lambda}_1 = (\mathbf{u}\cdot\mathbf{R})(\mathbf{R} - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot\mathbf{R}))|\mathbf{R}|\Big|_{\mathbf{R}_1}^{\mathbf{R}_2} - [\mathbf{Id} - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot)]\boldsymbol{\lambda}_1 + [\mathbf{Id} - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot)]|\mathbf{u}\times\mathbf{R}_1|^2\boldsymbol{\eta}_1 \tag{4.35}$$

In this equation $\boldsymbol{\lambda}_1(\mathbf{R}_1, \mathbf{R}_2)$ appears on both sides; to solve for it we need to invert the matrix $2\,\mathbf{Id} - \frac{1}{2}\mathbf{u}(\mathbf{u}\cdot)$. It is straightforward to verify that $(2\mathbf{Id} - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot))^{-1} = \frac{1}{2}\mathbf{Id} + \frac{1}{6}\mathbf{u}(\mathbf{u}\cdot)$ and therefore we obtain the desired formula

$$\boldsymbol{\lambda}_1(\mathbf{R}_1, \mathbf{R}_2) = \big(\frac{1}{2}\mathbf{Id} + \frac{1}{6}\mathbf{u}(\mathbf{u}\cdot)\big)\Big[(\mathbf{u}\cdot\mathbf{R})(\mathbf{R} - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot\mathbf{R}))|\mathbf{R}|\Big|_{\mathbf{R}_1}^{\mathbf{R}_2} + [\mathbf{Id} - \frac{\mathbf{u}}{2}(\mathbf{u}\cdot)]|\mathbf{u}\times\mathbf{R}_1|^2\boldsymbol{\eta}_1(\mathbf{R}_1, \mathbf{R}_2)\Big] \tag{4.36}$$

4.7.4 Vector calculus identities

Gradient theorem:

$$\int_{\mathbf{R}_1}^{\mathbf{R}_2} (\mathbf{u} \cdot \nabla) f(\mathbf{R})\, dl = f(\mathbf{R}_2) - f(\mathbf{R}_1) \tag{4.37}$$

Decomposition of a vector into components parallel and orthogonal to a unit vector $\mathbf{u}$ with $|\mathbf{u}| = 1$:

$$\mathbf{a} = \mathbf{u}(\mathbf{u} \cdot \mathbf{a}) - \mathbf{u} \times (\mathbf{u} \times \mathbf{a}) \tag{4.38}$$

Stokes' theorem (alternative form):

$$\int_S (\mathbf{n}_s \times (\mathbf{n}_s \times \nabla))\, a(\mathbf{R})\, ds = \int_{\partial S} a(\mathbf{R}) \mathbf{u}_l\, dl \tag{4.39}$$

4.8 Summary

We presented a method to compute the energy of the magnetostatic interaction between linearly magnetized polyhedrons. The magnetostatic energy integral (4.3) is computed using a hybrid procedure where four out of six integration steps are performed analytically resulting in a nonsingular 2d integral (4.17) which is then computed numerically.

The method can be used in finite element micromagnetics to compute the demagnetizing energy with a high degree of accuracy (for instance, as a reference value in comparison to fast, less accurate traditional methods such as FEM/BEM). Combined with a suitable long range approximation for the magnetostatic integral (4.3), it can allow an implementation of energy-based fast multpole method (FMM) or tree-code algorithms for the computation of the demagnetizing field.

Numerical testing showed excellent agreement between the hybrid computation using the new analytical formula (4.17)–(4.25) and the reference finite difference simulation.

Chapter 5

Ferromagnetic resonance and normal modes

This Chapter describes a proposal for a standard reference problem for the simulation of ferromagnetic resonance; the Chapter is based on the collaborative paper [3], currently in preparation for submission to a journal. The author's contribution included the computation of normal mode frequencies, shapes, damping times, and the FMR spectrum using the eigenvalue approach described in Section 2.5.

5.1 Introduction

Computational micromagnetics is a well developed field that sees widespread use in both modern physics and magnetic device engineering communities [66, 67, 68]. With the advancement of micromagnetic models, simulation techniques, and processing power, the list of topics that can be studied has grown substantially and includes such diverse fields as the spin transfer torque [69] and spin wave dispersion in magnonic crystals [70]. An essential equation in the most of micromagnetic system models [71] is a well defined Landau-Lifshitz-Gilbert (LLG) equation – a differential equation governing the magnetization dynamics. However, this equation can be analytically solved only for a very limited number of systems and, because of that, the complexity of common problems require the use of micromagnetic simulation packages such as OOMMF [72], LLG Micromagnetics [73], Micromagnum [42], and Mumax [74] using the Finite Difference (FD) approach, and Nmag [75] and Magpar [76], employing the Finite Element (FE) approach to spatial discretization. To compare this range of numerical solvers, as well as to evaluate their validity and reliability, NIST's Micromagnetic Modelling Activity Group (μMag) publishes standard problems [77, 78, 79]. Recent additions have included the spin transfer torque [69] and the spin wave dispersion [80] standard problems. In the light of this, it is natural to extend the coverage of standard problems in order to include the FerroMagnetic Resonance (FMR), a technique closely associated with many practical uses ranging from material characterization to the study of spin dynamics.

FMR probes the magnetization dynamics in samples using microwave fields. The absorption of applied microwave field is at its maximum when the microwave's frequency matches the one of the

studied system's resonant modes. By analyzing the resonance modes as a function of an applied magnetic field, some material parameters, such as the Gilbert damping and magnetic anisotropy constants can be determined [32]. This makes FMR a powerful technique in the characterization of ferromagnetic nanostructures; including measurements of spin pumping [33] and exchange coupling [34]. In a typical experiment, microwaves are directed across the sample using a coplanar waveguide, and their transmission is measured as a function of both external bias field and excitation frequency [35].

In terms of computational micromagnetics, there are at least three methods that can be used to simulate the FMR:

1. Application of a time-dependent periodic sinusoidal magnetic microwave field of fixed frequency f to determine the magnetization precession amplitude in response to the system. If the precession amplitude is small, the power absorption of the microwave field would be small as the excitation's frequency does not couple well to the set of natural frequencies of the system. This method is conceptually simple but computationally very demanding as, for every frequency f, the micromagnetic simulation needs to compute the system's magnetization time evolution after the transient dynamics has been damped and steady magnetization precession is reached. This will only provide one point on the frequency-absorption curve and only a micromagnetic simulation software that supports a time dependent external magnetic field can be used.

2. Ringdown method [36]: the system is perturbed from its equilibrium state by applying a short-lived and sufficiently weak excitation, followed by simulation and recording of the magnetization dynamics. Resonance frequencies and corresponding modes are extracted by performing the Fourier transform on the recorded data. This is an efficient way to determine the eigenmodes of the system.

3. Eigenvalue method [37]: instead of simulating the time evolution of the system's magnetization as in the methods above, the problem is represented as an eigenvalue problem, whose solutions provide the frequencies (eigenvalues) and mode shapes (eigenvectors) of the system. This method requires specialist software that is not widely available.

Our goal is to establish a standard problem to serve as a benchmark against which future simulation tools and computational studies of the FMR can be compared and validated. In this standard problem proposal, we will follow the second (ringdown) method, which is supported by most micromagnetic packages and compare its output with the third (eigenvalue) method. We provide a detailed standard problem description and specification as well as the complete set of computational steps in order to make it easily reproducible and accessible to a wide community. It is hoped that this work will aid the development of micromagnetic simulations of systems undergoing FMR and support and drive experimental efforts.

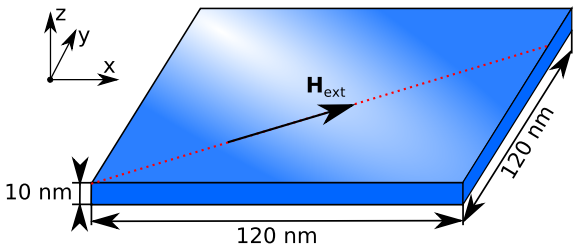

Figure 5.1: Geometry of the thin film sample, showing the static bias field $\mathbf{H}_{\text{ext}}$. The field is slightly off-diagonal to break the symmetry of the system and thus avoid degenerate eigenmodes.

Parameter	Value	Unit
saturation magnetization (M_{s})	800	kA/m
exchange constant (A)	1.3×10^{-11}	J/m
anisotropy constant (K)	0	J/m^3
gyromagnetic ratio (γ^*)	2.21×10^5	m/(As)
Gilbert damping (α), relaxation	1.0	
Gilbert damping (α), dynamic	0.008	
DC bias field magnitude ($\lvert\mathbf{H}_0\rvert$)	80	kA/m
DC bias field direction ($\mathbf{e}$), relaxation	[1, 0.715, 0]	
DC bias field direction ($\mathbf{e}$), dynamic	[1, 0.7, 0]	

Table 5.1: External magnetic fields and material (permalloy) parameters used. Where these change between the initial relaxation stage of the simulation, and the subsequent dynamic stage, both values are shown.

5.2 Selection and definition of standard problem

5.2.1 Problem definition

We choose a thin film permalloy cuboidal sample measuring $120 \times 120 \times 10\,\text{nm}^3$, as shown in Fig. 5.1. The choice of a cuboid is important as it ensures that the finite difference method employed by OOMMF does not introduce errors due to irregular boundaries that cannot be discretized well [14]. On the other hand, we choose the thin film geometry to be thin enough so that the variation of magnetization dynamics along the out-of-film direction can be neglected, and the demagnetization field keeps the magnetic moments in-plane. Material parameters are shown in Table 5.1. An external magnetic bias field $\mathbf{H}_{\text{ext}}$ with $H_{\text{ext}} = 80\,\text{kA/m}$ magnitude is applied along $\mathbf{e} = (1, 0.715, 0)$ direction (at 35.56° to the x-axis), i.e. $\mathbf{H}_{\text{ext}} = H_{\text{ext}}(\mathbf{e}/|\mathbf{e}|) = (65.1, 46.5, 0)\,\text{kA/m}$ as shown in Fig. 5.1. We do not choose the external magnetic field direction along the sample diagonal in order to break the system's symmetry and thus avoid degenerate eigenmodes.

First, we initialize the system with a uniform out-of-plane magnetization $\mathbf{m}_0 = (0, 0, 1)$. The system is allowed to relax for 5 ns, which was found to be sufficient time to obtain a well-converged equilibrium magnetisation configuration. We refer to this stage of simulation as the *relaxation*

stage, and its final relaxed magnetization configuration is saved to serve as the initial configuration for the next *dynamic stage*. Conceptually, what is required to find the relaxed state is to minimize the system's energy in the presence of an external magnetic bias field, taking into account exchange and demagnetisation energy contributions. We note that there are other ways of obtaining this configuration, including energy minimization (as for example supported by OOMMF), or solution of the LLG without the precession term (as supported by Nmag). Because we want to use a well defined method that is supported by all simulation tools, we minimize the system's energy by integrating the LLG equation with a large, quasistatic Gilbert damping $\alpha = 1$ for $5\,\text{ns}$. The use of any of these methods is expected to lead to the same relaxed equilibrium magnetization configuration.

In the next step (*dynamic stage*), a simulation is started using the equilibrium magnetization configuration from the relaxation stage as the initial configuration. Now, the direction of an external magnetic field is altered to $\mathbf{e} = (1, 0.7, 0)$, i.e. $\mathbf{H}_{\text{ext}} = H_{\text{ext}}(\mathbf{e}/|\mathbf{e}|) = (65.5, 45.8, 0)\,\text{kA/m}$. This corresponds to a rotation of the bias field to 35° to the x-axis. Due to the change in x and y components of the external magnetic field, the initial magnetization configuration is now out of equilibrium. Consequently, the system tends to relax towards the lowest energy configuration in the presence of new external magnetic field. This simulation stage runs for $T = 20\,\text{ns}$ and time- (and space-) resolved magnetization y-component $M_y(t)$ is recorded every $\Delta t = 5\,\text{ps}$. The Gilbert damping in this simulation stage is reduced to $\alpha = 0.008$. Using the recorded data, Fourier transform is performed to produce the FMR spectrum and obtain eigenfrequencies (and the eigenmodes). Spatially resolved transformations allow examination of the shapes of the modes (see Section 5.2.3). Simulation parameters for both stages of the simulation are given in Tab. 5.1.

5.2.2 Problem Selection

In this section, we address the selection criteria for the standard problem, and give an explanation of how each is met within the proposed framework:

1. *Initial magnetization configuration*. This standard problem is defined in two stages: (i) relaxation stage and (ii) dynamic stage. In terms of the initial magnetisation configuration, in the relaxation stage, all simulation tools are expected to produce the same uniform out-of-plane $\mathbf{m}_0 = (0, 0, 1)$ initial magnetisation configuration. On the other hand, the dynamic stage uses the resulting relaxed equilibrium state from the first stage as an initial configuration. Therefore, the dynamic stage depends on the quality of relaxed magnetisation configuration from the relaxation stage, which may vary between different simulation tools.

2. *Excitation of system.* The perturbation or excitation field must, apart from being reproducible, be sufficiently large to excite magnetization dynamics, yet be small enough so that the system remains in the linear regime. This is achieved by altering the direction of the bias field, as a simple practical approach that does not require time-dependent applied fields. The power spectrum obtained is specific for the chosen excitation, and thus the excitation is a key part of the problem definition. All simulations tools, even the ones that do not support

time-dependent external magnetic fields, are expected to be able to excite the system in the same manner.

3. *Computation time*. Standard problems, apart from being simple and reproducible, require as short as possible computation time. In micromagnetic simulations, the computational time depends mostly on the number of degrees of freedom in the discretized problem. Accordingly, the spatial discretization of $5\,\mathrm{nm}$ is chosen as a balance between computational time and accuracy. Although the second simulation stage is performed with realistic Gilbert damping value $\alpha = 0.008$ over a limited simulation time, in the first (relaxation) stage, we set $\alpha = 1$ to ensure the magnetization reaches a well converged state within the allotted time.

4. *Verification of results*. Ideally, results should be verified against other methods of obtaining them. In this work, we use different simulation packages (including finite difference and finite element discretization schemes) that have been developed by different groups. Furthermore, we use a completely different computational (eigenvalue) based method to obtain the power density spectrum and excited normal modes separately. The magnetization data is saved as a function of time and space, then transformed into FMR data.

5.2.3 Data Analysis

We outline two different ways to compute the power spectrum of the simulated system.

5.2.3.1 Spectrum $S_y(f)$ of spatially averaged magnetization

In this case, an observable we use is the spatially averaged magnetization $\langle \mathbf{M} \rangle_{\mathbf{r}}(t)$, because it is easily accessible in all known simulation tools. Using a discrete Fourier transform [81], we can obtain the power spectrum of the average magnetization in the frequency domain. As the dynamic simulation progresses, at uniform time steps t_k, we record the spatially averaged magnetization $\langle \mathbf{M} \rangle_{\mathbf{r}}(t_k)$, where $t_k = k\Delta t$ with $\Delta t = 5\,\mathrm{ps}$, and $k = 1, 2, \ldots, N$, with $N = 4000$ being the number of time steps. However, we only consider the y-component of spatially averaged magnetization $\langle M_y \rangle_{\mathbf{r}}(t_k)$ to compute the power spectrum $S_y(f) = |\mathcal{F}_y(f)|^2$ using

$$\mathcal{F}_y(f) = \sum_{k=1}^{N} \langle M_y \rangle_{\mathbf{r}}(t_k) e^{-i2\pi f t_k}. \tag{5.1}$$

According to the chosen parameter values, the sampling frequency is $f_\mathrm{s} = 1/\Delta t = 50\,\mathrm{MHz}$, which implies that the maximum frequency that can be sampled (Nyquist frequency) is $f_\mathrm{N} = 2f_\mathrm{s} = 100\,\mathrm{GHz}$. This method requires that the discrete Fourier transform is performed only once in order to compute the power spectrum $S_y(f)$ of the average magnetization.

5.2.3.2 Local power spectrum and $\tilde{S}_y(f)$

Equation (5.1) uses the spatially averaged magnetization to compute its frequency spectrum. Following McMichael and Stiles approach [36], we compute a local power spectrum over the extent

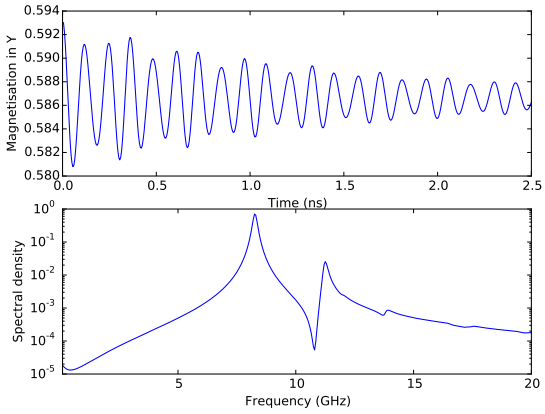

Figure 5.2: (a): Spatially averaged y-component of the magnetization, $M_y(t)$, as determined by ringdown method in by OOMMF. (b) Fourier transform of this data, calculated using Eq. 5.1.

of the sample and introduce a second method to compute the system's power spectrum. In contrast to the first method, this requires computation of discrete Fourier transforms at all spatial sampling points.

We analyze $n = n_x n_y$ scalar time-dependent signals: for every recording time t_k we sample the magnetization on a two-dimensional grid of positions $\mathbf{r}_{m,p}$ where n_x and n_y are the number of sampling points in x and y directions, respectively. More precisely, $\mathbf{r}_{m,p} = ((m - \frac{1}{2})\frac{L_x}{n_x}, (p - \frac{1}{2})\frac{L_y}{n_y}, 2.5\,\mathrm{nm})$ with $m = 1, 2, \ldots, n_x$, $p = 1, 2, \ldots, n_y$, and $L_x = L_y = 120\,\mathrm{nm}$. In the remainder of this work, we have used $n_x = 24$ and $n_y = 24$. For simplicity and generality, we label the sampling points $\mathbf{r}_{m,p}$ as $\mathbf{r}_j$, with $j = 1, 2, \ldots, n_x n_y$.

Now, we can compute the *local* power spectrum

$$S_y(\mathbf{r}_j, f) = |\mathcal{F}_y(\mathbf{r}_j, f)|^2 \tag{5.2}$$

for each of the recorded signals (i.e. for each position $\mathbf{r}_j$), with

$$\mathcal{F}_y(\mathbf{r}_j, f) = \sum_{k=1}^{N} M_y(\mathbf{r}_j, t_k) e^{-i2\pi f t_k}. \tag{5.3}$$

Now, by averaging the local power spectra $S_y(\mathbf{r}_j, f)$, we obtain

$$\tilde{S}_y(f) = \frac{1}{n_x n_y} \sum_{j=1}^{n_x n_y} S_y(\mathbf{r}_j, f). \tag{5.4}$$

Both entities $S_y(f)$ and $\tilde{S}_y(f)$ are shown in Fig. 5.3 in log-scale, and strong resonance peaks are observed at $f_1 = 8.1\,\mathrm{GHz}$ and $f_2 = 11.1\,\mathrm{GHz}$. This method allows us to spatially resolve information about the normal modes of the system.

5.2.3.3 Phase information

In order to understand the precession of a resonance mode q across the extent of a thin film, we need to extract the phase information from the spatially resolved Fourier transform. We start with the

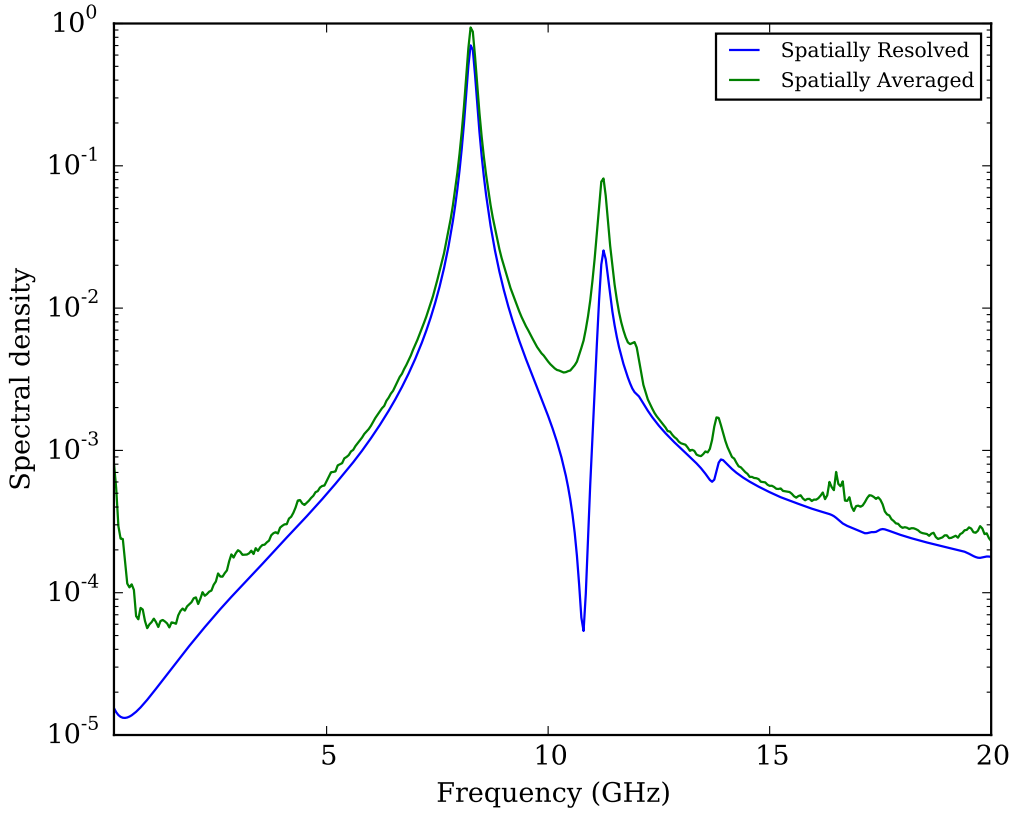

Figure 5.3: Power spectrum S_y from equation 5.1 (green line) and $\tilde{S}_y$ from equation 5.4 (blue line) from ringdown method in OOMMF.

complex Fourier coefficient $\mathcal{F}_y(\mathbf{r}_j, f_q)$ which represents the contribution of the frequency f_q to the time series of the magnetization y-component $\mathbf{M}_y(\mathbf{r}_j, t)$ of the magnetization dynamics at position $\mathbf{r}_j$. In our discrete Fourier transform, we have a set of N complex Fourier coefficients $\mathcal{F}_y(\mathbf{r}_j, f)$ at discrete frequencies f_k. The Fourier coefficient modulus contains the information about the amplitude, whereas its argument contains the phase information. Consequently, the information about the resonance mode q phase can be extracted as the complex Fourier coefficient argument as a function of position $\mathbf{r}_j$, which allows us to identify the relative phases between different spatial domains in a normal mode.

5.3 Results and discussion

5.3.1 Standard Problem Simulation Results

Figures 5.2 and 5.3 show the main results from the standard problem, as outlined in section 5.2.1, obtained using the OOMMF simulation tool. Time evolution of the average magnetisation y-component for the first $2.5\,\text{ns}$ of dynamic stage is shown in Fig. 5.2 (a), and the associated ferromagnetic resonance spectrum (Fourier transform of $\langle M_y \rangle_{\mathbf{r}}(t)$ over the entire $20\,\text{ns}$ dynamic simulation) is shown in Fig. 5.2 (b). Performing the Fourier transform of spatially averaged magnetization produces a slightly different result in comparison to the spatially resolved approach, which is shown in Fig. 5.3.

Using the spatially resolved approach, one can plot $S_y(\mathbf{r}_j, f_q)$ as a function of position $\mathbf{r}_j$ for the normal mode frequency f_q to represent both the power amplitude and phase of the normal mode q, as described in Sec. 5.2.3.3. Figure 5.4 shows the spatially resolved resonance mode at $f_1 = 8.1\,\text{GHz}$ with both the amplitude $|S|(\mathbf{r}_j, f_1)$ and phase information $\arg(S)(\mathbf{r}_j, f_1)$ for x, y and z magnetization components that was calculated from the OOMMF simulation using Eq. (5.4). The magnetisation precession is present in all three directions, with the highest amplitude in y-direction as expected since the largest external bias field perturbation is performed along the y-direction. Figure 5.3 shows that the FT spectrum is dominated by two modes. The low frequency mode extends across the middle of the sample; this corresponds to the mode of uniform precession observed in macroscopic samples.

The largest precession amplitude of the normal mode at $11\,\text{GHz}$ (spatially resolved plot shown in Fig. 5.5) is located at the corners of the sample and is dominated by the demagnetization energy associated with magnetization canting at the sample boundaries. In terms of the normal mode phase representation, an abrupt phase shift occurs as one moves away from the sample corner to the sample center. This normal mode is associated with the particular shape and size of the sample. Note that the precession amplitude in Fig. 5.4 (top) and 5.5 (top) is generally small where the phase changes: these oscillation nodes separate domains that show out of phase precession relative to each other. Similar effects have been observed, for example, in permalloy nanodisks: Guo *et al.* [82] used ferromagnetic resonances force microscopy to spatially resolve the resonance modes. They observed the same mode shapes simulated here, and demonstrated a strong relationship between the size of the disk and the relative strength of the modes. Section 5.4 details the results of

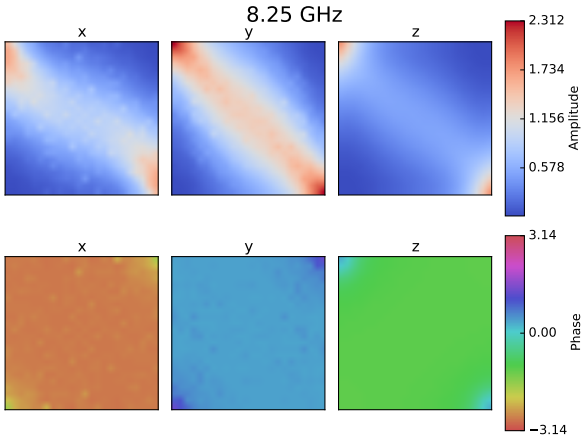

Figure 5.4: Spatially resolved resonance modes in all three Cartesian directions plotted over the extent of the sample at $f_1 = 8.25\,\text{GHz}$ from ringdown method in OOMMF. Top row: logarithmic scale of power spectra with base 10 for x-, y- and z-component, respectively. Bottom row: corresponding phase distributions for three components.

simulations performed without the demagnetization energy contribution (only one resonance is observed, corresponding to a macrospin model of uniform coherent precession).

A resonance mode also exists in the z-direction, $\mathcal{F}_z(\mathbf{r}_j, f)$. The precession of the moments describes an ellipse around the bias field, this has greatest amplitude in the $x - y$ plane, with the component in z being smaller due to the demagnetization field.

5.3.2 Eigenvalue method results

An alternative approach to calculating the normal modes is to linearize the LLG equation for the studied system around its equilibrium state; the normal modes of the resulting linear system of equations can then be determined by solving an eigenvalue problem. This approach does not require running and post-processing of a dynamic micromagnetic simulation, and is thus a good way to check the veracity of the results. A detailed description of this method can be found in [37]; a summary of the method is shown in Section 2.5.1. Now, we use the eigenvalue method to compute the resonance frequencies, the normal mode shapes, as well as the FMR spectrum of the simulated system.

Table 5.2 shows the first fifteen resonance frequencies, calculated using a finite difference discretization with cell size $5 \times 5 \times 5\,\text{nm}$ — this matches the simulation parameters used by OOMMF. The spatial distribution of these modes are plotted in Fig. 5.6. The power density spectrum, and thus the amplitude of each mode excited during the simulation, is dependent upon the perturbation of the system. Using the method described in Sectino 2.5.1.4, we compute the coupling of the used excitation to each mode and reconstruct the spectrum shown in Fig. 5.7, demonstrating an excellent agreement between the ringdown method and the eigenvalue method.

Finally, we show the comparison of the spatial profiles generated by the ringdown and eigenvalue methods. Figure 5.8 shows a comparison of the three lowest frequency modes, demonstrating

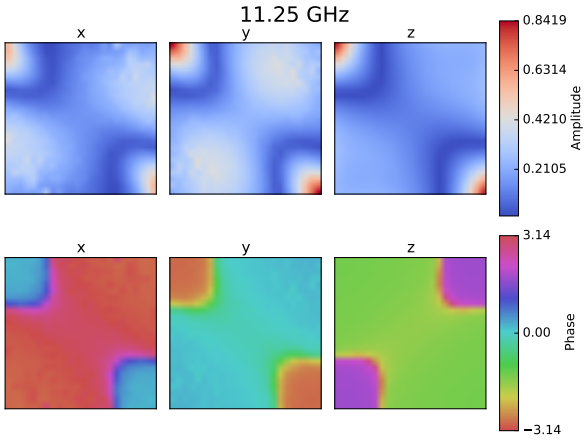

Figure 5.5: Spatially resolved resonance modes in all three Cartesian directions plotted over the extent of the sample at $f_2 = 11.25\,\mathrm{GHz}$ from ringdown method in OOMMF. Top row: logarithmic scale of power spectra with base 10 for x-, y- and z-component, respectively. Bottom row: corresponding phase distributions for three components.

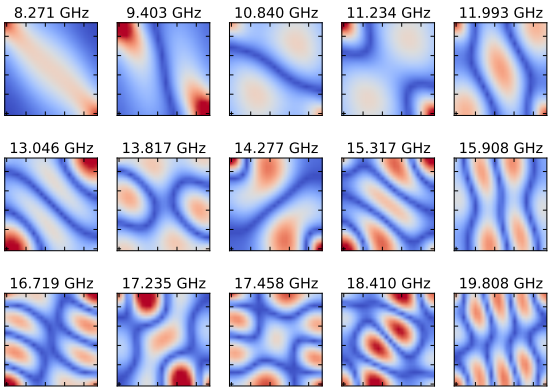

Figure 5.6: The spatial power spectrum of y-component for 15 lowest frequency modes.

excellent agreement for the two modes visible in Fig. 5.3. This agreement gets worse as the frequency of the normal modes increases and their amplitude in the ringdown method decreases; above $14\,\mathrm{GHz}$ the data quality is not sufficient to make a meaningful comparison. Nevertheless, the close agreement of results demonstrates the equivalence of these two approaches.

5.3.3 Falsification Properties

In defining a standard problem, it is useful to investigate how changing the parameters of the simulation will distort the results. This is intended to allow users to isolate inconsistencies within their own simulations when attempting to reproduce the output of this problem.

Mode Frequency (GHz)	Damping Time (ns)
8.270	1.549
9.402	1.639
10.839	1.437
11.233	1.452
11.992	1.401
13.045	1.345
13.816	1.292
14.276	1.253
15.316	1.191
15.907	1.156
16.718	1.126
17.234	1.094
17.457	1.094
18.409	1.030
19.806	0.963

Table 5.2: Frequency and damping time of the 15 lowest frequency modes, calculated using an eigenvalue problem approach.

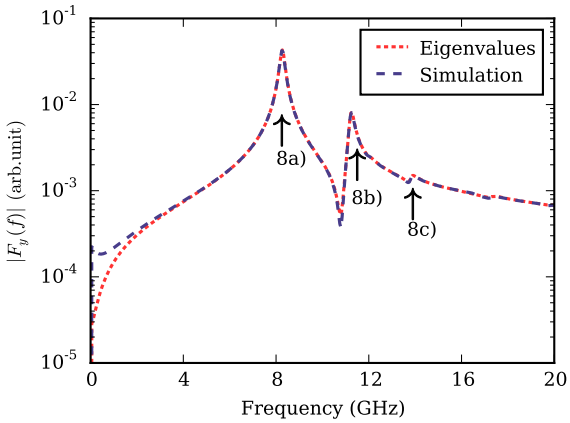

Figure 5.7: Comparison of the spatially averaged resonance spectrum ($|F_y(f)|$, from equation 5.1) calculated by simulation in OOMMF (dotted blue line) and from the eigenvalue problem formulation (solid red line). Excellent agreement is observed over the whole range, although the peak heights are slightly different. Arrows denote the positions of the modes plotted in Fig. 5.8

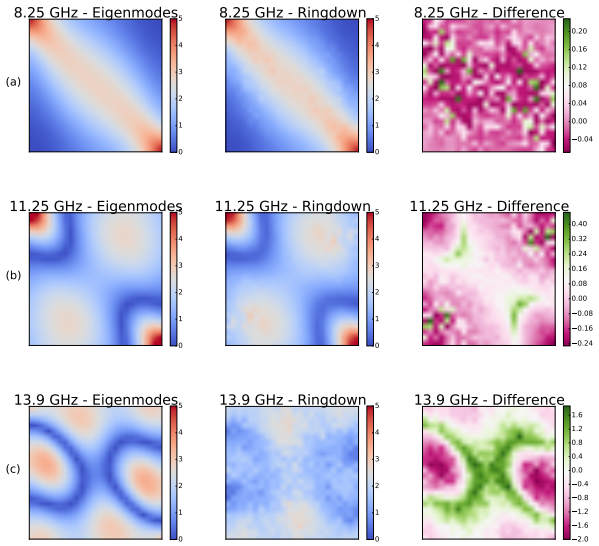

Figure 5.8: Comparison of the spatially resolved power spectrum given by the ringdown method (from OOMMF)and the eigenvalue problem for the y-component of the 3 lowest frequency modes (*top row:* $8.25\,\text{GHz}$, *middle row:* $11.25\,\text{GHz}$, *bottom row:* $13.9\,\text{GHz}$). Excellent agreement is observed for $8.25\,\text{GHz}$ and $11.25\,\text{GHz}$. The agreement gets worse as the amplitude of the mode generated by ringdown decreases, leading to a less well resolved spatial plot.

5.3.3.1 Damping Parameter

The magnitude of the Gilbert damping parameter during ringdown method determines the time taken for the system to reach its stable configuration. However, this did not affect the resonance frequencies produced by the Fourier transform, except in the strongly damped case where $\alpha \geq 0.1$. Figure 5.9 shows the power spectrum produced by the simulation for $\alpha = 10^{-1}$, 10^{-2}, 10^{-3} and 10^{-4}. As the damping parameter is decreased, the peaks become narrower and taller as expected. For the highest damping the spectrum is heavily suppressed, showing only two broad features, with the $11.25\,\text{GHz}$ mode barely visible above the tails of the $8.25\,\text{GHz}$ mode. In this case the system is approaching overdamping; if we set the higher $\alpha = 1$ then no precession occurs and the Fourier transform shows no peaks. As damping decreases extra peaks begin to form, for example at $f \approx$ 12 GHz and 13.5 GHz. At lower dampings the intensity of these features increases, but never surpasses that observed for the two dominant modes.

5.3.3.2 Relaxation Time

Figure 5.10 shows that starting the dynamic simulation stage from an improperly converged configuration from the relaxation stage causes significant instability within the system. Although there are still peaks at the resonance frequencies, there are also many other peaks corresponding to domains aligned in other directions relaxing back to align with the bias field. The frequency of the normal modes does not change, but the strength of the contributions from spurious modes is too large to make a meaningful analysis. The importance of allowing the first stage sufficient time to

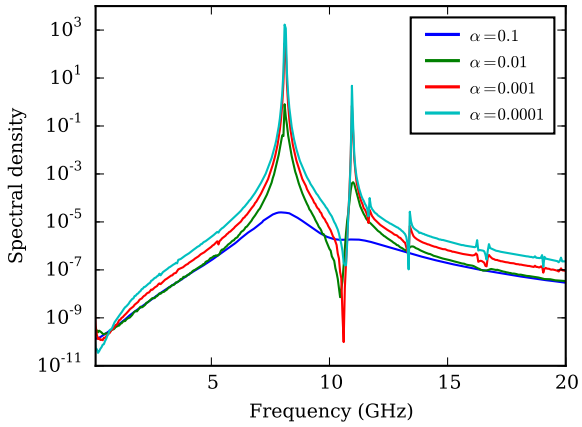

Figure 5.9: Normalized FMR spectrum for systems in the dynamic stage with a range of damping constants. At $\alpha \geq 0.5$ the system is over-damped, not producing resonance modes. As the damping decreases the peaks become taller and sharper.

reach a converged state is clear; the difference occurs because the the system dynamics contains the components that exist as a consequence of the system tending to reach the equilibrium state during the dynamic stage of simulation. The suggested time of 1 ns is a compromise between this requirement and computational brevity. We can see that the curves for 500 ps and 5000 ps produce very similar results.

5.3.3.3 Relaxation Stage Perturbation Angle

Fig. 5.11 shows the effect changing the perturbation angle between bias field in the relaxation and dynamic stages has on the resonance frequencies. Changing the perturbation angle of the bias field changes the amount of energy supplied to the system in the initial excitation, which manifests as a greater area under the power spectrum curve. If the perturbation angle is too small ($< 0.1°$) no peaks are observed above the noise level of the power spectrum. Conversely, if the perturbation angle is too large the system deviates significantly from the equilibrium state, and additional modes form, leading to a distorted power spectrum.While the resonance frequency does not significantly change, the spectrum is eventually dominated by these other features and modes. At such high perturbations both Nmag and OOMMF show a slight drop in resonant frequency. We note that there is a slight asymmetry in the effects of perturbations in the positive and negative y-directions, which we ascribe to the fact that a bias field angle of 35° is not symmetric in this direction.

5.3.3.4 Spatial Discretization

The effects of changing the cell size are shown in Fig. 5.12. Decreasing the resolution of the mesh (increasing the size of the tetrahedra in FE or cuboids in FD) causes the divergence between FD and FE codes. This is to be expected, as the differing approach to calculation of demagnetization field is one of the key differences between the two approaches. In OOMMF the frequency of the

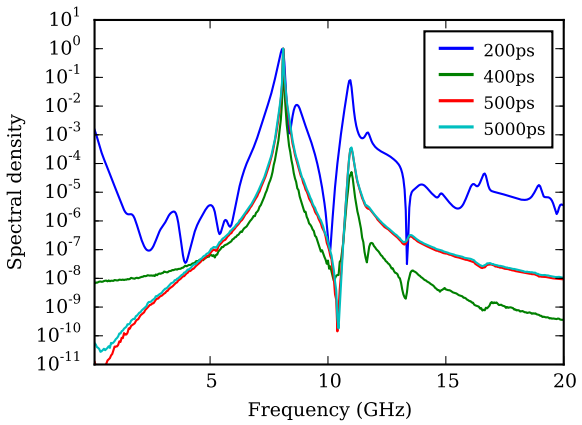

Figure 5.10: Normalized FMR spectrum as calculated for systems entering the dynamic stage after varying the time in the relaxation stage. Allowing more time to relax leads to a lower amplitude, less noise and more well-defined peaks.

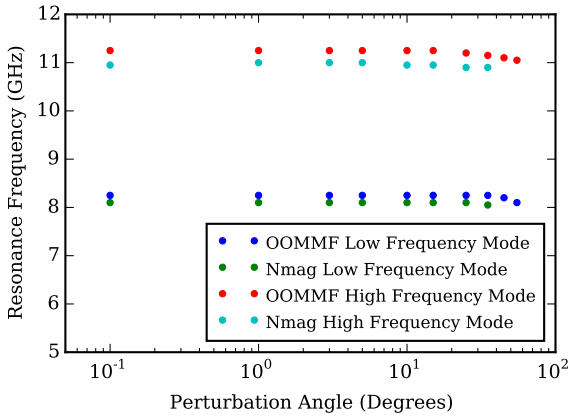

Figure 5.11: Changes to the resonance frequency of the main modes in the FMR spectrum as the magnitude of the initial perturbation is altered. The FD method is relatively unaffected by this change for small angles, at high angles the spectrum becomes noisy and resonance frequency drops. FE behaves in the same way at high angles, but also shows a decrease in resonance frequency for small perturbations.

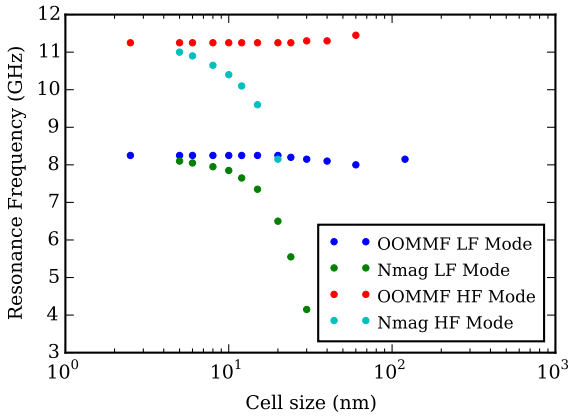

Figure 5.12: Location of the main resonance modes in the FMR spectrum as a function of resolution of the mesh. The FE method shows greater deviation from standard results with changes to the parameters, due to its more sensitive handling of demagnetization effects.

low frequency mode decreases, while the frequency of the main edge mode increases. This comes about due to changes in relative importance of demagnetization effects from the edge of the sample, with less nodes near the boundaries the sample becomes more like an idealized infinite thin film.

In Nmag, however, the frequencies of both normal modes drop sharply. When the elements in the mesh are approximately four times larger than they are in the standard simulation, the spectrum becomes too noisy to identify peaks. At the same time the reduced resolution of the spatial plots makes it difficult to determine the character of the modes. It is worth noting that the frequency of the low frequency mode appears to be approaching the value obtained in simulations without demagnetization (see Sec. 5.4 for more details) before it disappears.

In both codes, high frequency features are suppressed with increasing element size. These correspond to higher order modes that cannot form if there are too few elements to support their spatial variation. It is well known that choice of an appropriate mesh discretization is crucial in computational micromagnetics, an aphorism that is well supported by these results. The deviation of resonant frequency with mesh resolution therefore suggests that a resolution comparable to the exchange length in permalloy ($\sim 5\,\mathrm{nm}$) is indeed the most appropriate.

5.3.4 Comparison of Simulation Methodologies

For the standard problem defined above, the deviations between finite difference and finite element methods for both the resonance frequencies are significant, as shown in Fig. 5.13. A smaller cell size $2 \times 2 \times 1\,\mathrm{nm}^3$ will reduce the deviations significantly, that is, the convergence of finite element method (Nmag) is slower than the finite difference method (OOMMF). The corresponding comparison for the average magnetization (y-component) evolution is shown in Fig. 5.14. It is obvious that two micromagnetic package produce different simulation results when the cell size is $5 \times 5 \times 5\,\mathrm{nm}^3$. Note that in the case of the tetrahedra used in the finite element method this means that space is divided into 6 tetrahedra that together form a cube of dimensions $5 \times 5 \times 5\,\mathrm{nm}^3$.

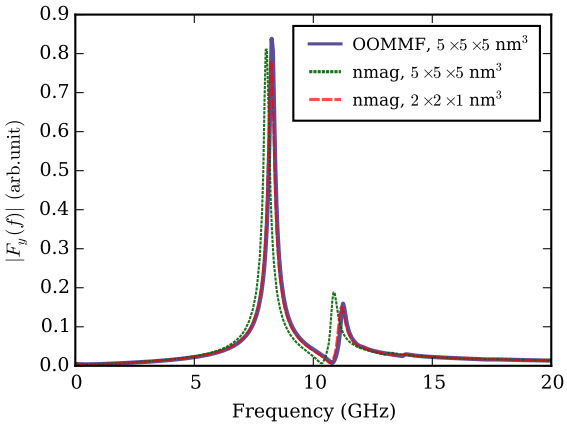

Figure 5.13: A comparison of the resonance spectrum $|F_y(f)|$ between OOMMF and Nmag with different resolutions of the mesh. A cell size of $5 \times 5 \times 5\,\mathrm{nm}^3$ is used for OOMMF.

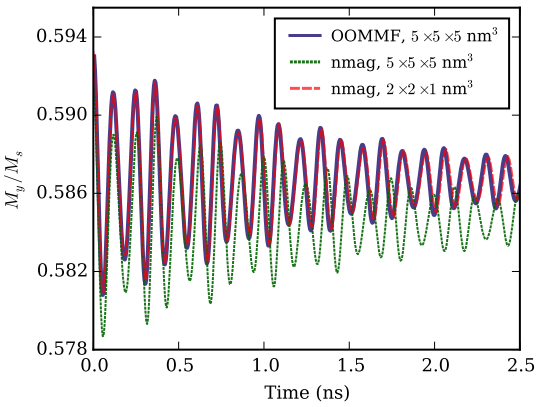

Figure 5.14: Comparison of average magnetization (y-component) evolution between OOMMF and Nmag with different resolutions of the mesh.. Note the phase shift that develops between different spatial resolutions in Nmag, corresponding to a different mode frequency in Fig. 5.13.

This behaviour can be understood by considering the different methods to compute the effective field based on the different approaches to spatial discretization employed by finite difference and finite element packages. Excellent agreement is obtained when a finer mesh discretization is used in Nmag.

Computing the demagnetization field is a computationally expensive step using the FE method (FEM), converging much more slowly than when the FD method (FDM) is used. In FDM the computation takes place at the center of a series of cuboids used to build the sample, while in FEM it takes place at the nodes of the mesh tetrahedra. While this gives significantly better approximations to irregular shapes than the cuboids, it is problematic when the values of the demagnetization tensor sharply vary. If the mesh is not fine enough to accurately resolve the change the effective fields will be calculated incorrectly, and spurious results will be produced. In this simulation, the error arises from contributions from the top and bottom surfaces of the film, and a fourfold increase in resolution in the z-direction brings the FDM simulations into agreement with the FEM, at the

cost of significantly increasing the runtime. This problem could also be alleviated through the use of a spatially varying mesh density, placing more mesh nodes in the regions near the surfaces to accurately sample the demagnetization tensor.

5.4 Simulations without Demagnetization

In this section, we show the results of this standard problem in a setup where only the exchange and the applied Zeeman effective fields are considered. In particular, the demagnetization energy has been ignored. These results are therefore directly comparable with those presented in Fig. 5.13. Figure 5.15 shows the power spectrum of a simulation carried out with demagnetization effects disabled in OOMMF and Nmag. The data has been obtained using the ringdown method. It can be seen that the two packages are in excellent agreement, producing only one mode at 2.8 GHz. In the absence of the demagnetization energy, we obtain this single mode corresponding to coherent precession of the magnetization as shown in Fig. 5.15. This matches the result from the Kittel equation for a material when demagnetization energy contribution is neglected, for which [83]:

$$f = \frac{\gamma}{2\pi} \times \mu_0 \times H_{\text{applied}} \tag{5.5}$$

yields $f = 2.81$ GHz.

As the simulation starts from a uniform, well-converged state only the lowest order, uniform, mode is observed. Modes located at the edge of the sample are suppressed due to the absence of demagnetization.

As discussed in Sec. 5.3.4 differences can arise between simulations performed using the finite difference and finite element approaches due to their handling of demagnetization effects at the film boundaries. The data above shows that both approaches produce very similar spectra in the absence of demagnetization effects. We stress that this information is presented for comparative purposes only - it does not have physical meaning. Running simulations without demagnetization is, however, a useful tool in the debugging process or to analyze specific effects without the additional complications of magnetostatic energy.

5.5 Nmag tolerances

By analysing the time evolution of the average magnetization z-component, obtained by running Nmag simulation with default time integration tolerances, we observe that a numerical noise is present after approximatelly 0.8 ns. By simply performing the Fourier transform on this data, this numerical noise can be interpreted as a particular eigenmode of certain frequency. Because of that, although this does not affect any results presented in this work, we provide the improved demagnetization field computation:

```
1 ksp_tols = {"DBC.rtol":1e-7,
2              "DBC.atol":1e-7,
3              "DBC.maxits":1000000,
4              "NBC.rtol":1e-7,
```

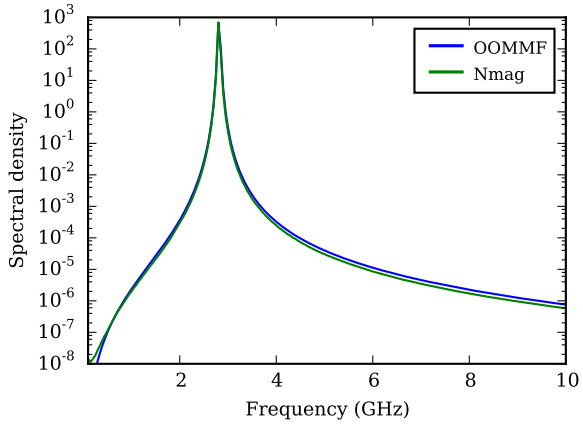

Figure 5.15: Power spectrum for the proposed standard problem with the demagnetization field disabled. This removes all but one peak, wherein the entire sample is in resonance together. Finite element and finite different codes produce the same result under these conditions.

```
            "NBC.atol":1e-7,
            "NBC.maxits":1000000,
            "PC.rtol":1e-3,
            "PC.atol":1e-6,
            "PC.maxits":1000000}
```

and time integration tolerances:

```
sim.set_params(stopping_dm_dt=0.0, ts_abs_tol=1e-7, ts_rel_tol=1e-7)
```

in the dynamic simulation stage. The improved tolerances remove the numerical noise from the average magnetization time evolution, but increases the running time.

5.6 Summary

A standard problem for micromagnetic simulations of ferromagnetic resonance in a thin film has been introduced. FMR is a technique that is used widely. While micromagnetic simulations are able to provide insightful analysis and prediction of FMR experiments, it is not trivial to conduct those simulations. With this chapter, we provide step by step instructions and specific parameters and results that can be used to validate simulation tools before they are applied to new problems.

We provide performance data from two popular micromagnetics packages (OOMMF and Nmag), thus providing data for the deviations that can be expected between different discretization and computation strategies. This standard problem will serve as a good introduction to the procedures involved and would also allow benchmarking and testing of new simulation packages.

Chapter 6

Joule heating in nanowires

In this Chapter, we study the effect of Joule heating from electric currents flowing through ferromagnetic nanowires using system configurations from typical spin transfer torque experiments: (thin) silicon nitride membranes, (thick) silicon wafers, and (thick) diamond wafers.

The work described in this Chapter was published in [4]. The author's contribution included the simulations done using ANSYS and the analytical estimate (i.e. Section 6.3.3 and onwards).

6.1 Introduction

Recently, there has been much interest both in fundamental studies of spin torque transfer [84, 85, 86, 87, 88] and in efforts to realize devices such as the race-track memory exploiting the spin torque transfer [89, 90]. In either case, at the present very large current densities have to be used to move domain walls and, more generally, to modify the ferromagnetic patterns. Associated with these large current densities is a substantial amount of Joule heating that increases the temperature of the sample or device. It is a crucial question to understand how strongly the temperature increases as this can affect the observed physics considerably, e.g. [91, 92, 93], and may even lead to a temporary breakdown of ferromagnetism if the Curie temperature is exceeded. The depinning of a domain wall could be due to a strong spin-current torque transfer, or as a result of the extreme heating of the material due to reduced magnetic pinning at elevated temperatures, or due to the intermittent suppression of ferromagnetism.

You *et al.* [94] have derived an analytical expression that allows to compute the increase of temperature for a nanowire (extending to plus and minus infinity in y-direction) with height h (in z-direction) and width w (in x-direction). The nanowire is attached to a semi-infinite substrate (which fills all space for $z \leq 0$). The heating is due to an explicit term $S(x,t)$ which can vary across the width of the wire and as a function of time.

Meier *et al.* [87] use energy considerations to estimate the total amount of energy deposited into the nanowire and substrate system to show that – for their particular parameters – the heating and associated temperature increase stays below the Curie temperature.

In this work, we use a numerical multi-physics simulation approach which allows to determine the temperature distribution $T(\mathbf{r},t)$ for all times t and positions $\mathbf{r}$. Starting from a given geometry and an applied voltage (or current), we compute the resulting current density, the associated heat

generation, and the temperature distribution. While such a numerical approach provides less insight than an analytical approximation, it allows us to exactly determine the temperature distribution for nanowires of finite length, nanowires with constrictions and thin substrates for which the assumption of an infinite thickness is inappropriate. Thus it evinces the limits of applicability of analytical approximations.

While the simulation and analytical techniques used and developed in this work are not limited to ferromagnetic nanowires on substrates, we have chosen materials, geometries and current densities that are typical for spin-torque driven domain wall motion studies in ferromagnetic nanowires in order to illuminate the role of Joule heating in this active research area [84, 85, 86, 87, 88, 89, 90, 91, 92, 93].

Section 6.2 introduces the method underlying the work. Section 6.3 reports results from a number of case studies. Starting from the heating of a nanowire of infinite length without constrictions (Sec. 6.3.1), we introduce a symmetric constriction in a finite-length wire where a cuboidal part of material has been removed (Sec. 6.3.2) to demonstrate the additional heating that results from an increased current density in close proximity to the constriction and in the constriction (as in [95]). Section 6.3.3 studies a nanowire with a notch-like constriction (triangular shape removed from the wire on one side only) that is placed on a silicon nitride substrate of 100 nm thickness (as in [96, 97, 98]). The same system is studied in Sec. 6.3.4 where the silicon nitride membrane is replaced with a silicon wafer with a thickness of the order of 500 μm. A zigzag wire on the same silicon wafer, as experimentally investigated in [86], is studied in Sec. 6.3.5. We simulate a straight nanowire without constrictions placed on a diamond substrate as in [99] in Sec. 6.3.6. In Sec. 6.4.1 we investigate and discuss the applicability of the analytical temperature calculation model of [94], and deduce an analytical model valid for quasi two-dimensional systems such as membrane substrates in Sec. 6.4.2. We briefly discuss free-standing and perpendicular nanowires in Sec. 6.5, before we close with a summary in Sec. 6.6.

6.2 Method

A current density $\mathbf{j}(\mathbf{r}, t)$ can be related to the change of temperature $T(\mathbf{r}, t)$ of a material as a function of position $\mathbf{r}$ and time t using

$$\frac{\partial T}{\partial t} = \frac{k}{\rho C}\nabla^2 T + \frac{Q}{\rho C} = \frac{1}{\rho C}(k\nabla^2 T + Q) \tag{6.1}$$

with k the thermal conductivity (W/(K m)), ρ the density (kg/m^3), C the specific heat capacity (J/(kg K)), $\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$ the Laplace operator, T the temperature (K), and Q a heating term (W/m^3). We assume that there is no heat flow ($\partial T/\partial \mathbf{n} = 0$) at the material-air (and substrate-air) boundary. The Joule heating of a current density $\mathbf{j}$ in an electrical field $\mathbf{E}$ is given by

$$Q = \mathbf{j} \cdot \mathbf{E} = \frac{1}{\sigma}\mathbf{j}^2 \tag{6.2}$$

where σ is the electrical conductivity (S/m = 1/(Ω m)) and we have used $\mathbf{j} = \sigma\mathbf{E}$.

Equation (6.1) becomes trivial to solve if we assume a uniform current density in a slab of one material with constant density, constant thermal conductivity, and constant specific heat capacity (see Sec. 6.3.1). In general, for samples with geometrical features or spatially inhomogeneous material parameters, the problem becomes quite complex and can often only be solved using numerical methods. For the work presented here we have used the simulation software suites ANSYS 12.0 [100], Comsol multi-physics [101], and the Nsim multiphysics simulation library [102] that underpins the Nmag [103] micromagnetic simulation package. All three tools were used for case study 2 (Sec. 6.3.2) and produce identical results for a given desired accuracy within their error tolerance settings. The majority of the other case studies was simulated using ANSYS. We have taken material parameters (see Tab. 6.1) appropriate for room temperature, and have treated each material parameter as a constant for each simulation, *i.e.* here a temperature dependence of the material parameters is not taken into account.

6.3 Simulation results

Case studies 1 and 2 (Sec. 6.3.1 and 6.3.2) investigate a nanowire without a substrate. Case study 3 (Sec. 6.3.3 studies a nanowire on (2d) silicon nitride substrate membrane. Case studies 4 and 5 (Sec. 6.3.4 and Sec. 6.3.5) investigate nanowires on (3d) silicon wafer substrates, and case study 6 (Sec. 6.3.6) reports from a nanowire on a (3d) diamond substrate.

We refer to the substrate as two-dimensional where the wire length is much greater than the substrate thickness (but still carry out numerical calculations by discretising space finely in all three dimensions). We call the substrate three-dimensional if the substrate thickness is much greater than the wire length.

6.3.1 Case study 1: Uniform current density

Initially, we study the extreme case of no cooling of the ferromagnetic conductor: neither through heat transfer to the surrounding air, nor to the substrate and the contacts. This allows to estimate an upper limit of the heating rate and the consequent change in temperature over time.

Assuming a uniform current density $\mathbf{j}$, uniform conductivity σ and initially uniform temperature distribution T, equation (6.1) simplifies to

$$\frac{\mathrm{d}T}{\mathrm{d}t} = \frac{j^2}{\rho C \sigma} \tag{6.3}$$

All the parameters on the right hand side are constant, and thus the temperature T will change at a constant rate of $j^2/(\rho C\sigma)$. As only the ferromagnetic conducting nanowire can store the heat from the Joule heating, the temperature has to increase proportionally to the heating term Q which is proportional to j^2.

Using material parameters for Permalloy ($C = 430\,\mathrm{J/(kg\,K)}$, $\rho = 8700\ \mathrm{kg/m^3}$, $\sigma = 1/(25 \cdot 10^{-8}\ \Omega\mathrm{m}) = 4 \cdot 10^6\ (\Omega\mathrm{m})^{-1}$, $j = 10^{12}\,\mathrm{A/m^2}$), we obtain a change of temperature with time

$$\frac{\mathrm{d}T}{\mathrm{d}t} = \frac{j^2}{\rho C \sigma} = 6.683 \cdot 10^{10}\,\mathrm{K/s} = 66.83\,\mathrm{K/ns}. \tag{6.4}$$

Case Study	Parameter	Value	Unit	Reference
1–4	Py σ^{-1}	25	$\mu\Omega\,\mathrm{cm}$	measured [104] in experiment [97, 98] (case study 3), $R = 280\,\Omega$
5	Py σ^{-1}	42	$\mu\Omega\,\mathrm{cm}$	measured in experiment [86], $R = 5000\,\Omega$
6	Py σ^{-1}	39	$\mu\Omega\,\mathrm{cm}$	measured in experiment [99], $R = 675\,\Omega$
1–6	Py C	0.43	J/(g K)	Ref. [105] p. 252, $6.00\,\mathrm{cal}/(\mathrm{mol\,K})$
1–6	Py k	46.4	W/(K m)	Ref. [106] p. 1140, $T = 300\,\mathrm{K}$
1–6	Py ρ	8.7	g/cm^3	Ref. [107] Tab. I, lattice constant 3.54
3	Si_3N_4 C	0.7	J/(g K)	Ref. [108]
3	Si_3N_4 k	3.2	W/(K m)	Ref. [108]
3	Si_3N_4 ρ	3	g/cm^3	Ref. [108]
4, 5	Si C	0.714	J/(g K)	Ref. [109], $C_\mathrm{p} = 20.05\,\mathrm{J}/(\mathrm{mol\,K})$
4, 5	Si k	148	W/(K m)	Ref. [110] p. I-588
4, 5	Si ρ	2.33	g/cm^3	Ref. [111], Tab. III
6	Diamond C	0.53	J/(g K)	Ref. [112], $C_\mathrm{p} = 6.37\,\mathrm{J}/(\mathrm{mol\,K})$
6	Diamond k	1400	W/(K m)	Ref. [113], Fig. 3, Type Ib
6	Diamond ρ	3.51	g/cm^3	Ref. [114]

Table 6.1: Material parameters used in the simulations (σ electric conductivity, σ^{-1} electric resistivity, C specific heat capacity, k thermal conductivity, ρ density).

The parameters for this case study 1 and all other case are summarised in Tab. 6.1. For the permalloy wire, we have chosen parameters for case studies 1 to 4 to match the experimental work in [96, 97, 98]. Where possible, parameters measured as part of the experiments have been used and have been complemented with literature values (see Tab. 6.1 for details).

The immediate conclusion from this is that the temperature of the sample cannot increase by more than 66.83 K per nanosecond if the current density of $10^{12}\,\mathrm{A/m^2}$ is not exceeded and if the current density is uniform within the whole sample for the chosen material parameters.

A current pulse over 15 nanoseconds has the potential to push the temperature up by just over 1000 degrees Kelvin, and thus potentially beyond the Curie temperature.

The substrate on which the ferromagnetic conductor has been grown will absorb a significant fraction of the heat generated in the conductor, and thus reduce the effective temperature of the magnetic material. The contacts play a similar role. On the other hand, any constrictions will result in a locally increased current density, which – through the j^2 term in Q in Eq. (6.2) and (6.1) – results in significantly increased local heating. We study the balance of these additional heating and cooling terms in the following sections in detail.

6.3.2 Case study 2: Constrictions

The effect of a constriction will vary strongly depending on the given geometry. The resulting current density distribution and spatial and temporal temperature distributions are non-trivial.

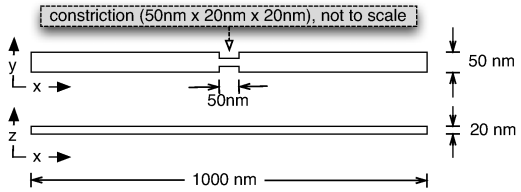

Figure 6.1: Sample geometry for case study 2: slab with constriction (Sec. 6.3.2).

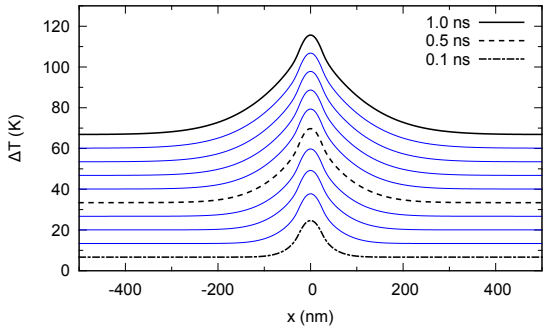

Figure 6.2: Temperature profile $\Delta T(x)$ in the constricted geometry shown in Fig. 6.1 at positions [x,0,0] for $t = 1$ ns and a current density of 10^{12} A/m^2 in the unconstricted ends of the slab.

Fig. 6.1 shows the geometry used for this study (as in [95, 115, 116]): a bar with dimensions $L_\mathrm{x} = 1000$ nm, $L_\mathrm{y} = 50$ nm, and $L_\mathrm{z} = 20$ nm. The origin is located in the center of the slab, i.e. the two opposite corners of the geometry are at $[-500, -25, -10]$ nm and $[500, 25, 10]$ nm. The constriction is placed at the center of the bar, and reduces the dimensions to $L_\mathrm{y}^\mathrm{constrict} = 20$ nm over a length of $L_\mathrm{x}^\mathrm{constrict} = 50$ nm. In comparison to case study 1 (Sec. 6.3.1), the current distribution is non-uniform in this geometry and therefore the local Joule heating and the resulting temperature field will be non-uniform. We thus need the thermal conductivity for $Ni_{80}Fe_{20}$ permalloy[106] $k = 46.4$ W/(K·m) for these calculations because the $\nabla^2 T$ term in equation (6.1) is non-zero.

Fig. 6.2 shows a temperature profile $\Delta T(x)$ through the constricted shape after application of a current density of 10^{12}A/m^2 at ten different times t as a function of position x. We use the notation $\Delta T(x)$ instead of $T(x)$ to indicate that this is the change of the temperature T relative to the initial temperature, for which we assume room temperature (≈ 300 K). The top thick black line shows the temperature distribution after 1 ns, the other 9 lines show earlier moments in time in successive time steps of 0.1 ns.

The peak around $x = 0$ is due to the constriction: the reduced cross section (in the y-z plane) results in an increased current density in the constriction. The Joule heating term (6.2) — which scales proportional to j^2 — is increased accordingly in the constriction. During the 1 ns, diffusion of heat takes place and results in the increase of temperature outside the constriction, and simultaneously a reduction of the speed of increase of the temperature in the constriction. The importance of heat diffusion can be seen by the width of the peak around $x = 0$ in Fig. 6.2.

The increase of temperature ΔT after $t = 1$ ns at the end of the nanowire (i.e. $x = 500$ nm and $x = -500$ nm) is 66.90 K. This is about 0.07 K higher than in the previous case study in Sec. 6.3.1 where a nanowire without a constriction was studied and the corresponding value of temperature

increase is $66.83\,\mathrm{K}$. This difference of 0.07K after 1 ns at the end of the wire originates from the extra heating in the constriction around $x = 0$ and diffusion of this heat through the wire.

We can also estimate the maximum possible temperature increase in the constriction by using equation (6.4) under the assumption that there was no diffusion of heat (*i.e.* $k = 0$), and that we obtain a current density of $j = 2.5 \cdot 10^{12}\mathrm{A/m^2}$ in the constriction (due to the reduced cross section from $L_\mathrm{z}L_\mathrm{y} = 1000\,\mathrm{nm}^2$ to $L_\mathrm{z}L_\mathrm{y}^\mathrm{constrict} = 400\,\mathrm{nm}^2$). The resulting rate for temperature increase in the constriction is here 417 K/ns, and thus much higher than the maximum temperature of 115.6 K that is reached after 1ns when taking into account the diffusion of heat from the constriction into the unconstricted parts of the nanowire. This comparison shows the drastic influence of diffusion of heat on the temperatures in the nanowire.

6.3.3 Case study 3: Nanowire with a notch on a silicon nitride substrate membrane

In this section we investigate a more realistic system following the work by Im *et al.* [96] for which we take into account the heat dissipation through the substrate. Im *et al.* study critical external fields for domain wall pinning from constrictions, and subsequent works [97, 98] study spin-torque driven domain wall motion for this geometry.

Here, we choose the geometry where the domain wall de-pinning field was most clearly defined: a permalloy nanowire with dimensions $L_\mathrm{x} = 5000\,\mathrm{nm}$, $L_\mathrm{y} = 150\,\mathrm{nm}$, $L_\mathrm{z} = 30\,\mathrm{nm}$ (see Fig. 1(b) in [96] and lower left subplot of Fig. 3 therein).

The resistivity of a permalloy thin film strongly depends on its thickness; in this simulation we scale the permalloy resistivity σ so that the total resistance of the device matches the resistance value $\approx 280\,\Omega$ reported [104] for the experiments [97, 98]. This leads to a resistivity $\sigma^{-1} = 25\,\mu\Omega\mathrm{cm}$, in line with published data on permalloy resistivity (Fig. 1(b) in [117], $t = 30\,\mathrm{nm}$). We use the same value for σ in case study 1 (Sec. 6.3.1), case study 2 (Sec. 6.3.2) and case study 4 (Sec. 6.3.4) to allow better comparison between the results.

This nanowire is lithographically defined centrally on a silicon nitride membrane that is 100 nm thick (as in [96]). The membrane used in the experiment [96] was purchased from Silson Ltd [118]. According to the manufacturer the silicon nitride membrane was grown using low pressure chemical vapour deposition (LPCVD). The thermal properties of LPCVD grown silicon nitride films depend on the details of the growth process as can be seen in the range of parameters being cited in the literature [108, 119, 120, 121, 122]. We assume values of $k = 3.2\,\mathrm{W/(m\,K)}$, $C = 0.7\,\mathrm{J/(g\,K)}$, and $\rho = 3\,\mathrm{g/cm^3}$ as in Ref. [108].

For the modelling of the membrane we use the shape of a disk with 0.5 mm radius and 100 nm thickness. Preparation of nanostructures on membranes is required for experiments using synchrotron light, that has to transmit through the sample. Such experiments give access to simultaneous time- and space resolution on the nanometer and the sub-nanosecond scale.[96, 84] Perfect thermal contact between the wire and the silicon nitride substrate disk is assumed. The center of the wire contains a 45 nm wide triangular notch on one side, and the geometry is sketched in Fig. 6.3(b). A current density of $10^{12}\,\mathrm{A/m^2}$ is applied at the ends of the wire over a time of 20 ns, similar to recent experiments such as in Refs. [97, 98, 123].

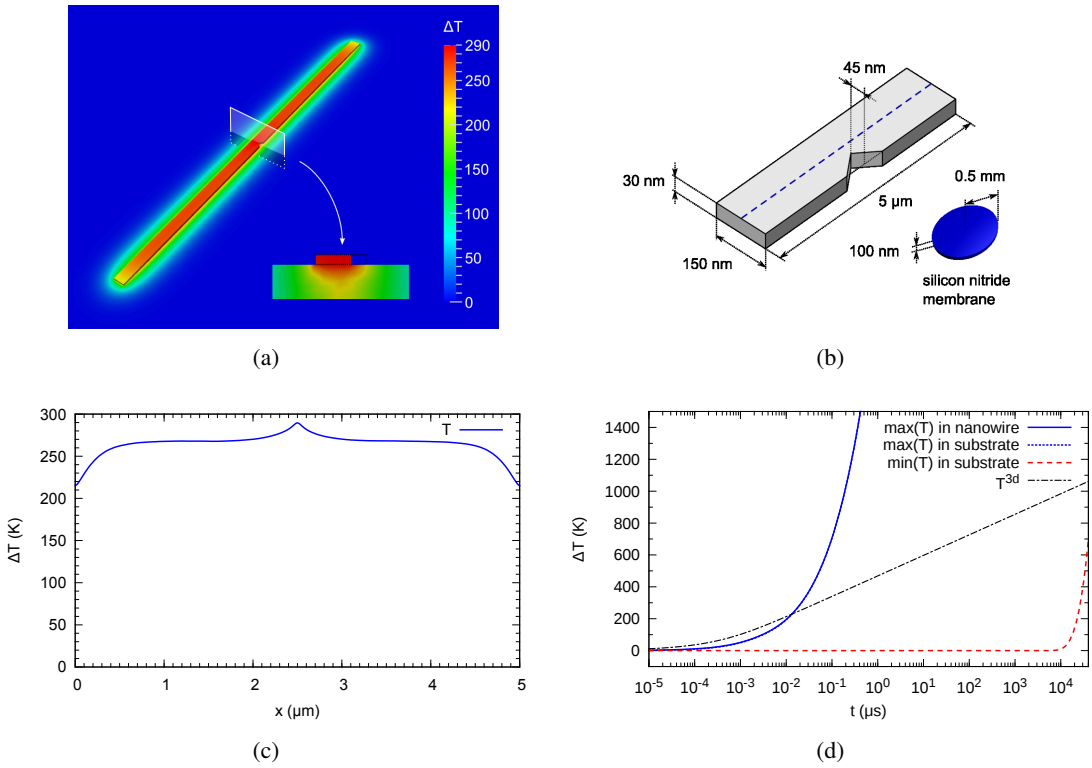

Figure 6.3: Joule heating in a permalloy nanowire with a notch on a silicon nitride membrane [96] as discussed in Sec. 6.3.3. The membrane is modelled as a disk with height 100 nm and radius 0.5 mm. The current density is $j = 10^{12}\,\mathrm{A/m^2}$; for material parameters see Tab. 6.1. (a) temperature distribution $\Delta T(\mathbf{r})$ in the nanowire and substrate after $t = 20\,\mathrm{ns}$, (b) geometry of the model (not to scale) and the plotting path through the nanowire (dashed line) used in Fig. 6.3(c), (c) temperature profile $\Delta T(x)$ after $t = 20\,\mathrm{ns}$ along the length of the wire (following the plotting path shown in Fig. 6.3(b)) with a silicon nitride membrane substrate, (d) maximum (dotted line) and minimum (dashed line) temperature in the silicon nitride membrane and maximum temperature (solid line) in the wire as a function of current pulse length. The solid and dotted line coincide at this scale. The dash-dotted line $T^{3\mathrm{d}}$ is the prediction of the analytical You model discussed in Sec. 6.4.1. See also Sec. 6.4.2 and Fig. 6.7 for further discussion.

Figure 6.3(a) shows an overview of the geometry and the computed temperature distribution after 20 ns. The 30% notch is just about visible on the right hand side of the nanowire halfway between the ends of the wire. The cut-plane shown in the right hand side of Fig. 6.3(a) as an inset shows the temperature distribution in the nanowire in the y-z plane in the center of the constriction (as indicated by a white line and semi-transparent plane in the main plot). Figure 6.3(b) shows the notch geometry in more detail (not to scale).

Fig. 6.3(c) shows the temperature profile $\Delta T(x)$ after 20 ns taken along a line at the top of the nanowire (the same data is encoded in the colours of Fig. 6.3(a) although more difficult to read quantitatively). We see that most of the wire is at an increased temperature around 270 K. The maximum temperature is found at the constriction: as in Sec. 6.3.2 the current density is increased here due to the reduced cross section in the y-z plane.

In contrast to the previous example (Sec. 6.3.2 and Fig. 6.2), the temperature peak at the constriction is less pronounced, as a smaller amount of material is absent and consequently the current density and the associated increase in Joule heating is smaller.

The maximum temperature increase reaches 290 K. The zero level in the simulation corresponds to the temperature at which the experiment is started: if the experiment is carried out at 300 K we expect the maximum temperature to be 590 K after 20 ns, which is below the Curie temperature (≈ 840 K) for Permalloy.[124]

From Figs. 6.3(a) and 6.3(c) we can see that the temperature at the ends of the wire is lower than near the middle: for $x \lesssim 1\,\mu$m and $x \gtrsim 4\,\mu$m temperature decreases towards $\approx 215\,K$ at the ends of the wire. This is due to more efficient cooling through the substrate: at the ends of the wire there is substrate to three sides rather than two as in the middle parts of the wire.

The importance of the substrate in cooling the nanowire can be seen if we use equation (6.4) to compute the temperature after 20 ns for a wire with the same geometry but without the notch and without the substrate: the heating rate is $\mathrm{d}T/\mathrm{d}t = j^2/(\rho C\sigma) = 66.83\,\mathrm{K/ns}$ as in Sec. 6.3.1 because the geometry does not enter that calculation. Without the substrate, we would have a temperature increase of ≈ 1336 K after 20 nano seconds (even without taking the extra heating from the notch constriction into account).

In Fig. 6.3(d), the solid line shows the maximum temperature in the nanowire, the dotted (dashed) line shows the maximum (minimum) value of the temperature in the silicon nitride membrane substrate and the dash-dotted line shows data computed using the model by You *et al.* [94] as a function of time over which the current pulse is applied. Note that the You model has not been derived to be used for such a thin substrate and that the large deviation for large values of t is thus expected. We discuss this in detail in Sec. 6.4.

We can use the maximum wire temperature graph to determine the length of the current pulse that can be maintained until the temperature is pushed up to the Curie temperature, or to the material's evaporation temperature (note that the plot shows the temperature increase since the start of the experiment, not absolute temperature). If the experiment is carried out at room temperature (≈ 300 K), then the Curie temperature (≈ 840 K) is attained with a pulse duration of approximately 60 ns. The material starts to melt and evaporate for $t \gtrsim 300$ ns, once the melting temperature of $\approx 1450\,^\circ$C is exceeded.

Asymptotically, the maximum temperature in the wire is proportional to the logarithm of pulse duration for a 2d substrate and a point-like heating source. This regime is entered when the nanowire and the constriction have been heated up to a steady state, and from there on the temperature in the nanowire increases logarithmically with time while the heat front propagates from the center of the substrate disk towards the disk's boundary.

The maximum temperature in the substrate (dotted line) is assumed at the interface between the nanowire and the substrate, in the location where the nanowire is hottest. Because of the assumption of perfect thermal contact between wire and the membrane substrate, the temperature at the top of the membrane is the same as the temperature at the bottom of the wire. The difference between the maximum wire temperature and maximum substrate temperature thus provides an indication for

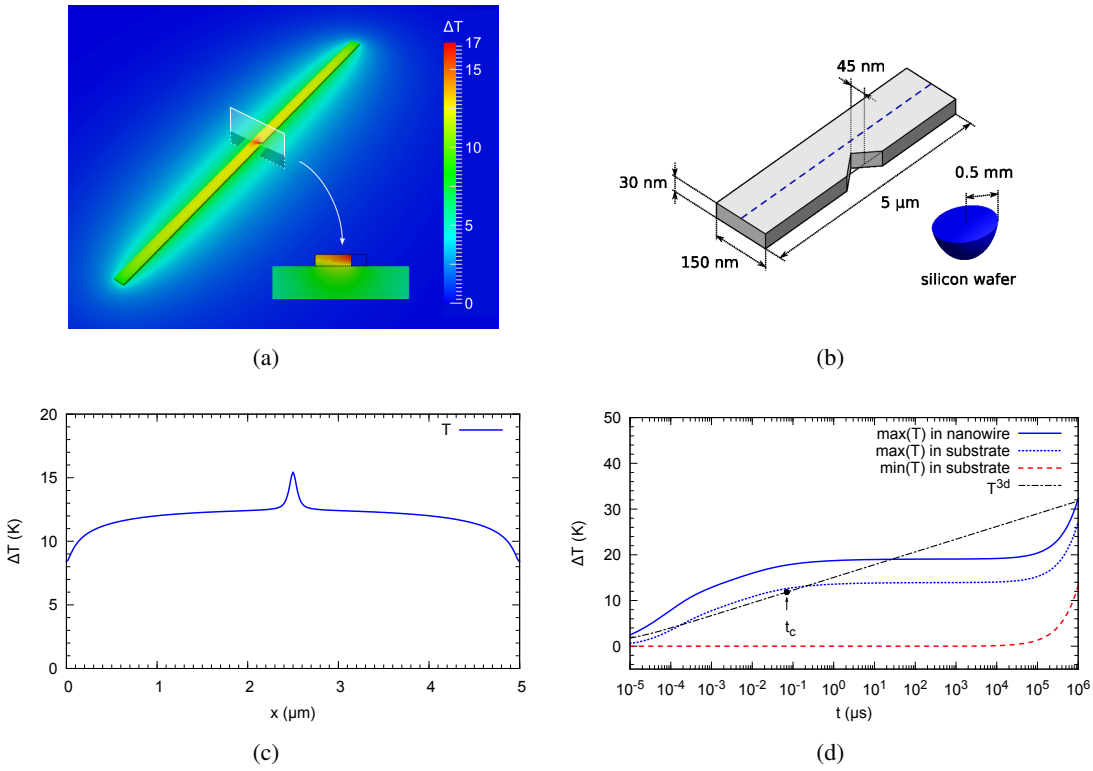

Figure 6.4: Joule heating in a permalloy nanowire with a notch on a silicon wafer as discussed in Sec. 6.3.4 (geometry of wire from Refs. [97, 98]): (a) temperature distribution $\Delta T(\mathbf{r})$ in the nanowire and substrate after $t = 20\,\mathrm{ns}$, (b) geometry of the model (not to scale) and the plotting path (dashed line) used in Fig. 6.4(c), (c) temperature profile $\Delta T(x)$ after $t = 20\,\mathrm{ns}$ along the length of the wire on a silicon wafer substrate, (d) maximum and minimum temsperatures of the silicon substrate, maximum temperature of the wire, and the prediction of the You model $T^{3\mathrm{d}}$ as a function of pulse length; t_c is the characteristic time as described in Sec. 6.4.1. The silicon wafer is modelled as a half sphere (flat side attached to nanowire with notch) with radius 0.5 mm. The current density is $j = 10^{12}\,\mathrm{A/m^2}$ at the end of the wire. For material parameters see Tab. 6.1.

the temperature gradient found in the wire. In Fig. 6.3d the two lines nearly coincide.

The dashed line in Fig. 6.3(d) shows the minimum value of the temperature taken across the combined system of nanowire and substrate. It starts to deviate from zero when the heat front has propagated from the center of the silicon nitride substrate disk to the boundary. In our example, that happens after $\approx 10^4\,\mu$s when the simulated temperature increase of the wire would be over 14,700 K. The asymptotic logarithmic behaviour of the temperature increase in the wire is maintained only until the heat front reaches the boundary of the substrate.

6.3.4 Case study 4: Nanowire with a notch on a silicon wafer substrate

A possible way to increase the maximum pulse length is to use a silicon wafer instead of the silicon nitride membrane, which is typically done in high frequency transport experiments. Silicon is a better conductor of heat than silicon nitride. Furthermore, the extra thickness of the wafer

gives more space for heat to dissipate. In this section, we model the same nanowire with a notch geometry as above in case study 3 (Sec. 6.3.3) but place the nanowire on a silicon half-sphere of radius 0.5 mm instead of a 100 nm thin silicon nitride disk membrane.

Fig. 6.4(c) shows the temperature profile of the nanowire with notch placed on a silicon wafer after a 20 ns pulse (current density at the wire ends is 10^{12} A/m^2). The maximum temperature increase is 17 K and should be compared with Fig. 6.3(c) where the nanowire was placed on a (much thinner) silicon nitride substrate and the maximum temperature increase was 290 K.

Clearly, compared to the (2d) silicon nitride membrane, the (3d) silicon wafer is much more efficient at diffusing heat: the temperature increase is approximately a factor 17 smaller.

Fig. 6.4(d) shows the minimum and maximum temperature for the silicon wafer substrate, and the maximum temperature of the nanowire, analog to Fig. 6.3(d) for the nanowire on the silicon nitride membrane. We can identify three regimes: (i) for small t the maximum temperature in the wire increases approximately proportional to the logarithm of time. (ii) For $1\,\mu\text{s} \lesssim t \lesssim 10,000\,\mu\text{s}$, the maximum temperature stays constant in the wire. This is the (3d) steady-state regime where the heat front propagates from the center of the silicon substrate half-sphere to the surface of the sphere. (iii) Approximately for $t \gtrsim 10,000\,\mu\text{s}$, the maximum wire and substrate temperatures and the minimum substrate temperature start to increase again simultaneously. This indicates that the heat front has reached the surface of the half-sphere shaped wafer substrate, and that the heat from the nanowire cannot be carried away through the propagating heat front anymore.

The maximum temperature increase of ≈ 19 K is reached after $\sim 1\,\mu$s and is maintained until $t = 10\,\text{ms}$, when the heat front reaches the boundary of the wafer and the whole wafer begins to heat up. The temperature gradient from top to bottom of the wire is of the order of 5 K (difference of maximum wire and maximum substrate temperature).

For a system with an infinite 3d substrate we expect in general that the nanowire temperature stays constant once regime (ii) of the heat front propagation has been reached. Regime (iii) would not exist for an infinite substrate, or a substrate that is efficiently cooled at its boundaries.

6.3.5 Case study 5: Zigzag-shaped nanowire on a silicon wafer substrate

In this section we investigate a nanowire geometry and experimental set up as reported in [86] where current densities of $2.2 \times 10^{12}\,\text{A}/\text{m}^2$ are applied for $10\,\mu$s.

Figure 6.5(b) shows the permalloy nanowire consisting of three straight $20\,\mu\text{m}$ segments connected by 45° bends of radius $2\,\mu\text{m}$ (see also Fig. 1 in [86]). The wire is 500 nm wide and 10 nm thick, and is placed on a silicon wafer substrate which is modelled as a silicon half-sphere of radius 0.5 mm as in the previous example in Sec. 6.3.4. For the simulation we scale the resistivity so that the total resistance of the device matches the value 5 kΩ reported in [86]. This value is achieved with resistivity $\sigma^{-1} = 42\,\mu\Omega\text{cm}$, which agrees well with published data on permalloy resistivity (Fig. 1(b) in [117], $t = 10$ nm).

Figure 6.5(a) and 6.5(c) shows the temperature profile after $10\,\mu$s. The maximum temperature increase does not exceed 133 K despite the large current density and the long pulse duration. The ends of the zig-zag wire are significantly cooler (below 80 K) than the center which can be attributed

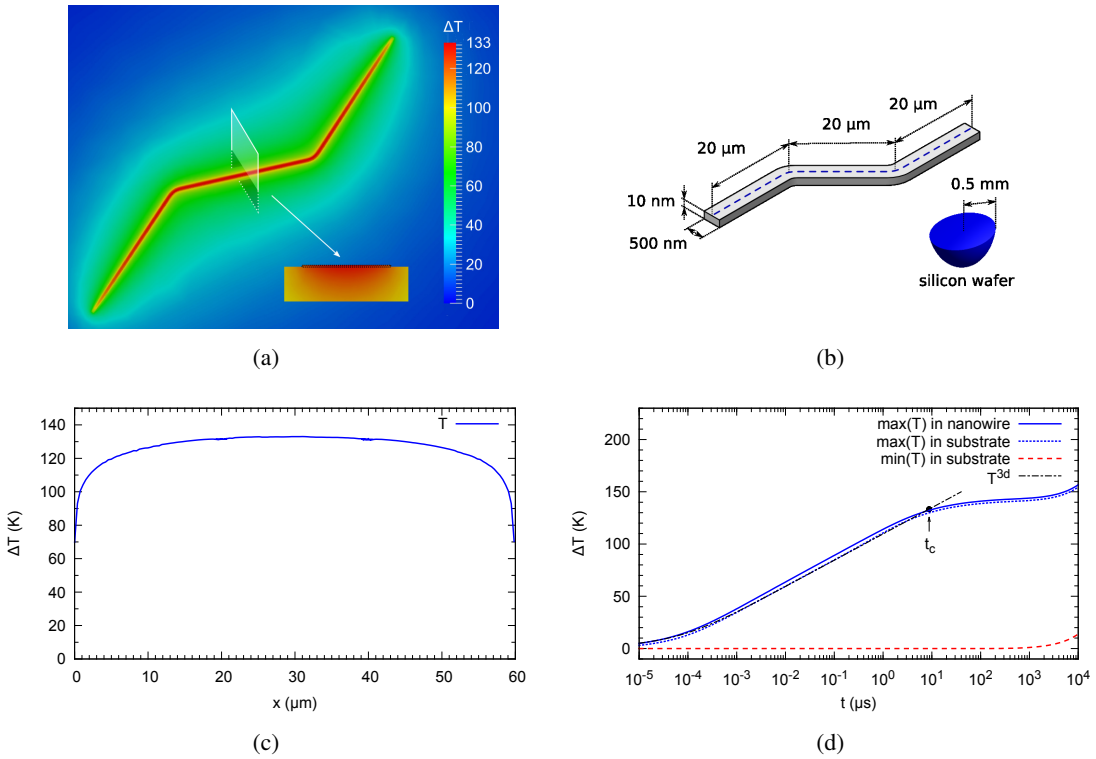

Figure 6.5: Joule heating in zigzag permalloy nanowire on a silicon wafer[86] (Sec. 6.3.5): (a) temperature distribution $\Delta T(\mathbf{r})$ in the nanowire and substrate after $t = 10\,\mu$s, (b) geometry of the model (following Ref. [86]) and the plotting path (dashed) for Fig. 6.5(c), (c) temperature profile $\Delta T(x)$ after $t = 10\,\mu$s along the length of the wire (following the plotting path shown in Fig. 6.5(b)), (d) maximum and minimum temperatures of the silicon substrate, maximum temperature of the wire, and the prediction of the You model as a function of pulse length; t_c is calculated based on the distance between the opposite ends of the wire $L = 56\,\mathrm{nm}$. The current density is $j = 2.2 \times 10^{12}\,\mathrm{A/m^2}$; the resistivity of permalloy $\sigma^{-1} = 42\,\mu\Omega\mathrm{cm}$. For other material parameters see Tab. 6.1.

to more silicon wafer substrate accessible to carry away the heat that emerges from the nanowire.

Figure 6.5(d) shows how the maximum wire temperature (which is located in the middle segment of the zig-zag wire) and substrate minimum and maximum temperatures change over time. As before, we can identify three regimes: (i) initial heating of wire and substrate, (ii) steady state with heat front propagating through substrate and (iii) general heating of wire and substrate when heat front reaches the substrate boundary in the simulation model. Due to the larger nanowire geometry, region (ii) is less pronounced here than in Fig. 6.4(d).

6.3.6 Case study 6: Straight nanowire on diamond substrate

Recently, realizations of very high current densities have been reported when the permalloy nanowire is placed on a diamond substrate [99] instead of silicon nitride or silicon. This finding is useful for time-integrating experiments where current excitations exceeding the microsecond timescale are required [125, 126]. In this section, we investigate the geometry and experiment

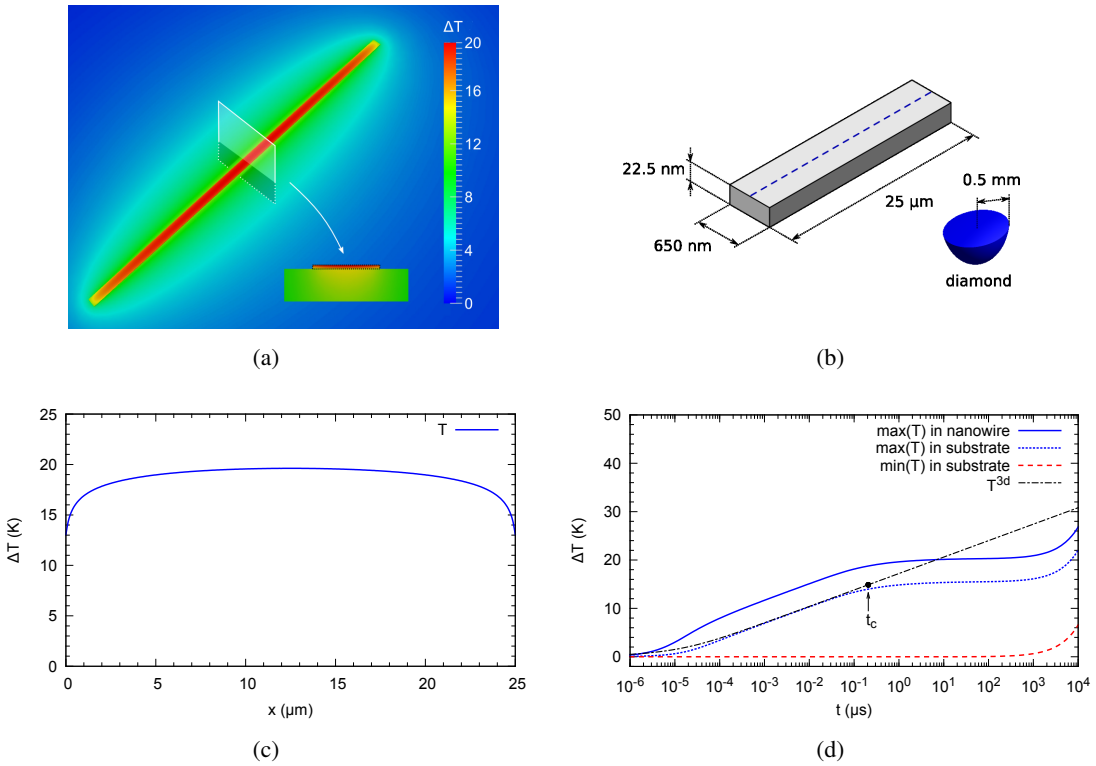

Figure 6.6: Joule heating in a permalloy nanowire on a diamond crystal substrate [99] (Sec. 6.3.6): (a) temperature distribution in the nanowire and substrate T after $t = 1\,\mu$s, (b) geometry of the model and the plotting path, (c) temperature profile $\Delta T(x)$ after $t = 1\,\mu$s along the plotting path, (d) maximum and minimum temperatures of the diamond substrate, maximum temperature of the wire, and the prediction of the You model as a function of pulse length. The current density is $j = 1.5 \times 10^{12}\,\mathrm{A/m^2}$; the resistivity of permalloy $\sigma^{-1} = 39\,\mu\Omega\mathrm{cm}$; for other material parameters see Tab. 6.1.

described in [99].

Figure 6.6(b) shows a rectangular permalloy nanowire with dimensions $25\,\mu\mathrm{m} \times 650\,\mathrm{nm} \times 22.5\,\mathrm{nm}$ which is grown on a diamond crystal substrate. The thermal conductivity of diamond depends on the purity of the crystal. For the simulation, we conservatively assume a relatively impure synthetic type Ib crystal with thermal conductivity $k = 1400\,\mathrm{W/(K\,m)}$ (Fig. 3 in [113]). A purer crystal would have higher thermal conductivity and would be more efficient at cooling the nanowire. The density and heat capacity of diamond are given in Tab. 6.1. As before, perfect thermal contact between the wire and the substrate is assumed. In the original experiment [99], the whole device was placed in a cryo bath. The area of contact between the diamond crystal substrate and the bath is large compared to the size of the nanowire, therefore the temperature difference at the contact layer is likely to be small. In the simulation, the effect of the bath can be represented by using an infinite medium of diamond. In line with the case studies 4 and 5 in the previous sections, we use a half-sphere shape for the diamond substrate (radius of 0.5 mm). The resistivity of the wire material is scaled to $39\,\mu\Omega\mathrm{cm}$ so that the total resistance of the wire is $675\,\Omega$ as in [99].

Figure 6.6(a) and 6.6(c) show the temperature profile after application of a current density of

$1.5 \times 10^{12}\,\mathrm{A/m^2}$ for 1 μs, and Fig. 6.6(d) shows the maximum and minimum temperatures in the system as a function of time.

From Fig. 6.6(d) we can see the three regimes as in case study 4 and 5: regime (i) shows logarithmic increase of temperature with time up to approximately $0.2\,\mu$s. For larger t the temperature remains constant during regime (ii). For t greater than approximately $1000\,\mu$s, we reach regime (iii) where the heatfront reaches the substrate boundary. The maximum temperature increase in the wire is not exceeding 21 K for times smaller than $1000\,\mu$s, and the substrate temperature increase stays below ≈ 16 K in that period. The plot shows that the precise current density pulse duration is not so critical: if t is between 1 μs and $1000\,\mu$s the nanowire attains approximately the same temperature. It is this steady state temperature increase value that should be compared with the experiment.

The difference between the maximum substrate temperature and the maximum wire temperature reflects a temperature gradient from top to bottom in the nanowire. This is just about visible in the inset in Fig. 6.6(a).

For larger times than $1000\,\mu$s, we find that the maximum and minimum temperature of the substrate and the maximum wire temperature start to increase simultaneously. This is an effect of the finite size of the substrate in the model, or, equivalently, the lack of the modelling of heat transfer away from the diamond substrate through the cryo bath. There are other, in comparison to the cryo bath less important cooling contributions, such as electric contacts, substrate holder, and surrounding gas, that have not been considered here. For the interpretation of the simulation results for the experiment in [99], we need to ignore the regime for $t \gtrsim 1000\,\mu$s.

The shape of the temperature profile (Fig. 6.6(c)) is very similar to the zigzag wire temperature profile (Fig. 6.5(c)).

In the original experiment, a continuous current corresponding to a current density of $1.5 \times 10^{12}\,\mathrm{A/m^2}$ heated the wire by about 230 K (Fig. 4 in [99]).

In our simulation the nanowire changed temperature by less than 21 K. A possible explanation for this discrepancy is that the contact and heat transfer between the wire and the diamond was imperfect.

6.4 Analytical models

6.4.1 Model by You, Sung, and Joe for a nanowire on a (3d) substrate

You, Sung, and Joe [94] have provided an analytic expression $T^{3\mathrm{d}}(t)$ to approximate the temperature $T(t)$ of the current-heated nanowire on a three-dimensional substrate as a function of time t (Eq. (16) in [94]):

$$T^{3\mathrm{d}}(t) = \frac{whj^2}{\pi k \sigma} \operatorname{arcsinh}\left(\frac{2\sqrt{tk/(\rho C)}}{\alpha w}\right) \tag{6.5}$$

where w and h are the width and height of the wire, σ is the wire conductivity, j is the current density, and k, ρ, and C are the thermal conductivity of the substrate, mass density of the substrate, and specific heat capacity of the substrate. We have used their adjustable parameter $\alpha = 0.5$ for calculations shown in Figs. 6.3 to 6.7.

We use name $T^{3\mathrm{d}}(t)$ for the equation from You, Sung and Joe – which is applicable to 3d-substrates – to emphasize the difference to the similar looking equation for $T^{2\mathrm{d}}(t)$ that is derived in Sec. 6.4.2 and which is applicable to 2d-substrates.

For derivation of Eq. (6.5), it is assumed that the nanowire is infinitely long, and attached to a semi-infinite substrate. While the thickness h and width w of the wire enter the derivation to compute the Joule heating due to a given current density, the model does not allow for a temperature variation within the nanowire nor does the nanowire have a heat capacity in the model. Within this model, a heat front of (half-)cylindrical shape (cylinder axis aligned with the wire) will propagate within the substrate when the wire is heated. Thus, there is translational invariance along the direction of the wire.

We start our discussion with the zigzag nanowire as shown in Fig. 6.5. Figure 6.5(d) shows the temperature prediction of the You model as a dash-dotted line. It follows the maximum temperature in the substrate very closely for times up to approximately $2\,\mu$s.

At short times t below 1 ns we can see the You model slightly overestimating the temperature in the nanowire in Fig. 6.5(d). As the model does not allow for a finite heat capacity of the wire, this is expected. As the heat capacity of the nanowire is insignificant in comparison to the substrate, this overestimation disappears if sufficient heat has been pumped into the system.

The difference between the maximum temperature in the wire and the maximum temperature in the substrate comes from a temperature gradient within the wire: the maximum temperature in the wire is at the top of the wire (which is furthest away from the cooling substrate) and the maximum temperature of the substrate is found at the top of the substrate just where the wire reaches its maximum temperature. Due to the assumption of perfect thermal contact, the bottom of the wire is exactly at the same temperature as the top of the substrate within the model description.

Since the You model does not allow for a temperature gradient within the wire, we expect its temperature prediction to follow the maximum temperature increase in the substrate. This is visible in Fig. 6.5(d) for $t \gtrsim 2\,$ns.

Regarding the deviation between the You model and the simulation results for $t \gtrsim 2\,\mu$s, we need to establish whether the required assumptions for the model are fulfilled. The You model is derived for an infinitely long wire on an infinite substrate, whereas the segments of the zigzag wire studied here have finite length. In the initial stage of heating, the temperature front in the substrate will move away from the wire sections with heat fronts aligned parallel with the wire. The heat front forms a half-cylinder (for each zigzag segment) whose axis is aligned with the wire. As long as the diameter of this half cylinder is small relative to the segment length, the wire appears locally to be infinitely long and the heat front propagates in a direction perpendicular to the wire. This is the regime where the You model is applicable, and which we have labeled as "regime (i)" in the discussion of case studies 4 (Sec. 6.3.4) to 6 (Sec. 6.3.6). When the heat front has propagated sufficiently far from the nanowire to change its shape from a cylinder surface to a spherical surface, the You model is not applicable anymore. This happens approximately after $t \gtrsim 2\,\mu$s. We have referred to the spherical heat front propagation in the discussion above as "regime (ii)".

For the zigzag wire study the agreement of the model by You *et al.* [94] with the simulation

is thus very good within the time range where the model is applicable. We note that the wire is relatively long (20 μm per segment) and has no constriction. The You model cannot be applied for $t \gtrsim 10\,\mu$s because the finite wire length becomes important.

For the nanowire without a constriction on the diamond substrate as studied in Sec. 6.3.6 and shown in Fig. 6.6, the agreement is similarly good; the You model temperature follows the substrate temperature very accurately up to $t \approx 0.1\,\mu$s. For larger t, the model becomes inaccurate as the finite length of the wire becomes important at that point.

Figure 6.4(d) shows for the nanowire with a notch on the silicon substrate that for $t \gtrsim 0.1\,\mu$s the gradients of both maximum temperature curves approach zero which indicates the onset of regime (ii) and implies that the You model cannot be applied for $t \gtrsim 0.1\,\mu$s. For smaller $t \lesssim 0.1\,\mu$s, the You model temperature roughly follows the maximum substrate temperature with a maximum absolute deviation of less than 3 K.

We have carried out additional simulations (data not shown) which have demonstrated that the maximum substrate temperature line (dotted line in Fig. 6.4(d)) is shifted down by a few degree Kelvin if the notch is removed from the geometry. The You model temperature and the maximum substrate temperature curves then coincide for $2 \lesssim t \lesssim 10$ ns. If, furthermore, we increase the wire length from 5 μm to 30 μm, the two curves coincide for $2 \lesssim t \lesssim 300$ ns.

Both the notch and the relative shortness of the wire decrease the accuracy of the prediction of the You model: the notch roughly shifts all temperature curves up by a few degrees whereas the length of the wire determines the time when regime (ii) is entered.

In contrast to the previous examples the nanowire with a notch in Fig. 6.3 is attached to a relatively thin silicon nitride membrane of thickness 100 nm (the silicon and diamond substrates for the discussion above are of the size order of 500 μm). The You model should not be used in the regime of membranes as the model expects an infinite substrate.

However, one could argue that the You model should be applicable for very small t until the heat front emerging from the wire has propagated through the 100 nm thick substrate. Additional simulations (data not shown) reveal that this is the case after $t \approx 0.1$ ns.

We summarize that the You model cannot be expected to provide accurate temperature predictions for thin substrates. The deviation of the You curve in Fig. 6.3(d) originates in the inappropriate application of the model to a system with a thin, effectively two-dimensional, substrate.

For a 'thick', effectively three-dimensional, substrate we find that the applicability of the model is limited by the finite length of the wire. The maximum time t_c up to which the You model is appropriate, can be estimated by calculating the characteristic time scale of the heat conduction equation (6.1). For a wire of length L, this characteristic time t_c is

$$t_c \sim \left(\frac{L}{2}\right)^2 \frac{\rho C}{k} \tag{6.6}$$

where k, ρ, and C are the thermal conductivity, density and specific heat capacity of the substrate material. The greater the nanowire length, the longer it takes for the heat front to assume spherical shape around the nanowire heating source. The greater the heat capacity (ρC) and the smaller the

thermal conductivity, the slower is the propagation of the heat front in the substrate. The current density does not enter the equation as it only affects the temperature and not the time or length scale.

Substituting the corresponding parameters for each case study, we compute the characteristic time t_c for the nanowire with a notch in case study 4 (Sec. 6.3.4) to be 70 ns, for the zigzag nanowire in case study 5 (Sec. 6.3.5) to be $8.8\,\mu$s, and for the nanowire on diamond in case study 6 (Sec. 6.3.6) to be $0.2\,\mu$s. We have added these values to the figures Fig. 6.4(d), 6.5(d), and 6.6(d) and they are in good agreement with the corresponding finite element results.

The time t_c marks the transition from regime (i) to regime (ii).

In summary, we find that the You model provides an accurate description of the maximum substrate temperature if used within its bounds of applicability, i.e. during regime (i) for three-dimensional substrates. In the You model, the nanowire has no heat capacity and this results in the model slightly overestimating the temperature for very small t (visible for example in Fig. 6.5(d) for $t \lesssim 1$ ns). The temperature within the nanowire can show a gradient (hotter at the top, cooler at the interface to the substrate), and the You model computes the smaller temperature in the wire. For the studies carried out here we find this temperature difference to be less than 10 K in all cases although this difference depends on material parameters and wire thickness. The You model cannot be applied for thin, effectively two-dimensional, substrates such as the membrane substrate case study in Sec. 6.3.3 and Fig. 6.3.

6.4.2 Analytic expression for nanowire on a membrane (2d substrate)

Equation (6.5) is valid for pulse durations up to the critical duration t_c (Eq. (6.6)) for effectively three-dimensional substrates, i.e. substrates whose thickness is sufficiently large so that the heat front in the substrate does not reach the substrate boundary for $t < t_\mathrm{c}$. This condition is fulfilled for the case studies in Sec. 6.3.4, 6.3.5, and 6.3.6. If the substrate is effectively two-dimensional (such as in Sec. 6.3.3), then Eq. (6.5) cannot be applied. In this section, we show how Eq. (6.5) can be adapted to the 2d case.

The assumptions made in the derivation of Eq. (6.5) require a system of nanowire and substrate that is translationally invariant in one direction. In Sec. 6.3.4–6.3.6 this direction was along the long axis of the wire, and we refer to this axis as x, and assume that the height h of the wire extends along the z axis.

For the case of a nanowire on a thin membrane substrate, we can regard the system as two-dimensional by assuming invariance in the perpendicular direction. In other words, to apply a modified version of Eq. (6.5), we assume that the temperature distribution in the membrane system is invariant along the z axis. We model the nanowire as embedded in the substrate (not grown on top of the substrate as in the real system) and imagine an increase of thickness of both the wire and the substrate such that they expand from $-\infty$ to $+\infty$ in z-direction. The cooling and heating in each slice (in the x-y plane) of the nanowire and substrate system is not affected by cooling and heating from slices above and below; it does not matter whether we consider only one isolated slice

(as in the real system) or imagine an infinite stack of slices closely packed on top of each other.

In more detail, we first convert the system of the nanowire of height h and substrate of thickness d to a 2d system of equal height. We increase the height of the nanowire by a factor $c = d/h$ so that the wire and the substrate are now both of height d (assuming that generally $d > h$ but the derivation also holds for $d < h$). This increases the volume of the wire by a factor of c, and thus we will have to correct down the heat emerging from the wire by the same factor at a later point.

Second, to obtain translational invariance in the z-direction we imagine a stack of such identical 2d systems on top of each other. Using the substitutions $w \to L$ and $h \to w$, we obtain

$$T^{\mathrm{2d}}(t) = \frac{Lwj^2}{\pi k\sigma}\frac{1}{c}\frac{1}{2}\,\mathrm{arcsinh}\left(\frac{2\sqrt{tk/(\rho C)}}{0.5L}\right) \tag{6.7}$$

$$= \frac{whLj^2}{2d\pi k\sigma}\,\mathrm{arcsinh}\left(\frac{2\sqrt{tk/(\rho C)}}{0.5L}\right) \tag{6.8}$$

The first fraction in Eq. (6.7) is based on Eq. (6.5) and includes the substitutions $w \to L$ and $h \to w$, and we also substitute $w \to L$ in the denominator of the $\mathrm{arcsinh}$ argument, and use $\alpha = 0.5$. The second fraction $(1/c)$ reduces the temperature increase by c to compensate for the increase of heating by the factor c above when we increased the thickness of the nanowire to the thickness of the substrate. The third fraction $(1/2)$ in Eq. (6.7) is a correction because the nanowire is now surrounded by substrate in all directions, and not only in one half-space as in Eq. (6.5), thus the cooling is twice as efficient and the temperature increase is halved.

Equation (6.8) can be used to compute the maximum temperature increase $T^{\mathrm{2d}}(t)$ for a nanowire of length L, width w, and height h on a two-dimensional substrate of thickness d.

In contrast to $T^{\mathrm{3d}}(t)$ there is no upper time limit t_{c} for the validity of Eq. (6.8) as the emerging heat-front will always stay translationally invariant.

The comparison of $T^{\mathrm{2d}}(t)$ with the finite element simulation results from case study 3 (Sec. 6.3.3) is shown in Fig. 6.7. The overall agreement with the simulation results is good for all times t. The heat capacity of the wire is not considered in the model for $T^{\mathrm{2d}}(t)$ which is reflected in the overestimation of the temperature in Fig. 6.7. From comparative simulations with different material parameters, we find that the agreement of the two curves is better for reduced heat capacity of the wire, and better for increased thermal conductivity of the substrate. The effect of the finite heat capacity of the wire becomes less important for longer current pulses, and the two curves in Fig. 6.6 become closer for larger t (not shown here).

6.5 Perpendicular nanowires

Recent progress in sample growth has allowed to create free-standing nanowires which are grown perpendicular to their substrate (for example [127, 128, 129]). While it is outside the scope of this work to investigate these systems in detail, we comment briefly on possible analytical approximations. For such a free standing nanowire, the analytical expression (6.4) is a good first approximation to compute its temperature increase as a function of applied current density duration.

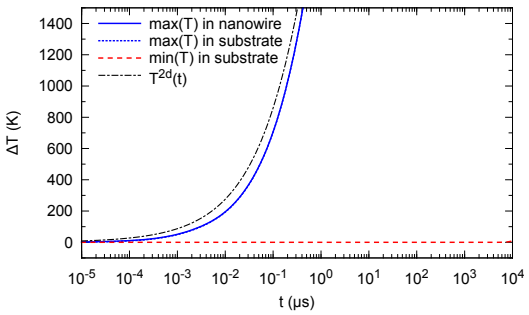

Figure 6.7: Comparison of simulated wire temperature from case study 3 (Sec. 6.3.3) with substrate temperatures obtained using the analytical expression $T^{2\mathrm{d}}(t)$ from Eq. (6.8), as a function of pulse length.

For the temperature increase $T_{\perp}^{3\mathrm{d}}(t)$ of a perpendicular nanowire which is completely embedded in a substrate material (such as an Al_2O_3 matrix), a variation of equation (6.5) can be employed:

$$T_{\perp}^{3\mathrm{d}}(t) = \frac{1}{2} T^{3\mathrm{d}}(t) \tag{6.9}$$

In contrast to the nanowire mounted on top of a half-space filling substrate (6.5) as studied in section 6.3, the substrate is here space-filling, and thus twice as effective in cooling the system.

6.6 Summary

We have carried out detailed numerical simulations of the current distribution, Joule heating, and dissipation of temperature and heat through the nanowire and the substrate for a number of experiments and three different substrate types. We find that the silicon nitride membrane (thickness 100 nm) is the least efficient in cooling a nanowire that experiences Joule heating. Due to the quasi-two-dimensional nature of the membrane, the temperature in the nanowire will keep increasing proportionally to the logarithm of time for longer current pulses while the heat front (forming a circle in the membrane substrate) propagates away from the nanowire, which is located in the center of the heat front circle.

Using a (effectively three-dimensional) silicon wafer substrate instead of the (effectively two-dimensional) silicon nitride membrane, there is a qualitative change: once the steady state is entered, the heat front propagates in three dimensions and keeps the temperature in the heated nanowire constant. In addition to the better cooling through the three-dimensionality of a silicon wafer, the cooling is improved further by the thermal conductivity of silicon which is more than one order of magnitude greater than that of silicon nitride (see Tab. 6.1).

If we replace silicon in the 3d substrate with diamond, the cooling is improved again significantly: diamond's thermal conductivity is about an order of magnitude greater than that of silicon (see Tab. 6.1).

In addition to these generic insights, we have worked out the temperature increase quantitatively for a series of recent experimental publications [86, 96, 99, 97, 98]. The model simulations

presented here show for all of them that the temperature increase due to the Joule heating did not result in the temperature exceeding the Curie temperature.

We compare these results to the approximating but analytical model expression provided by You, Sung, and Joe [94] and investigate the limits of its applicability. We provide an estimate for the characteristic time t_c over which the You model is valid for three-dimensional substrates.

Finally, we provide a new analytical expression that allows to compute the temperature for a nanowire on a two-dimensional substrate in the presence of Joule heating. This expression should be of significant value in the design and realization of spin-torque transfer studies on membranes where experimenters are often operating very close to the Curie temperature or even the melting temperature, and where no estimate of the wire's temperature has been possible so far.

Chapter 7

Summary and conclusions

The thesis describes several new computational methods for the micromagnetic model and their applications.

Chapter 3 presented a novel method to compute the demagnetising tensor in finite difference micromagnetics via numerical integration. This method aims to resolve the inaccuracy introduced by numerical cancellation when the analytical formula for the demagnetising tensor is used. Compared to existing approaches, the method achieves improved accuracy at intermediate distances from the origin.

Chapter 4 presented a novel semi-analytical method for the computation of the energy of the magnetostatic interaction between linearly magnetised polyhedrons. This method can be used for the direct computation of the demagnetising energy (and field, in the energy-based formulation) in the fast multipole method (as described in Section 2.3.4).

Chapter 5 presented a proposal for a standard problem for modelling ferromagnetic resonance in micromagnetic systems. The author implemented and applied the eigenvalue method (Section 2.5) to computing the resonance modes, frequencies, and spectrum for the reference system.

Chapter 6 studied the effect of Joule heating in ferromagnetic nanowires used in typical spin torque transfer experiments, via numerical and analytical modelling. The numerical results were compared to the approximate analytical estimate for 3d substrates, and a new analytical expression for the computation of the temperature increase in 2d substrates was derived.

Bibliography

[1] W. F. Brown, *Micromagnetics*. Interscience, 1963.

[2] D. Chernyshenko and H. Fangohr, "Computing the demagnetizing tensor for finite difference micromagnetic simulations via numerical integration," *Journal of Magnetism and Magnetic Materials*, vol. 381, pp. 440–445, 2015.

[3] A. Baker, M. Beg, G. Ashton, W. Wang, M. Albert, D. Chernyshenko, S. Zhang, M.-A. Bisotti, M. Franchin, C. L. Hu, R. Stamps, T. Hesjedal, and H. Fangohr, "Proposal of a micromagnetic standard problem for ferromagnetic resonance simulations." In preparation for submission, 2015.

[4] H. Fangohr, D. S. Chernyshenko, M. Franchin, T. Fischbacher, and G. Meier, "Joule heating in nanowires," *Phys. Rev. B*, vol. 84, p. 054437, Aug 2011.

[5] L. Landau and E. Lifshits, "On the theory of the dispersion of magnetic permeability in ferromagnetic bodies," *Phys. Z. Sowjetunion*, vol. 8, pp. 153–169, 1935.

[6] T. Gilbert, "A Phenomenological Theory of Damping in Ferromagnetic Materials," *IEEE Transactions on Magnetics*, vol. 40, pp. 3443–3449, Nov. 2004.

[7] A. Aharoni, *Introduction to the theory of Ferromagnetism*. Oxford University Press, 2001.

[8] J. Miltat and M. Donahue, "Numerical micromagnetics: Finite difference methods," in *Handbook of Magnetism and Advanced Magnetic Materials*, 2007.

[9] J. Fidler and T. Schrefl, "Micromagnetic modelling — the current state of the art," *Journal of Physics D: Applied Physics*, vol. 33, pp. R135–R156, Aug. 2000.

[10] M. Schabes and a. Aharoni, "Magnetostatic interaction fields for a three-dimensional array of ferromagnetic cubes," *IEEE Transactions on Magnetics*, vol. 23, pp. 3882–3888, Nov. 1987.

[11] A. J. Newell, W. Williams, and D. J. Dunlop, "A generalization of the demagnetizing tensor for nonuniform magnetization," *Journal of Geophysical Research*, vol. 98, no. B6, pp. 9551–9555, 1993.

[12] M. Maicus, E. Lopez, M. Sanchez, C. Aroca, and P. Sanchez, "Magnetostatic energy calculations in two- and three-dimensional arrays of ferromagnetic prisms," *IEEE Transactions on Magnetics*, vol. 34, no. 3, pp. 601–607, 1998.

[13] H. Fukushima, Y. Nakatani, and N. Hayashi, "Volume average demagnetizing tensor of rectangular prisms," *IEEE Transactions on Magnetics*, vol. 34, no. 1, pp. 193–198, 1998.

[14] M. J. Donahue, "Accurate computation of the demagnetization tensor," in *6th International Symposium on Hysteresis Modeling and Micromagnetics*, 2007. `http://math.nist.gov/~MDonahue/talks.html`, accessed Mar 2014.

[15] M. Donahue and D. Porter, "OOMMF User's Guide, Version 1.0." Interagency Report NISTIR 6376, National Institute of Standards and Technology, Gaithersburg, MD (Sept 1999).

[16] T. Schrefl, G. Hrkac, S. Bance, D. Suess, O. Ertl, and J. Fidler, *Numerical Methods in Micromagnetics (Finite Element Method)*. 2007.

[17] A. Knittel, *Micromagnetic simulations of three dimensional core-shell nanostructures*. PhD thesis, University of Southampton, UK, 2011.

[18] D. Fredkin and T. Koehler, "Hybrid method for computing demagnetizing fields," *IEEE Transactions on Magnetics*, vol. 26, pp. 415–417, Mar. 1990.

[19] A. van Herk, *Three-dimensional analysis of magnetic fields in recording head configurations*. PhD thesis, Tech. Hogeschool, Delft, 1980.

[20] D. Lindholm, "Three-dimensional magnetostatic fields from point-matched integral equations with linearly varying scalar sources," *IEEE Transactions on Magnetics*, vol. 20, pp. 2025–2032, Sept. 1984.

[21] C. Garcia-Cervera and A. Roma, "Adaptive mesh refinement for micromagnetics simulations," *IEEE Transactions on Magnetics*, vol. 42, pp. 1648–1654, June 2006.

[22] P. Visscher and D. Apalkov, "Simple recursive implementation of fast multipole method," *Journal of Magnetism and Magnetic Materials*, vol. 322, pp. 275–281, Jan. 2010.

[23] G. Brown, T. Schulthess, D. Apalkov, and P. Visscher, "Flexible fast multipole method for magnetic simulations," *IEEE Transactions on Magnetics*, vol. 40, pp. 2146–2148, July 2004.

[24] W. Zhang and S. Haas, "Adaptation and performance of the Cartesian coordinates fast multipole method for nanomagnetic simulations," *Journal of Magnetism and Magnetic Materials*, vol. 321, pp. 3687–3692, Nov. 2009.

[25] V. Rokhlin, "Rapid solution of integral equations of scattering theory in two dimensions," *Journal of Computational Physics*, vol. 86, pp. 414–439, Feb. 1990.

[26] R. Coifman, V. Rokhlin, and S. Wandzura, "The fast multipole method for the wave equation: a pedestrian prescription," *IEEE Antennas and Propagation Magazine*, vol. 35, pp. 7–12, June 1993.

[27] L. Greengard and V. Rokhlin, "A new version of the Fast Multipole Method for the Laplace equation in three dimensions," *Acta Numerica*, vol. 6, p. 229, Nov. 1997.

[28] W. Dehnen, "A Hierarchical O(N) Force Calculation Algorithm," *Journal of Computational Physics*, vol. 179, pp. 27–42, June 2002.

[29] R. Yokota, "An FMM Based on Dual Tree Traversal for Many-core Architectures," *Journal of Algorithms & Computational Technology*, vol. 7, pp. 301–324, Sept. 2013.

[30] M. D'Aquino, C. Serpico, and G. Miano, "Geometrical integration of Landau-Lifshitz-Gilbert equation based on the mid-point rule," *Journal of Computational Physics*, vol. 209, pp. 730–753, Nov. 2005.

[31] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward, "SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers," *ACM Transactions on Mathematical Software*, vol. 31, pp. 363–396, Sept. 2005.

[32] M. Farle, "Ferromagnetic resonance of ultrathin metallic layers," *Reports on Progress in Physics*, vol. 61, no. 7, p. 755, 1998.

[33] B. Heinrich, Y. Tserkovnyak, G. Woltersdorf, A. Brataas, R. Urban, and G. E. Bauer, "Dynamic exchange coupling in magnetic bilayers," *Physical Review Letters*, vol. 90, no. 18, p. 187601, 2003.

[34] B. Heinrich, Z. Celinski, J. Cochran, W. Muir, J. Rudd, Q. Zhong, A. Arrott, K. Myrtle, and J. Kirschner, "Ferromagnetic and antiferromagnetic exchange coupling in bcc epitaxial ultrathin Fe(001)/Cu(001)/Fe(001) trilayers," *Physical Review Letters*, vol. 64, no. 6, p. 673, 1990.

[35] H. T. Nembach, T. J. Silva, J. M. Shaw, M. L. Schneider, M. J. Carey, S. Maat, and J. R. Childress, "Perpendicular ferromagnetic resonance measurements of damping and Landé g-factor in sputtered $(Co_2Mn)_{1-x}Ge_x$ thin films," *Phys. Rev. B*, vol. 84, p. 054424, Aug 2011.

[36] R. McMichael and M. Stiles, "Magnetic normal modes of nanoelements," *Journal of Applied Physics*, vol. 97, May 2005.

[37] M. d'Aquino, C. Serpico, G. Miano, and C. Forestiere, "A novel formulation for the numerical computation of magnetization modes in complex micromagnetic systems," *Journal of Computational Physics*, vol. 228, no. 17, pp. 6130–6149, 2009.

[38] W. F. Brown, Jr., *Micromagnetics*. New York: Wiley Interscience, 1963.

[39] "MicroMagus." `http://www.micromagus.de/`, accessed Mar 2014.

[40] "LLG Micromagnetics Simulator." `http://llgmicro.home.mindspring.com/`, accessed Mar 2014.

[41] A. Vansteenkiste and B. Van de Wiele, "Mumax: a new high-performance micromagnetic simulation tool," *Journal of Magnetism and Magnetic Materials*, vol. 323, no. 21, pp. 2585–2591, 2011.

[42] "MicroMagnum." `http://micromagnum.informatik.uni-hamburg.de`, accessed Mar 2014.

[43] M. Labrune and J. Miltat, "Wall structures in ferro/antiferromagnetic exchange-coupled bilayers: A numerical micromagnetic approach," *Journal of Magnetism and Magnetic Materials*, vol. 151, pp. 231–245, Nov. 1995.

[44] G. Albuquerque, *Magnetization Precession in Confined Geometry: Physical and Numerical Aspects*. PhD thesis, Paris-Sud University, Orsay, 2002.

[45] J. Toussaint, A. Marty, N. Vukadinovic, J. Youssef, and M. Labrune, "A new technique for ferromagnetic resonance calculations," *Computational Materials Science*, vol. 24, pp. 175–180, May 2002.

[46] S. A. Smolyak, "Quadrature and interpolation formulas for tensor products of certain classes of functions," *Dokl. Akad. Nauk SSSR*, vol. 4, pp. 240–243, 1963.

[47] K. Petras, "Smolyak cubature of given polynomial degree with few nodes for increasing dimension," *Numerische Mathematik*, vol. 93, pp. 729–753, Feb. 2003.

[48] T. N. L. Patterson, "The optimum addition of points to quadrature formulae," *Mathematics of Computation*, vol. 22, pp. 847–856, 1968.

[49] L. Fousse, G. Hanrot, V. Lefèvre, P. Pélissier, and P. Zimmermann, "MPFR: A multiple-precision binary floating-point library with correct rounding," *ACM Transactions on Mathematical Software*, vol. 33, June 2007.

[50] "Online Encyclopaedia of Cubature Formulas." `http://nines.cs.kuleuven.be/ecf/`, accessed Mar 2014.

[51] T. Sørevik, *Full-symmetriske integrasjonsregler for enhets 4-kuben*. Master's thesis, Institutt for Informatikk, Universitetet i Bergen, 1985.

[52] R. Cools, "An encyclopaedia of cubature formulas," *Journal of Complexity*, vol. 19, pp. 445–453, June 2003.

[53] R. Cools, "Monomial cubature rules since Stroud: a compilation — part 2," *Journal of Computational and Applied Mathematics*, vol. 112, pp. 21–27, Nov. 1999.

[54] R. Cools and P. Rabinowitz, "Monomial cubature rules since Stroud: a compilation," *Journal of Computational and Applied Mathematics*, vol. 48, pp. 309–326, Nov. 1993.

[55] D. Wilton, S. Rao, A. Glisson, D. Schaubert, O. Al-Bundak, and C. Butler, "Potential integrals for uniform and linear source distributions on polygonal and polyhedral domains," *IEEE Transactions on Antennas and Propagation*, vol. 32, pp. 276–281, Mar. 1984.

[56] R. Graglia, "Static and dynamic potential integrals for linearly varying source distributions in two- and three-dimensional problems," *IEEE Transactions on Antennas and Propagation*, vol. 35, pp. 662–669, June 1987.

[57] R. Graglia, "On the numerical integration of the linear shape functions times the 3-D Green's function or its gradient on a plane triangle," *IEEE Transactions on Antennas and Propagation*, vol. 41, no. 10, pp. 1448–1455, 1993.

[58] M. Fabbri, "Magnetic Flux Density and Vector Potential of Uniform Polyhedral Sources," *IEEE Transactions on Magnetics*, vol. 44, pp. 32–36, Jan. 2008.

[59] M. Fabbri, "Magnetic flux density and vector potential of linear polyhedral sources," *COMPEL: The International Journal for Computation and Mathematics in Electrical and Electronic Engineering*, vol. 28, no. 6, pp. 1688–1700, 2009.

[60] A. Masserey, J. Rappaz, R. Rozsnyo, and M. Swierkosz, "Numerical integration of the three-dimensional Green kernel for an electromagnetic problem," *Journal of Computational Physics*, vol. 205, pp. 48–71, May 2005.

[61] M. Beleggia and M. De Graef, "General magnetostatic shape-shape interactions," *Journal of Magnetism and Magnetic Materials*, vol. 285, pp. L1–L10, Jan. 2005.

[62] M. Beleggia, S. Tandon, Y. Zhu, and M. De Graef, "On the magnetostatic interactions between nanoparticles of arbitrary shape," *Journal of Magnetism and Magnetic Materials*, vol. 278, pp. 270–284, July 2004.

[63] P. Arcioni, M. Bressan, and L. Perregrini, "On the evaluation of the double surface integrals arising in the application of the boundary integral method to 3-D problems," *IEEE Transactions on Microwave Theory and Techniques*, vol. 45, pp. 436–439, Mar. 1997.

[64] S. Wandzurat and H. Xiao, "Symmetric quadrature rules on a triangle," *Computers & Mathematics with Applications*, vol. 45, pp. 1829–1840, June 2003.

[65] H. Xiao and Z. Gimbutas, "A numerical algorithm for the construction of efficient quadrature rules in two and higher dimensions," *Computers & Mathematics with Applications*, vol. 59, pp. 663–676, Jan. 2010.

[66] D. V. Berkov, C. T. Boone, and I. N. Krivorotov, "Micromagnetic simulations of magnetization dynamics in a nanowire induced by a spin-polarized current injected via a point contact," *Phys. Rev. B*, vol. 83, p. 054420, Feb 2011.

[67] S. Erokhin, D. Berkov, N. Gorn, and A. Michels, "Micromagnetic modeling and small-angle neutron scattering characterization of magnetic nanocomposites," *Phys. Rev. B*, vol. 85, p. 024410, Jan 2012.

[68] G. Finocchio, I. N. Krivorotov, X. Cheng, L. Torres, and B. Azzerboni, "Micromagnetic understanding of stochastic resonance driven by spin-transfer-torque," *Phys. Rev. B*, vol. 83, p. 134402, Apr 2011.

[69] M. Najafi, B. Krüger, S. Bohlens, M. Franchin, H. Fangohr, A. Vanhaverbeke, R. Allenspach, M. Bolte, U. Merkt, D. Pfannkuche, D. P. F. Moller, and G. Meier, "Proposal for a standard problem for micromagnetic simulations including spin-transfer torque," *J. Appl. Phys.*, vol. 105, no. 11, p. 113914, 2009.

[70] F. Ma, H. Lim, Z. Wang, S. Piramanayagam, S. Ng, and M. Kuok, "Micromagnetic study of spin wave propagation in bicomponent magnonic crystal waveguides," *Appl. Phys. Lett.*, vol. 98, no. 15, pp. 153107–153107, 2011.

[71] W. F. Brown Jr., *Micromagnetics*. New York: Wiley, 1963.

[72] M. Donahue and D. Porter, *Interagency Report NISTIR 6376*. National Institute of Standards and Technology, Gaithersburg, MD.

[73] M. Scheinfein, *LLG Micromagnetics Simulator*.

[74] A. Vansteenkiste, J. Leliaert, M. Dvornik, M. Helsen, F. Garcia-Sanchez, and B. Van Waeyenberge, "The design and verification of MuMax3," *AIP Advances*, vol. 4, no. 10, pp. –, 2014.

[75] T. Fischbacher, M. Franchin, G. Bordignon, and H. Fangohr, "A systematic approach to multiphysics extensions of finite-element-based micromagnetic simulations: Nmag," *IEEE Trans. Magn.*, vol. 43, p. 2896, 2007.

[76] W. Scholz, J. Fidler, T. Schrefl, D. Suess, R. Dittrich, H. Forster, and V. Tsiantos, "Scalable parallel micromagnetic solvers for magnetic nanostructures," *Comp. Mater. Sci.*, vol. 28, p. 366, 2003.

[77] M. J. Donahue, D. G. Porter, R. D. McMichael, and J. Eicke, "Behavior of μMAG standard problem No. 2 in the small particle limit," *J. Appl. Phys.*, vol. 87, no. 9, pp. 5520–5522, 2000.

[78] R. Hertel and H. Kronmüller, "Finite element calculations on the single-domain limit of a ferromagnetic cube—a solution to μMAG Standard Problem No. 3," *J. Magn. Magn. Mater.*, vol. 238, no. 2, pp. 185–199, 2002.

[79] V. D. Tsiantos, D. Suess, T. Schrefl, and J. Fidler, "Stiffness analysis for the micromagnetic standard problem No. 4," *J. Appl. Phys.*, vol. 89, no. 11, pp. 7600–7602, 2001.

[80] G. Venkat, D. Kumar, M. Franchin, O. Dmytriiev, M. Mruczkiewicz, H. Fangohr, A. Barman, M. Krawczyk, and A. Prabhakar, "Proposal for a standard micromagnetic problem: Spin wave dispersion in a magnonic waveguide," *IEEE Trans. Magn.*, vol. 49, no. 1, pp. 524–529, 2013.

[81] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, 2007.

[82] F. Guo, L. M. Belova, and R. McMichael, "Spectroscopy and imaging of edge modes in permalloy nanodisks," *Physical Review Letters*, vol. 110, no. 1, p. 017601, 2013.

[83] C. Kittel, *Introduction to solid state physics*. Wiley, 2005.

[84] T. Kamionka, M. Martens, K. W. Chou, M. Curcic, A. Drews, G. Schütz, T. Tyliszczak, H. Stoll, B. Van Waeyenberge, and G. Meier, "Magnetic antivortex-core reversal by circular-rotational spin currents," *Phys. Rev. Lett.*, vol. 105, no. 13, p. 137204, 2010.

[85] A. Yamaguchi, T. Ono, S. Nasu, K. Miyake, K. Mibu, and T. Shinjo, "Real-space observation of current-driven domain wall motion in submicron magnetic wires," *Phys. Rev. Lett.*, vol. 92, no. 7, p. 077205, 2004.

[86] M. Kläui, P.-O. Jubert, R. Allenspach, A. Bischof, J. A. C. Bland, G. Faini, U. Rüdiger, C. A. F. Vaz, L. Vila, and C. Vouille, "Direct observation of domain-wall configurations transformed by spin currents," *Phys. Rev. Lett.*, vol. 95, no. 2, p. 026601, 2005.

[87] G. Meier, M. Bolte, R. Eiselt, B. Krüger, D.-H. Kim, and P. Fischer, "Direct imaging of stochastic domain-wall motion driven by nanosecond current pulses," *Phys. Rev. Lett.*, vol. 98, no. 18, p. 187202, 2007.

[88] S. Lepadatu, M. C. Hickey, A. Potenza, H. Marchetto, T. R. Charlton, S. Langridge, S. S. Dhesi, and C. H. Marrows, "Experimental determination of spin-transfer torque nonadiabaticity parameter and spin polarization in permalloy," *Phys. Rev. B*, vol. 79, no. 9, p. 094402, 2009.

[89] S. S. P. Parkin, M. Hayashi, and L. Thomas, "Magnetic domain-wall racetrack memory.," *Science (New York, N.Y.)*, vol. 320, pp. 190–4, Apr. 2008.

[90] S. S. P. Parkin tech. rep., 2004. U.S. Patent No. 309 6 834 005.

[91] S. Urazhdin, N. O. Birge, W. P. Pratt, and J. Bass, "Current-driven magnetic excitations in permalloy-based multilayer nanopillars," *Phys. Rev. Lett.*, vol. 91, no. 14, p. 146803, 2003.

[92] M. Laufenberg, W. Bührer, D. Bedau, P.-E. Melchy, M. Kläui, L. Vila, G. Faini, C. A. F. Vaz, J. A. C. Bland, and U. Rüdiger, "Temperature dependence of the spin torque effect in current-induced domain wall motion," *Phys. Rev. Lett.*, vol. 97, no. 4, p. 046602, 2006.

[93] F. Junginger, M. Kläui, D. Backes, U. Rüdiger, T. Kasama, R. E. Dunin-Borkowski, L. J. Heyderman, C. A. F. Vaz, and J. A. C. Bland, "Spin torque and heating effects in current-induced domain wall motion probed by transmission electron microscopy," *Appl. Phys. Lett.*, vol. 90, no. 13, p. 132506, 2007.

[94] C. You, I. Sung, and B. Joe, "Analytic expression for the temperature of the current-heated nanowire for the current-induced domain wall motion," *Appl. Phys. Lett.*, vol. 89, p. 222513, 2006.

[95] P. Bruno, "Geometrically constrained magnetic wall," *Phys. Rev. Lett.*, vol. 83, no. 12, pp. 2425–2428, 1999.

[96] M.-Y. Im, L. Bocklage, P. Fischer, and G. Meier, "Direct observation of stochastic domain-wall depinning in magnetic nanowires," *Phys. Rev. Lett.*, vol. 102, no. 14, p. 147204, 2009.

[97] L. Bocklage, B. Krüger, T. Matsuyama, M. Bolte, U. Merkt, D. Pfannkuche, and G. Meier, "Dependence of magnetic domain-wall motion on a fast changing current," *Phys. Rev. Lett.*, vol. 103, no. 19, p. 197204, 2009.

[98] H. H. Langner, L. Bocklage, B. Krüger, T. Matsuyama, and G. Meier, "Magnetic domain-wall depinning with reduced current density by short pulse rise time," *Applied Physics Letters*, vol. 97, no. 24, p. 242503, 2010.

[99] S. Hankemeier, K. Sachse, Y. Stark, R. Frömter, and H. Oepen, "Ultrahigh current densities in Permalloy nanowires on diamond," *Appl. Phys. Lett.*, vol. 92, p. 242503, 2008.

[100] Ansys Inc, "ANSYS, version 12.0," 2010. http://www.ansys.com.

[101] Multiphysics Modeling and Engineering Simulation Software, "COMSOL Multi-physics version 4.0," 2008. http://www.comsol.com/.

[102] T. Fischbacher and H. Fangohr, "Continuum multi-physics modeling with scripting languages: the Nsim simulation compiler prototype for classical field theory," *ArXiv e-prints*, p. 0907.1587, 2009. http://arxiv.org/abs/0907.1587.

[103] T. Fischbacher, M. Franchin, G. Bordignon, and H. Fangohr, "A systematic approach to multiphysics extensions of finite-element based micromagnetic simulations: Nmag," *IEEE Transactions on Magnetics*, vol. 43, no. 6, p. 2896, 2007. http://nmag.soton.ac.uk.

[104] G. Meier. personal communication, 2010.

[105] D. Bonnenberg, K. A. Hempel, and H. P. J. Wijn, *SpringerMaterials - The Landolt-Börnstein Database*, ch. 1.2.1.2.10 Thermomagnetic properties, thermal expansion coefficient, specific heat, Debye temperature, thermal conductivity, p. 252. Springer.

[106] C. Ho, M. Ackerman, K. Wu, S. Oh, and T. Havill, "Thermal conductivity of ten selected binary alloy systems," *Journal of Physical and Chemical Reference Data*, vol. 7, no. 3, pp. 959–1178, 1978.

[107] E. A. Owen, E. L. Yates, and A. H. Sully, "An x-ray investigation of pure iron-nickel alloys. part 4: the variation of lattice-parameter with composition," *Proceedings of the Physical Society*, vol. 49, no. 3, p. 315, 1937.

[108] C. H. Mastrangelo, Y.-C. Tai, and R. S. Muller, "Thermophysical properties of low-residual stress, Silicon-rich, LPCVD silicon nitride films," *Sensors and Actuators A: Physical*, vol. 23, pp. 856–860, Apr. 1990.

[109] E. Domalski and E. Hearing, *NIST Chemistry WebBook, NIST Standard Reference Database Number 69*, ch. Condensed Phase Heat Capacity Data. NIST.

[110] C. Y. Ho, R. W. Powell, and P. Liley, *Thermal Conductivity of the Elements: A Comprehensive Review*. AIP, 1978.

[111] H. Bettin, M. Glaser, F. Spieweck, H. Toth, A. Sacconi, A. Peuto, K. Fujii, M. Tanaka, and Y. Nezu, "International intercomparison of silicon density standards," *IEEE Transactions on Instrumentation and Measurement*, vol. 46, no. 2, pp. 556 –559, 1997.

[112] E. Domalski and E. Hearing, *NIST Chemistry WebBook, NIST Standard Reference Database Number 69*, ch. Condensed Phase Heat Capacity Data. NIST.

[113] Y. Yamamoto, T. Imai, K. Tanabe, T. Tsuno, Y. Kumazawa, and N. Fujimori, "The measurement of thermal properties of diamond," *Diamond and Related Materials*, vol. 6, no. 8, pp. 1057 – 1061, 1997.

[114] P. Ownby and R. Stewart, *Engineered Materials Handbook*, vol. 4, ch. Engineering Properties of Diamond and Graphite, pp. 821–834. ASM International, 1991.

[115] H. Fangohr, J. P. Zimmermann, R. P. Boardman, D. C. Gonzalez, and C. H. de Groot, "Numerical investigation of domain walls in constrained geometries," *J. Appl. Phys.*, vol. 103, no. 7, pp. 07–926, 2008.

[116] Y. Wang, C. H. de Groot, D. Claudio-Gonzalez, and H. Fangohr, "Magnetoresistance in a lithography defined single constrained domain wall spin-valve," *Applied Physics Letters*, vol. 97, no. 26, p. 262501, 2010.

[117] L. K. Bogart and D. Atkinson, "Domain wall anisotropic magnetoresistance in planar nanowires," *Appl. Phys. Lett.*, vol. 94, no. 4, p. 042511, 2009.

[118] Silson Ltd. Northampton, United Kingdom. http://www.silson.com.

[119] X. Zhang and C. P. Grigoropoulos, "Thermal conductivity and diffusivity of free-standing silicon nitride thin films," *Review of Scientific Instruments*, vol. 66, no. 2, p. 1115, 1995.

[120] A. Irace and P. M. Sarro, "Measurement of thermal conductivity and diffusivity of single and multilayer membranes," *Sensors and Actuators A: Physical*, vol. 76, pp. 323–328, Aug. 1999.

[121] B. Zink, "Specific heat and thermal conductivity of low-stress amorphous Si-N membranes," *Solid State Communications*, vol. 129, pp. 199–204, Jan. 2004.

[122] R. Sultan, a. D. Avery, G. Stiehl, and B. L. Zink, "Thermal conductivity of micromachined low-stress silicon-nitride beams from 77 to 325 K," *Journal of Applied Physics*, vol. 105, no. 4, p. 043501, 2009.

[123] L. Thomas, M. Hayashi, X. Jiang, C. Rettner, and S. S. P. Parkin, "Perturbation of spin-valve nanowire reference layers during domain wall motion induced by nanosecond-long current pulses," *Appl. Phys. Lett.*, vol. 92, no. 11, p. 112504, 2008.

[124] J. Crangle and G. C. Hallam, "The magnetization of face-centred cubic and body-centred cubic iron + nickel alloys," *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 272, no. 1348, pp. 119–132, 1963.

[125] L. Heyne, J. Rhensius, D. Ilgaz, A. Bisig, U. Rüdiger, M. Kläui, L. Joly, F. Nolting, L. J. Heyderman, J. U. Thiele, and F. Kronast, "Direct determination of large spin-torque nonadiabaticity in vortex core dynamics," *Phys. Rev. Lett.*, vol. 105, no. 18, p. 187203, 2010.

[126] B. Krüger, M. Najafi, S. Bohlens, R. Frömter, D. P. F. Möller, and D. Pfannkuche, "Proposal of a robust measurement scheme for the nonadiabatic spin torque using the displacement of magnetic vortices," *Phys. Rev. Lett.*, vol. 104, no. 7, p. 077201, 2010.

[127] K. Nielsch, F. Müller, A. Li, and U. Gösele, "Uniform nickel deposition into ordered alumina pores by pulsed electrodeposition," *Advanced Materials*, vol. 12, no. 8, pp. 582–586, 2000.

[128] L. Vila, P. Vincent, L. Dauginet-De Pra, G. Pirio, E. Minoux, L. Gangloff, S. Demoustier-Champagne, N. Sarazin, E. Ferain, R. Legras, L. Piraux, and P. Legagneux, "Growth and field-emission properties of vertically aligned cobalt nanowire arrays," *Nano Letters*, vol. 4, no. 3, pp. 521–524, 2004.

[129] K. Pitzschel, J. Bachmann, S. Martens, J. M. Montero-Moreno, J. Kimling, G. Meier, J. Escrig, K. Nielsch, and D. Görlitz, "Magnetic reversal of cylindrical nickel nanowires with modulated diameters," *Journal of Applied Physics*, vol. 109, no. 3, p. 033907, 2011.