

UNIVERSITY OF SOUTHAMPTON

Computational Methods and Studies in Nanomagnetics

by

Ryan Alexander Pepper

A thesis submitted for the degree of
Doctor of Philosophy

in the
Faculty of Engineering and Physical Sciences
Aeronautics, Astronautics and Computational Engineering

July 2020

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND PHYSICAL SCIENCES
AERONAUTICS, ASTRONAUTICS AND COMPUTATIONAL
ENGINEERING

Doctor of Philosophy

**COMPUTATIONAL METHODS AND STUDIES IN
NANOMAGNETICS**

by Ryan Alexander Pepper

Computational study of magnetic materials has been crucial for the development of new technologies in areas such as data storage. One challenge with current computational methods is that the dipolar field calculation dominates the computation time. In this work we show how the fast multipole method can be applied to this problem, and other long range force and potential calculations, through symbolic generation of operator functions in a generic fashion. We study the equilibrium states found in triangular and square samples of the helimagnetic material FeGe, in which skyrmions have been observed, by varying the size and applied field. We show that the equilibrium states of such systems is modified in comparison to previously studied disk systems of this material, with larger sample sizes required for skyrmions to form the ground state. We show the final states obtained from relaxation of a uniform magnetisation in order to provide data for experimental comparison. We then study the energy barriers between ferromagnetic and skyrmion states in Cobalt monolayers when triangular, square and Bezier edge defects are introduced, and show how this varies by size and by whether the dipolar field is included in the calculation. Finally, we study the equilibrium behaviour of Bloch points in FeGe disks and nanotacks made up of two layers in which the Dzyaloshinskii-Moriya interaction has opposing chirality. We then study the dynamic behaviour under an in-plane magnetic field, showing that the Bloch point reaches a velocity linearly proportional to the applied field.

Declaration of Authorship

I, Ryan Alexander Pepper, declare that this thesis entitled *Computational Methods and Studies in Nanomagnetics* 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:
 - Chapter 4 in Ryan Alexander Pepper, Hans Fangohr, fmmgen: Automatic Code Generation of Operators for Cartesian Fast Multipole and Barnes-Hut Methods, arXiv:2005.12351 (2020)
 - Chapter 5 in Ryan Alexander Pepper, et. al., Skyrmion states in thin confined polygonal nanostructures, Journal of Applied Physics 123, 093903 (2018).

Signed:

Date:

Contents

Declaration of Authorship	v
Nomenclature	ix
Acknowledgements	xii
1 Introduction	1
2 Magnetic Theory and Background	7
2.1 Introduction	7
2.2 Origins of Magnetism	7
2.3 Heisenberg Model and Micromagnetics	10
2.3.1 Heisenberg Exchange	12
2.3.2 Dzyaloshinskii-Moriya Exchange	14
2.3.3 Anisotropy	16
2.3.4 Dipolar Interaction	16
2.3.5 Competition between energies, characteristic length scales and emergent phenomena	17
2.3.6 Equation of Motion	19
3 Computational Methods in Nanomagnetics	23
3.1 Introduction	23
3.2 Atomistic Models	24
3.3 Micromagnetics	29
3.3.1 Finite Difference Method	29
3.3.2 Finite Element Method	33
3.3.2.1 Generalised Weighted Average Method and the Weak Form	33
3.3.2.2 Application to Micromagnetics	37
3.4 Evolution of the magnetisation	39
3.4.1 Time Integration	39
3.4.2 Energetics	41
3.5 Approach Choice: Micromagnetics or Atomistic?	44
4 Hierarchical Methods for Atomistic Long Range Field Calcula- tions	47

4.1	Mathematical Basis	50
4.2	Implementation	52
4.2.1	Operator Generation	52
4.3	Methods	56
4.3.1	Tree Decomposition	56
4.3.2	Tree Traversal	58
4.4	Testing	59
4.5	Performance comparison with Fast Fourier Transform Approach for Dipolar Field	63
4.6	Suggestions for Performance Improvements of Hierarchical Calculations	65
4.7	Discussion	68
5	Magnetic Skyrmions	71
5.1	Introduction	71
5.2	Skyrmions in Thin Polygonal Nanostructures	77
5.2.1	Method	77
5.2.2	Results	79
5.2.2.1	Equilibrium States	79
5.2.2.2	Ground States	84
5.2.2.3	Proposed experimental study	85
5.2.3	Conclusion	86
6	Edge Defects and Energy Barriers in Skyrmion Monolayers	89
6.1	Introduction	89
6.2	System Definition	92
6.3	Results	94
6.3.1	Triangular Defects	94
6.3.2	Rectangular and Bezier Defects	101
6.4	Conclusion	103
7	Bloch Points in Chiral Systems	105
7.1	Background	105
7.2	Field Induced Bloch Point Motion	109
7.2.1	Bloch Point in Equilibrium under an Applied Field	109
7.2.2	Bloch Point Velocity	112
7.3	Conclusion	118
8	Conclusion	121
	Bibliography	123

Nomenclature

r	Position
v	Velocity
μ	Magnetic Moment
J	Heisenberg Exchange Constant
d	Dzyaloshinskii-Moriya Exchange Constant
k	Anisotropy Constant
$\hat{\mathbf{u}}$	Anisotropy Axis
\mathbf{M}	Magnetisation
\mathbf{m}	Normalised Magnetisation
M_s	Saturation Magnetisation
A	Micromagnetic Exchange Constant
D	Micromagnetic Dzyaloshinskii-Moriya Constant
H	External Magnetic Field
\mathbf{H}_{eff}	Effective Field
α	Gilbert Damping Coefficient
λ	Landau-Lifshitz Damping Coefficient
γ	Gyromagnetic Ratio
W	Winding Number
S	Skyrmion Number
BP	Bloch Point
FMM	Fast Multipole Method
BH	Barnes-Hut Method
FFT	Fast Fourier Transform

To Mum

Acknowledgements

Firstly, I'd like to thank my supervisor Prof. Hans Fangohr, especially for teaching me how to be a 'real' software engineer, for the brilliant chances I've had to learn and develop as a scientist with him, and for the opportunities he has given me in terms of stepping outside of my comfort zone at conferences and meetings. I would also like to thank my cosupervisor Dr. Ondrej Hovorka for being a great source of entertainment and moral support. I would especially like to thank my close colleagues Dr. Marijan Beg and Dr. David Cortés-Ortuño for their help in understanding the subject matter and for suggesting interesting avenues for enquiry during my PhD. I would also like to thank my other colleagues in the Fangohr group who have since departed from the University of Southampton; Dr. Weiwei Wang, Ms. Rebecca Carey, Dr. Maximilian Albert, Dr. Jonathan Waters, Dr. Mark Vousden, Mr. Marc-Antonio Bisotti, Dr. Oliver Laslett, Dr. Thomas Kluyver, and Dr. Ondrej Hovorka. I'd also like to thank the friends and colleagues in the Centre for Doctoral Training in Next Generation Modelling, especially Dr. Craig Rafter, Dr. David Lusher and Mr. Hossam Ragheb.

The Iridis 4 and 5 High Performance Computing Facilities and their support facilities were used extensively in this work, and I would particularly like to thank David Baker, who gave me the opportunity to test the Iridis 5 cluster soon after its delivery, along with its fledgling visualisation service, which had an enormous amount of impact on the time taken to run and understand my simulation results.

I would like to thank my examiners, Prof. Wyn Williams from the University of Edinburgh and Dr. Marcus Newton from the University of Southampton for taking the time to read my thesis and for my viva voce examination. Their useful suggestions and critique have helped me to improve both the general readability of this thesis and the scientific content.

I would also like to thank my Dad and sister, Jessica, for their support throughout, and finally, to my girlfriend, fiancée and now wife, Coral, for understanding the stresses of PhD life, and putting up with me dashing off at odd hours of the night to the computer to check on simulation results.

I acknowledge financial support from EPSRC Doctoral Training Centre Grant EP/L015382/1, OpenDreamKit Horizon 2020 European Research Infrastructures project (676541), and the EPSRC Programme grant on Skyrmeionics (EP/N032128/1). I would particularly like to thank the Computational Physics Group of the Institute of Physics, who granted me a travel bursary which helped me attend the

2017 SolSkyMag conference in San Sebastian, Spain, and the additional conference budget granted by the NGCM Doctoral Training Center which enabled to me to attend the MMM 2019 conference in Washington DC.

Chapter 1

Introduction

The properties of magnetic materials are of great interest due to the wide range of applications in areas as diverse as power generation [1], electric motors [2], and in consumer electronic devices such as speakers and microphones. Probably the most ubiquitous use of magnets in everyday consumer technology currently is in magnetic storage devices. The magnetic hard drive, invented at IBM in the early 1950s [3], revolutionised data storage in the computer industry. The density of magnetic storage has increased dramatically since the conception of the hard drive, driven both by better understanding of the magnetic behaviour of the materials, and rapid developments in reading technology, a particularly notable example of which was the discovery of giant magneto-resistance (GMR) by Fert and Grunberg in the 1980s, for which the Nobel Prize in Physics was awarded [4, 5]. Today, magnetic recording densities stand at over 1Tb/in².

Engineering hard drives to increase the storage density further has proved difficult, however, due to the so-called ‘magnetic recording trilemma’ [6]. In present hard-drive technology, which is based around perpendicular magnetic recording, a single bit is written onto a region containing many magnetic grains [7]. As magnetic grain sizes have become smaller, the amount of energy needed to change the magnetisation orientation (and hence the bit value) is reduced. For a given material, then, an increase in bit-density leads to a shorter lifetime for data integrity. Increasing the anisotropy density correspondingly allows for sufficiently long data lifetimes, but brings additional problems. With an increased anisotropy, the magnetic field needed to write a value to a bit also becomes correspondingly larger, and increasing the write field at small scales is difficult [8]. A combination of the scaling restrictions, fragility and comparatively slow read/write performance of magnetic storage has led to NAND-gate based solid state drives (SSDs)

gaining increasing market share and usage, despite limitations on the number of write-cycles, lack of suitability for archival purposes, and the relative cost of such drives being higher [9].

Research into increasing the storage density of magnetic storage has been ongoing, and many proposals have been suggested. These can be broadly split into two categories; those which are based on existing technology with small modifications, and those which are based on very different principles for storing data. Heat Assisted Magnetic Recording (HAMR) and Microwave Assisted Magnetic Recording (MAMR) have been the most successful implementations so far, and rely on reducing the field needed to write to a grain by laser-induced heating and microwave respectively [10–12]. Materials constraints for these types of devices have been difficult to fulfil [13], and at the time of writing these technologies are only just becoming commercially available [14]. Proposals for bit-patterned media (BPM) take a different approach to increasing the storage density, by patterning individual magnetic bits onto a substrate with the spatial isolation of the bits leads to increased stability. However, the actual experimental procedures for producing such materials either through advanced lithography techniques or self-assembly are difficult, and it remains to be seen as to whether such devices are commercially viable [15, 16]. Combinations of these approaches have also been proposed as possibilities for future density advances [17, 18].

Magnetic storage based around the internal motion of a magnetic material is not a new concept, but research into such technologies underwent something of a revival in recent years. Bubble domain memory, in which the bubbles formed data carriers in tracks, was heavily researched and saw some utilisation as a solid-state alternative to hard-drives from the mid-1970s until the mid-1980s [19, 20]. It became superseded in use by transistor based memory such as flash, aside from in some specialised areas [21], in part because the read time of such devices was much slower. More recently, proposals of current driven domain-wall based storage on racetrack systems were intensively studied [22]. Proposals for such devices discussed the possibility of racetracks which could either be planar, or in a more complex geometry, with proposed high density storage being based around 3-D layouts of the track. Current magnetic storage devices are mechanical, and data is brought to the position of a reading head through the rotation of a spinning platter which holds the information. The concept of racetrack is in contrast solid state - the information itself is moved along the track towards reading heads using short pulses of current. While there has been a large amount of research into domain wall racetrack storage [23, 24], in practice there have been found to be

many engineering challenges, such as pinning of the domain walls and the break down of domain walls above the Walker threshold [25] and no commercial devices have ever been produced.

Recent proposed studies have suggested that as an alternative to domains forming data carriers in racetrack memories, a skyrmion could be used instead [26]. Skyrmions, which are particle like configurations of a field, have been shown to be stable configurations in certain types of ferromagnetic compounds, which lack an inversion symmetry [27–29]. These states have been shown to be the ground state configuration in some confined geometries [30–32]. Skyrmions are afforded some degree of ‘topological protection’, though to what extent this is true is a hotly debated topic, due to the discrete atomic origins of magnetism. Skyrmions were first observed in the B20 material class, but much of the current experimental work in this area has focused on multi-layered structures, where one or more magnetic layers are in contact with a non-magnetic heavy metal such as platinum [33–36].

Much of the work thus far on skyrmion based racetrack storage has been purely theoretical [37–41]. While the physics of skyrmions is in general fairly well understood from a theoretical standpoint, controlling them reliably seems to be proving a large issue in experimental research. This is because the skyrmion driven by a current undergoes a Topological hall effect, resulting in Coriolis-like motion which acts to move the Skyrmion texture perpendicular to the current direction, due to electrons picking up a Berry phase when passing through the non-uniform magnetisation of the Skyrmion [42, 43]. Recent research has in addition suggested that pinning will be an important issue [44]. Thus far, techniques for reliable electrical detection isolated individual skyrmions in racetrack type samples are limited in temperature range [45, 46]. One of the most fundamental questions that still remains to be answered is what effect the shape and boundaries of physical systems have on skyrmions. Whether real skyrmion memory devices are feasible is as yet an unanswered question, but as yet there is little to suggest that the skyrmion as an information carrier would offer much of an advantage over bubble or domain-wall technology. However, the skyrmion remains an interesting playground for fundamental research, as an example of an emergent property, with many open questions [47].

With the interest in chiral materials driven by the need for improved storage technologies, the study of other magnetic textures which can emerge in such systems has undergone a revival. Bloch points, which were widely studied in the 1970s and 1980s in domain walls, have been observed as mediators to skyrmion destruction.

Subsequent efforts based around domain walls were then proposed; in the late 1980s vertical Bloch-line technologies, in which a Bloch line propagated along a domain wall were proposed but never reached commercial viability [48, 49]. Recent interest in Bloch points has been driven by their presence in reversal mechanisms and other dynamic processes [50–53]. The study of Bloch points therefore provides another interesting, and less well studied, avenue of research.

The need to understand magnetic materials has led to the development of theoretical models which can explain their static and dynamic behaviour. Today, the computational study of magnetic systems uses a variety of techniques, which depends on the length scale of study. At the most fundamental level, density functional theory (DFT) is used to determine material parameters of systems, such as the electron density of atoms and the resulting magnetic parameters, though requires enormous amounts of computational power [54]. Classical models of magnets, whereby each atom is treated as an isolated magnetic dipole, are the next level on the hierarchy of modelling materials [55–58], though even here many problems of interest for devices require such computational power as to be intractable; the largest studies up to this point are still only able to treat relatively small system sizes [59]. Micromagnetism is a continuum field description of these classical models [60–62], and while this description is not adequate for some problems (such as the study of antiferromagnetism), it is nonetheless of great utility. In both discrete models and micromagnetism, the dominant part of any calculation is that of the non-local magnetostatic field, which limits the scaling of any calculation. In a narrow subset of problems, where the materials under study are ultra-thin films, the magnetostatic field can be treated as a weak anisotropy [63], but this is not applicable in general, and magnetostatic effects greatly change the energetics of phenomena such as domain walls and skyrmions [32, 64]. As a result, it is of critical importance to treat the magnetostatic calculations effectively, and a number of numerical methods for doing so have been used in the past. In particular, the most commonly used method requires the use of the Fast Fourier Transform (FFT) [65], but this restricts numerical treatment of micromagnetics and classical spin models to regular grids. In recent years, the Fast Multipole Method (FMM), with which the accuracy of calculation can be strictly controlled, has been proposed as a method for accelerating calculations in both atomistic and micromagnetic simulations [66, 67]. Despite this, it has not been widely applied to the study of nanomagnetic systems, and there are not suitable public implementations of the method which are straightforward to integrate into spin dynamics codes.

The layout of this thesis is as follows. In Chapter 2, a review of the theory behind

micromagnetism and classical spin models is shown. We discuss the magnetisation dynamics in ferromagnetic materials, and the energy terms most frequently used in micromagnetic and spin dynamics calculations. In Chapter 3, we review the common numerical methods used to solve micromagnetic and spin dynamics problems computationally. Chapter 4 discusses techniques for accelerating dipole field calculations and sets out accuracy and performance characteristics of these, and describes an open-source project fmmgen [68] which resulted from this work. In Chapter 5, we present a review of the literature on the phenomena of magnetic skyrmions, including discussion of their origin in the competition between magnetic energy terms. We discuss different types of system in which skyrmions have been observed, and review experimental observations of the behaviour of skyrmions. We show the results of work done on the energy landscape of skyrmions in confined polygonal geometries. Here, phase diagrams are produced showing the ground states for particular sizes of geometry of the material FeGe, and for varying applied fields. In Chapter 6, we show the result of energy barrier calculations in atomistic systems, by introducing different types of edge defect into a monolayer. We calculate quantitatively how the energy barrier to skyrmion destruction and creation is affected by such defects. In Chapter 7, we investigate the equilibrium properties of Bloch Points under applied fields in bi-layered magnetic nanodisks with opposing chirality. We then study the field-driven motion of Bloch Points in FeGe nanotacks, and calculate the Bloch Point velocity under such a field. In Chapter 8, we conclude and summarise the work.

Chapter 2

Magnetic Theory and Background

2.1 Introduction

On a macroscopic level, magnetism has been well understood for a relatively long time, with work by Maxwell [69] allowing predictive calculations on macroscopic scales to be made. However, the understanding of spontaneous ferromagnetism on an atomic scale necessitates the use of quantum theory, and was thus not developed until much later. Here, some basic exposition of the necessary concepts needed to understand magnetic modelling is made, detailing the origins of magnetism, the classical spin models which derive from the quantum origin of magnetism, and of micromagnetism, a continuum theory of magnetic materials. Within this, reference is made to the equations of motion governing the magnetisation, and to the computational techniques used to study magnetic behaviour on a range of length scales.

2.2 Origins of Magnetism

All atoms have electrons which are in motion around a central nucleus. Stern and Gerlach measured in 1922 the magnetic dipole moment of silver atoms by deflecting them through a magnetic field gradient [70]. It was found that the deflected atoms split into two bands. From quantum predictions, it was expected that the magnetic dipole moment of an atom μ would split into multiple bands along a measured component, according to the discrete orbital angular momenta

states as [71]:

$$\boldsymbol{\mu} = \gamma \mathbf{L} \quad (2.1)$$

where:

$$\mathbf{L} = \hbar \sqrt{l(l+1)} \quad (2.2)$$

$$L_z = \hbar m_l \quad (2.3)$$

and where $m_l = -l, -l+1, \dots, l-1, l$ and where the total moment is given by:

$$\mu^L = \mu_B \sqrt{l(l+1)} \quad (2.4)$$

Phipps and Taylor later showed in 1926 [72] that by scattering hydrogen at a low temperature such that each atom should sit in the ground state, the atom still showed an angular momentum component, which we call spin, and which transforms as the orbital angular momenta. For an electron, the spin $s = 1/2$, and the $m_s = -1/2, 1/2$ so that:

$$\mathbf{S} = \hbar \sqrt{s(s+1)} \quad (2.5)$$

$$S_z = \hbar m_s \quad (2.6)$$

and then the dipole moment of an electron in free space is given as:

$$\boldsymbol{\mu} = \mu_B (g_l \mathbf{L} + g_s \mathbf{S}) \quad (2.7)$$

where the dimensionless gyromagnetic ratios (or sometimes g-factors) g_l and g_s are constants which determine the relative contribution of a particle's angular momentum and spin to the dipole moment. Dirac showed in 1928, with the introduction of the equation that came to bear his name, that a quantum mechanical electron must have $g_s = 2$ and that the spin of the electron is $1/2$ [73]. Repeatedly, highly accurate experimental measurements have shown that for the free electron, $g_l = 1$ and $g_s \approx 2$, [74] which accords with more modern predictions from quantum electrodynamics.

As a starting point for many quantum mechanical calculations, the an approximation developed by Born and Oppenheimer is made, whereby the mass of nuclei is effectively infinite, and the motion of the nucleus ignored. When electrons are in motion around the nucleus, in the rest frame of the electron, the charged nucleus is in motion around the electron [71], which leads to the electron feeling a magnetic field due to its own angular momentum. The behaviour of the spin coupled with

this field generating from the electron orbit is known as a spin-orbit coupling. It was shown by Llewellyn Thomas in 1926 [75, 76] that the magnetic field felt by an electron due to the nucleus is given as:

$$\mathbf{B} = \frac{1}{emc^2 r} \frac{1}{r} \frac{dU(r)}{dr} \mathbf{L} \quad (2.8)$$

where $U(r)$ is the interaction between the nucleus and the electron (the Coulomb equation for a simple Hydrogen atom), and thus the energy of the spin-orbit interaction is given as:

$$\Delta E = -\boldsymbol{\mu}_s \cdot \mathbf{B} \quad (2.9)$$

$$= \frac{g_s \mu_B}{2emc^2 \hbar r} \frac{1}{r} \frac{dU(r)}{dr} \mathbf{L} \cdot \mathbf{S} \quad (2.10)$$

For even simple atoms, this leads to complex behaviour, with electron energy levels splitting due to relative orientations of \mathbf{S} and \mathbf{L} - a phenomena known as 'fine structure.'

In addition, further splitting of energy levels can occur upon the application of an external field; this is known as the Zeeman Effect after Zeeman who observed the splitting of spectral lines of light from a flame placed in an electromagnet in the late 1890s [77]. If a field B_z is applied such that it is along the measurement axis, the energy change is given as:

$$\Delta E_{\text{Zeeman}} = \gamma(g_L L_z + g_S S_z) B_z \quad (2.11)$$

If the spin-orbit field, applied external field, the total angular momentum for each electron can be significantly perturbed, and this leads to changes in the magnetic moment.

We have not yet discussed the influence of electron-electron interactions, and it is this which is key to understanding how atoms interact to cause ferromagnetism. In a simple quantum mechanical model, electrons around atoms interact through Coulombic repulsion, and through the quantum mechanical Pauli Exclusion principle cannot be interchanged. The interaction energy of two nearby atoms can be given in the form:

$$\tilde{E}_0 = E_0 \pm J_D \quad (2.12)$$

where E_0 is the ground state energy of an electron in the single atom state and J_D is the exchange coefficient, and

$$\Delta E = \frac{1}{2}(E_{\text{Ferromagnetic}} - E_{\text{Antiferromagnetic}}) \quad (2.13)$$

whereby if $J_D > 0$, an anti-ferromagnetic state is favoured, with the dipole moments of neighbouring atoms preferring to be antiparallel. On the contrary, with $J < 0$, the ferromagnetic state is preferred. By transformation, this can be described as the well known Heisenberg Exchange term:

$$E_{ij} = -J_{ij} \langle S_i \cdot S_j \rangle \quad (2.14)$$

which describes the interaction of two atoms.

The solution of the Schrödinger equation for a Hamiltonian describing the interactions between many electrons and nuclei is not solvable analytically. Computational techniques such as the Hartree-Fock method [78] and the later Density Functional Theory (DFT) [79, 80] attempt to find approximate solutions which describe the electronic structure of different types of materials, determining in the process the magnetic moment for each atomic species, and the strength of the Exchange interaction between atomic species, among other quantities. These methods can be highly sensitive to the choice of basis functions chosen to represent the electron wave function. For many materials of interest in magnetism, these models can reproduce phenomenological behaviour, but significantly under or overestimate values of the magnetic moment compared to experimental measurements [81]. In this case, other alternative methods must be used to make predictive calculations about magnetic materials, often parameterised from experimental measurements of material properties rather than through the quantum computations.

2.3 Heisenberg Model and Micromagnetics

It is possible to describe the properties and behaviour many magnetic materials by ignoring the quantum behaviour completely, and by using a completely classical model, whereby each atomic dipole is a classical vector. Hereafter, the term ‘spin’ is used interchangably to mean ‘atomic dipole moment’, as is common in the literature in this field. Lenz, starting from the assumption that preferential angles for magnetic dipoles on a crystal lattice would occur in some cases due to large

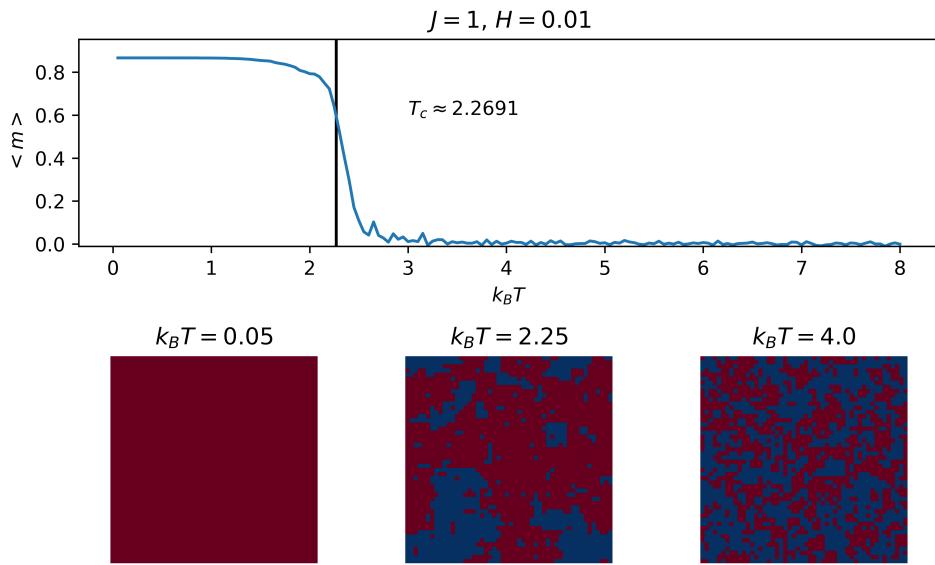


FIGURE 2.1: Spontaneous magnetism arising in the 2-D square-lattice Ising model at low temperatures for a 50×50 system with $J = 1$ and a low field of $H = 0.001$, using the Metropolis-Hastings algorithm and simulated annealing. For each temperature value, 5×10^5 trial steps are used, and the results are averaged over 30 runs. With the very weak field used simply to break the symmetry of the system, good agreement is shown with the transition temperature as calculated by Onsager, as $k_B T_c = \frac{2J}{\ln(1+\sqrt{2})} \approx 2.2691$.

potential energies [82], formulated what we now know as the classical ‘Ising Model’, due to the publication on this work by Ising, his student, who calculated the partition function for the model in one dimension [55]. In this model, each atomic dipole was coupled only to its nearest neighbour through a classical exchange interaction, and to an external field H such that the total energy of a system of spins was given as:

$$\mathcal{H} = -J_{ij} \sum_{i \neq j} S_i \cdot S_j - H \sum_i S_i \quad (2.15)$$

Where S_i has only two possible states - +1 or -1. Ising showed that in one-dimension lattice (i.e. a spin chain), ferromagnetic order could not be sustained for finite temperature, though it was later shown by Onsager that ferromagnetic order could be sustained in the two-dimensional square lattice Ising model [56].

2.3.1 Heisenberg Exchange

Generalising to three dimensions gives the classical Heisenberg model, where one considers sets of spin vectors $\boldsymbol{\mu}_i = \mu_i \mathbf{S}_i$. Each spin is again coupled to its neighbours through an exchange interaction, under the influence from an external field \mathbf{h} [57]:

$$\mathcal{H} = - \sum_{i \neq j} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j \quad (2.16)$$

An external field can be coupled to the spins, known as the Zeeman term, and now, this is also a vector quantity:

$$\mathcal{H}_{\text{Zeeman}} = - \sum_i \mu_i \mathbf{S}_i \cdot \mathbf{H} \quad (2.17)$$

It is straightforward to consider many atoms grouped into small volumes of space, such that:

$$\mathbf{m}(r) = \frac{\sum_i \mu_i}{V} \quad (2.18)$$

We can easily take the continuum limit and derive the Exchange interaction for a continuum magnetisation [61, 83]. Consider a one dimensional crystal of identical spins with an interatomic spacing a , and with the same coupling strength J between all spins. Assume that one can replace the spin unit vectors with a unit vector field $\mathbf{m}(r) = \mathbf{S}(r)/S$ that is identical to the original spin vectors at their locations. The Hamiltonian can be separated out so that under the sum is the total contribution to the energy for a given particle:

$$\mathcal{H} = -J \sum_i \mathbf{S}_{i-1} \cdot \mathbf{S}_i + \mathbf{S}_i \cdot \mathbf{S}_{i+1} \quad (2.19)$$

$$= -JS^2 \sum_i \mathbf{m}(x_i - a) \cdot \mathbf{m}(x_i) + \mathbf{m}(x_i) \cdot \mathbf{m}(x_i + a) \quad (2.20)$$

We now assume that $\mathbf{m}(x)$ varies slowly, noting that the assumption of slow variation required means that only ferromagnetic materials can be considered in the approximation; both antiferromagnetism and ferrimagnetism must be treated with the Heisenberg formulation directly. The small deviations from x can be expanded

out as a Taylor series, and odd powers of a cancel:

$$\begin{aligned}\mathcal{H} &= -JS^2 \sum_i \mathbf{m}_i \cdot \left[\mathbf{m}(x_i) + a \frac{\partial \mathbf{m}(x_i)}{\partial x} + \frac{a}{2} \frac{\partial^2 \mathbf{m}(x_i)}{\partial x^2} \right. \\ &\quad \left. + \mathbf{m}(x_i) - a \frac{\partial \mathbf{m}(x_i)}{\partial x} + \frac{a^2}{2} \frac{\partial^2 \mathbf{m}(x_i)}{\partial x^2} + \mathcal{O}(a^3) \right] \end{aligned}\quad (2.21)$$

$$= -JS^2 \sum_i \mathbf{m}(x_i) \cdot \left[2\mathbf{m}(x_i) + a^2 \frac{\partial^2 \mathbf{m}(x_i)}{\partial x^2} + \mathcal{O}(a^4) \right] \quad (2.22)$$

The constant term can be neglected, as it simply shifts the energy, and the exchange can then be approximated by taking dropping the higher order terms:

$$\mathcal{H} \approx -a^2 JS^2 \sum_i \mathbf{m}(x_i) \cdot \frac{\partial^2 \mathbf{m}(x_i)}{\partial x^2} dx \quad (2.23)$$

Along the chain, the number of atoms in a line segment dx is given as $\frac{n dx}{a}$, where n is the number atoms in a unit cell. Taking the continuum limit, we then get:

$$\mathcal{H} \approx -naJS^2 \int \mathbf{m}(x_i) \cdot \frac{\partial^2 \mathbf{m}(x_i)}{\partial x^2} dx \quad (2.24)$$

Twice differentiating the expression $\mathbf{m}(x) \cdot \mathbf{m}(x) = 1$ gives the expression $\mathbf{m} \cdot \frac{\partial^2 \mathbf{m}}{\partial x^2} = -\left(\frac{\partial \mathbf{m}}{\partial x}\right)^2$, and substituting this gives a more commonly used form:

$$\mathcal{H} = A_{1d} \int \left(\frac{\partial \mathbf{m}}{\partial x} \right)^2 dx \quad (2.25)$$

with $A_{1d} = anJS^2$. This expression can be generalised straightforwardly to higher dimensions by simply considering nearest neighbours in other directions, and replacing n with the number of atoms in the unit cell in that dimension; in two dimensions it is:

$$\begin{aligned}\mathcal{H} &= \frac{JS^2}{n} \iint \left(\frac{\partial \mathbf{m}}{\partial x} \right)^2 + \left(\frac{\partial \mathbf{m}}{\partial y} \right)^2 dx dy \\ &= A_{2d} \iint \left(\frac{\partial \mathbf{m}}{\partial x} \right)^2 + \left(\frac{\partial \mathbf{m}}{\partial y} \right)^2 dx dy \end{aligned}\quad (2.26)$$

with $A_{2d} = JS^2n$. For the square lattice, $n = 1$ and for the triangular lattice $n = 2$. In three-dimensions:

$$\begin{aligned}\mathcal{H} &= \frac{aJ}{n} \iiint \left(\frac{\partial \mathbf{m}}{\partial x} \right)^2 + \left(\frac{\partial \mathbf{m}}{\partial y} \right)^2 + \left(\frac{\partial \mathbf{m}}{\partial z} \right)^2 d^3r \\ &= A_{3d} \int (\nabla \mathbf{m})^2\end{aligned}\quad (2.27)$$

with $A_{2d} = \frac{nJS^2}{a}$. For a simple cubic unit cell, $n = 1$, while for a BCC crystal $n = 2$ and for FCC $n = 4$.

2.3.2 Dzyaloshinskii-Moriya Exchange

In many systems, more complex exchange behaviour can be observed, whereby spins align not parallel but with small twists relative to neighbouring spins. Dzyaloshinskii [84, 84] first formulated this type of interaction from symmetry considerations in micromagnetics, and this was later rationalised with quantum mechanics by Moriya [85]. Superexchange occurs between magnetic ions separated by non-magnetic atomic species, and is so called ‘super’ because of the distance over which it acts, and the behaviour is described by the Hubbard Hamiltonian. When superexchange and spin-orbit couplings are included in calculations, a non-negligible term known as Dzyaloshinskii-Moriya (DM) interaction leads to the observed anisotropic exchange. The resulting behaviour of competition between the Exchange and DM interaction leads to helices, cones and skyrmions, which will be discussed in later chapters. In general the form of the DM interaction can be written as:

$$\mathcal{H}_{DM} = \frac{1}{2} \sum_{i \neq j} \mathbf{d}_{ij} \cdot \mathbf{S}_i \times \mathbf{S}_j \quad (2.28)$$

The form of \mathbf{d}_{ij} depends on the type of material under study. In crystals, when a broken lattice symmetry occurs due to non-centrosymmetric unit cells, the DM tensor can be written as:

$$\mathbf{d}_{ij} = d\hat{\mathbf{r}}_{ij} \quad (2.29)$$

which is known as the Bulk DM.

The continuum description of this behaviour can be derived in a similar way to the case for the Exchange, except this time the even-ordered derivatives cancel because of the $\hat{\mathbf{r}}$ vector. In one-dimension:

$$\begin{aligned}
\mathcal{H} &= \frac{dS^2}{2} \sum_i (-\hat{\mathbf{x}}) \cdot [\mathbf{m}_i \times \mathbf{m}_{i-1}] + \hat{\mathbf{x}} \cdot [\mathbf{m}_i \times \mathbf{m}_{i+1}] \\
&= \frac{dS^2}{2} \hat{\mathbf{x}} \cdot \sum_i - \left[\mathbf{m}(x_i) - a \frac{\partial \mathbf{m}(x_i)}{\partial x} + \frac{a}{2} \frac{\partial^2 \mathbf{m}(x_i)}{\partial x^2} \right] \\
&\quad + \left[\mathbf{m}(x_i) + a \frac{\partial \mathbf{m}(x_i)}{\partial x} + \frac{a}{2} \frac{\partial^2 \mathbf{m}(x_i)}{\partial x^2} \right] \\
&= daS^2 \hat{\mathbf{x}} \cdot \sum_i \mathbf{m}(x_i) \times \frac{\partial \mathbf{m}(x_i)}{\partial x} \\
&= daS^2 \sum_i \left[\frac{m_y(x_i) \partial m_z(x_i)}{\partial x} \partial x \right] \\
&= daS^2 \int \left[m_y(x) \frac{\partial m_z(x)}{\partial x} - m_z(x) \frac{\partial m_y(x)}{\partial x} \right] \left(\frac{n}{a} \right) dx \\
&= D_{1d} \int m_y(x) \frac{\partial m_z(x)}{\partial x} - m_z(x) \frac{\partial m_y(x)}{\partial x} dx
\end{aligned} \tag{2.30}$$

where $D_{1d} = dS^2 n$. In three-dimensions, like terms can be collected in the components of \mathbf{m} , and the expression becomes:

$$\begin{aligned}
\mathcal{H} &= daS^2 \int \left[m_y \frac{\partial m_z}{\partial x} - m_z \frac{\partial m_y}{\partial x} + m_z \frac{\partial m_x}{\partial y} - m_x \frac{\partial m_z}{\partial y} + m_x \frac{\partial m_y}{\partial z} - m_y \frac{\partial m_x}{\partial z} \right] \left(\frac{n}{a^3} \right) d^3r \\
&= \frac{dnS^2}{a^2} \int \left[m_x \left(\frac{\partial m_y}{\partial z} - \frac{\partial m_z}{\partial y} \right) + m_y \left(\frac{\partial m_z}{\partial x} - \frac{\partial m_x}{\partial z} \right) + m_z \left(\frac{\partial m_x}{\partial y} - \frac{\partial m_y}{\partial x} \right) \right] d^3r \\
&= D_{3d} \int \mathbf{m} \cdot (\nabla \times \mathbf{m}) d^3r
\end{aligned} \tag{2.31}$$

with $D_{3d} = \frac{dnS^2}{a^2} n$.

For other systems, a DM interaction has its origin due to spin-orbit coupling between layers of different atoms [26, 30]. For this type of interaction, the form of the DM tensor is given as:

$$\mathbf{d}_{ij} = d\hat{\mathbf{r}}_{ij} \times \hat{\mathbf{z}} \tag{2.32}$$

For this interfacial DMI, the micromagnetic interaction can be derived similarly [30] and results in the form:

$$\mathcal{H} = D \iiint \left[\left(m_y \frac{\partial m_z}{\partial x} \right) - \left(m_x \frac{\partial m_z}{\partial y} - m_z \frac{\partial m_x}{\partial y} \right) \right] \tag{2.33}$$

It is important to note that between publications, the definition of the constant

D varies in sign, [86] which leads to the opposite sense of rotation for a given material to what one may expect.

2.3.3 Anisotropy

In many materials, inherent preferred directions exist within crystals for which the atomic dipoles align. In the simplest case, there can be a single axis along which this is felt, and this is known as uniaxial anisotropy. This can be described in the Heisenberg framework as:

$$\mathcal{H}_{\text{Anis}}^{\text{Uni}} = - \sum_i k_u (\mathbf{S}_i \cdot \mathbf{u})^2 \quad (2.34)$$

In other materials, cubic anisotropy [58] is the dominant form of anisotropy:

$$\mathcal{H}_{\text{Anis}}^{\text{Cub}} = \sum_i k_c (S_x^4 + S_y^4 + S_z^4) \quad (2.35)$$

which results in spins aligning preferentially along each of the three coordinate axes. The Uniaxial and Cubic anisotropies originate again in spin-orbit couplings between the electrons and the nuclei. Other forms of anisotropy such as pressure or strain induced anisotropy can be incorporated into models in a similar way, but result from the distance dependence of the Exchange integrals. The micromagnetic form is simple here, because the individual forms are not coupled to one another. In this case then, the continuum form is just given by integrating over a small volume:

$$\mathcal{H}_{\text{Anis}}^{\text{Uni}} = K_u \iiint \mathbf{m} \cdot \mathbf{u} d^3r \quad (2.36)$$

2.3.4 Dipolar Interaction

The magnetic behaviour of large systems of atomic dipoles is incomplete without the inclusion of a term describing the direct long-range interaction between the atomic dipoles moments, which are treated as point sources:

$$\mathcal{H}_{\text{Dipolar}} = \frac{\mu_0}{4\pi} \sum_i \sum_{j \neq i} \mu_i \mu_j \frac{3(\mathbf{S}_i \cdot \hat{\mathbf{r}}_{ij})(\mathbf{S}_j \cdot \hat{\mathbf{r}}_{ij}) - \mathbf{S}_i \cdot \mathbf{S}_j}{r_{ij}^3} \quad (2.37)$$

For micromagnetics, it is most straightforward to consider the magnetostatic part of Maxwell's equations to come to a form describing the internal dipolar field (normally known as the demagnetising field) for a volume, and formulate a boundary value problem rather than approach the problem from the other direction. Starting from:

$$\nabla \times \mathbf{H} = \mathbf{J} \quad (2.38)$$

$$\nabla \cdot \mathbf{B} = 0 \quad (2.39)$$

If there is no applied current in a system, then \mathbf{H} can then be written in terms of a magnetic scalar potential:

$$\mathbf{H} = -\nabla\Phi \quad (2.40)$$

If the material is linear, then $B = \mu_0(\mathbf{H} + \mathbf{M})$. Substituting into Equation 2.39, then:

$$\nabla^2\Phi = -\nabla \cdot \mathbf{M} \quad (2.41)$$

The full solution for this equation is given as:

$$\Phi(\mathbf{r}) = \frac{-1}{4\pi} \int_V \frac{\nabla' \cdot \mathbf{M}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d^3r' + \frac{1}{4\pi} \oint_S \frac{\mathbf{n}' \cdot \mathbf{M}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} dS \quad (2.42)$$

from which the gradient can be taken to give the field as in Eq. 2.40. The total micromagnetic energy is then:

$$\mathcal{H} = \int \mathbf{M} \cdot \mathbf{H} d^3r \quad (2.43)$$

2.3.5 Competition between energies, characteristic length scales and emergent phenomena

The combination of energies and their relative strengths in magnetic materials depends on their atomic constitution and the resulting electron density within the material. In magnetic systems, we observe not only ferromagnetism and antiferromagnetism, which occur when the exchange interaction is dominant, but a whole range of other phenomena over different length scales.

Perhaps the most fundamental of these length scales is that of the Exchange length. This is determined by the relative strengths of the exchange interaction and the dipolar interaction. This is usually given determined from the micromagnetic Exchange strength and the magnetostatic anisotropy [87, 88]

$$l_{\text{ex}} = \sqrt{\frac{A}{\frac{1}{2}\mu_0 M_s^2}} \quad (2.44)$$

The exchange length can give an indication of what features we may expect to observe in a system of a given size for a particular material. Below this size, the exchange interaction dominates over the magnetostatic field, and so we can expect systems on this order to be nearly uniformly magnetised. Small particles are therefore treated with Stoner-Wohlfarth model, assuming no exchange interaction between particles, and with the assumption that the magnetisation within such a particle rotates coherently (i.e. as a whole). Classical spin-dynamics can be used to describe the magnetic evolution of such particles as though they were individual spins [89]. Above the exchange length, it becomes costly to keep the entire system magnetised, and we expect that in the absence of applied fields, the magnetisation of a sample will form into multiple domains, separated by domain walls. These form one of two types; when the magnetisation varies across the domain wall in the out-of-plane direction relative to the local magnetisation, then the wall is classified as a Bloch wall, while if the magnetisation varies across the wall within the plane, it is a Ne  l wall, which are commonly observed in thin-film systems. The presence of a crystalline anisotropy acts to change the energy balance of such multi-domain structures, and leads to the narrowing of domain walls. Magnetic vortices are another phenomena which emerge from the combination of the magnetostatic energy density and exchange energy, and appear like a ‘tornado’ in the magnetisation with the magnetisation of the vortex core pointing along the core axis, and the magnetisation outside of the core rotating around it.

When the combination of Dzyaloshinskii-Moriya, Exchange and Demagnetising interactions are present, in both bulk and thin-film systems, lattices of magnetic skyrmions can form. These form whirls similar to vortices, but in contrast, undergo a full rotation from core to edge. Skyrmion physics is discussed in detail in Chapters 5 and 6.

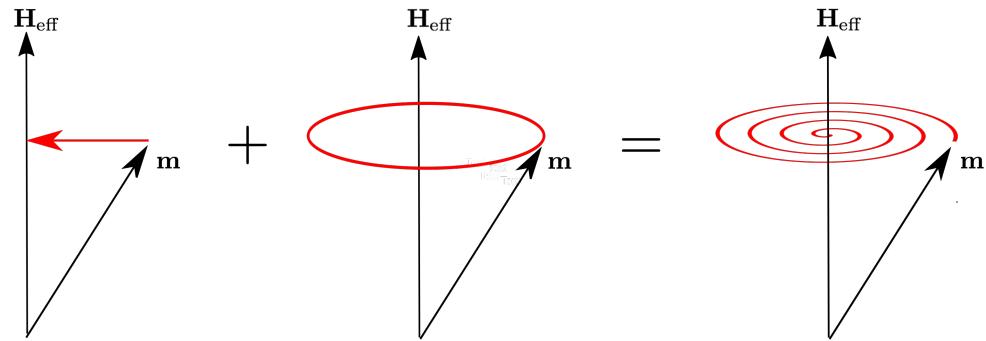


FIGURE 2.2: Figure showing the combined effects of (a) The damping term of the LLG equation and (b) The precessional term of the LLG equation on a magnetic vector \mathbf{m} under the influence of the effective field \mathbf{H}_{eff} . Over time, the magnetisation in a system tends towards the effective field and lose energy through the influence of the damping term.

2.3.6 Equation of Motion

A field acting on a magnetic moment acts to induce a torque on it, and so the equation of motion for a magnetic moment can be given as:

$$\frac{d\boldsymbol{\mu}}{dt} = -\gamma\boldsymbol{\mu} \times \mathbf{B} = -\gamma_0\boldsymbol{\mu} \times \mathbf{H} \quad (2.45)$$

where $\gamma_0 = \gamma\mu_0$, and γ is the gyromagnetic ratio. The torque acts perpendicular to the applied field and the moment itself, causing the moment to precess around the applied field; this is known as Larmor precession. The frequency can be straightforwardly determined and is given as $f_L = \frac{\mu_0 H}{2\pi}$. In the continuum limit, the equation is simply:

$$\frac{d\mathbf{M}}{dt} = -\gamma_0 [\mathbf{M} \times \mathbf{H}] \quad (2.46)$$

From each of the terms in the Heisenberg and micromagnetic energy formulations it is a useful concept to consider an ‘Effective Field’ which acts on each site, defined as a variational derivative of the energy:

$$\mathbf{H}_{\text{eff}} = -\frac{1}{\mu} \frac{\delta E}{\delta \boldsymbol{\mu}} \quad (2.47)$$

or in micromagnetics:

$$\mathbf{H}_{\text{eff}} = -\frac{1}{\mu_0 M_s} \frac{\delta E}{\delta \mathbf{m}} \quad (2.48)$$

The variational derivative can be computed with respect to a component of a field f_x using the Euler-Lagrange formula:

$$\frac{\delta E}{\delta f_x} = \frac{\partial E}{\partial f_x} - \frac{\partial}{\partial x} \left(\frac{\partial E}{\partial (\partial_x f_x)} \right) - \frac{\partial}{\partial y} \left(\frac{\partial E}{\partial (\partial_y f_x)} \right) - \frac{\partial}{\partial z} \left(\frac{\partial E}{\partial (\partial_z f_x)} \right) \quad (2.49)$$

The torque acting on each moment (or magnetised region) then feels is then a combination of terms from the local internal fields and external fields.

This is not, however, the full picture. In crystalline materials, the motion of magnetic moments is damped. Landau and Lifshitz [60] introduced in 1935 an additional term to describe damping in the micromagnetic formulation, to model damping due to spin-orbit coupling (in the original work this is labelled as a ‘relativistic’ contribution). This term acts to cause the magnetisation to relax towards the field direction over time:

$$\frac{\partial \mathbf{m}}{\partial t} = -\gamma_0 \mathbf{m} \times \mathbf{H}_{\text{eff}} - \lambda_{LL} \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}}) \quad (2.50)$$

This equation is known as the Landau-Lifshitz equation, and λ_{LL} is the phenomenologically determined Landau-Lifshitz damping constant.. Gilbert could not match experimental data in thin ferromagnetic sheets to the Landau-Lifshitz damping, and so introduced phenomenologically a Rayleigh Dissipation Function which damps the motion by way of a term which is proportional to the square of the rate of change of magnetisation [62, 90], resulting in the following form:

$$\frac{\partial \mathbf{m}}{\partial t} = -\gamma_0^* \mathbf{m} \times \mathbf{H}_{\text{eff}} + \alpha \frac{\partial \mathbf{m}}{\partial t} \quad (2.51)$$

This equation is known as the Landau-Lifshitz-Gilbert equation. Equivalence between this and the Landau-Lifshitz equation can be shown through taking the cross product of both sides of this equation with the magnetisation and simplifying using the vector triple product identity, giving:

$$\frac{\partial \mathbf{m}}{\partial t} = \frac{\gamma_0^*}{1 + \alpha^2} \mathbf{m} \times \mathbf{H}_{\text{eff}} - \frac{\gamma_0^* \alpha}{1 + \alpha^2} \mathbf{m} \times \mathbf{m} \times \mathbf{H} \quad (2.52)$$

with equivalence satisfied if we assume that:

$$\gamma_0 = \frac{\gamma_0^*}{1 + \alpha^2} \quad (2.53)$$

$$\lambda = \frac{\gamma_0^* \alpha}{1 + \alpha^2} \quad (2.54)$$

Many other contributions to the damping term aside from spin-orbit coupling have been considered, such as magnon decay [91, 92], but the full microscopic details and relative contributions of damping mechanisms are not yet well understood, with theoretical predictions of the damping largely over or underestimating the damping coefficient α . Because of this, the constant is generally determined experimentally for a given material from linewidths in ferromagnetic resonance experiments [93, 94].

Chapter 3

Computational Methods in Nanomagnetics

3.1 Introduction

Numerical solution of the atomistic and micromagnetic equations is absolutely essential to the modern study of magnetic materials. While some simple models can be solved analytically or semi-analytically, this is of little help when considering the large systems and dynamic behaviour of magnetic materials which are of interest in the experimental and industrial worlds.

To find the dynamics of spins atomistically, it is usual to consider regular lattices of spins of some form. In micromagnetics, however, the continuum field theory must be rediscretised in order to find solutions. There are two main approaches used to do this - the finite difference and finite element methods. The energetics in atomistic models are relatively straightforward to implement computationally, as each term can be directly implemented to act on column vectors representing each spin. The equations describing the micromagnetic energy equations, which are written in terms of continuous fields, must be rediscretised before they can be computed, however, by dividing any magnetic domain under study into small pieces. This may seem pointless, given that atomistic models are already discrete, but the beauty is that the length scales over which the continuous fields are discretised are much larger than individual spins, which allows the study of much larger systems. In general, the approach taken for all algorithms, both atomistic and micromagnetic is to find the effective field \mathbf{H}_{eff} . From this, the field is updated,

either through direct integration of the LLG equation, or through a standard energy minimisation algorithm such as the steepest descent or conjugate gradient methods.

In both micromagnetism and atomistic modelling, the dipolar field is a major complication, which requires special handling. This is because of both the long range character, and the length scales over which atomistic and micromagnetic simulations are carried out; it is not a good approximation to cut off the interaction beyond a certain distance as is sometimes done in molecular dynamics. The long range character means that all simulations are limited in scaling by the dipolar field calculation, because all magnetic regions interact with all other magnetic regions.

3.2 Atomistic Models

In order to compute the effective field in atomistic models numerically, it is usual (though not necessary) to consider a regular grid of atoms, placed at equal intervals. In order to store this information, and the information associated with each lattice site, it is useful to consider a cubic grid and the indexing scheme which must be used. Generally, it is useful to store the required information as a column vector. There are many ways to implement this, and here the method used in the software Fidimag is described.

For atoms on a three-dimensional atomic cubic lattice, we can consider each atom with a lattice index of (i, j, k) which acts as a useful proxy as to their position in real space. In three dimensions, with $N = N_x N_y N_z$ atoms, an index of:

$$I_{\text{Scalar}} = kN_x N_y + jN_x + i \quad (3.1)$$

can be used for indexing into arrays of scalar quantities, such as the spin magnitude. This can be applied to two dimensional systems simply by setting $k = 0$ everywhere.

For a vector quantity of fixed length, such as the magnetisation which has $L = 3$ components, the l^{th} component for the can be found in a length NL column vector at the location:

$$I_{\text{Vector}}^l = L(kN_x N_y + jN_x + i) + l \quad (3.2)$$

In this way then, a series of blocks of memory can be constructed to represent each of the scalar and vectors for each atom. How then, are interactions between atoms evaluated, once the magnetic spin directions? It is straightforward to construct a list of neighbouring spins by considering the $+x, -x, +y, -y, +z, -z$ directions. For an atom at (i, j, k) on a cubic lattice, there are a maximum of six nearest neighbours, which can be found at locations:

$$\begin{aligned}
 I_{-x} &= kN_x N_y + jN_x + (i - 1) \\
 I_{+x} &= kN_x N_y + jN_x + (i + 1) \\
 I_{-y} &= kN_x N_y + (j - 1)N_x + i \\
 I_{+y} &= kN_x N_y + (j + 1)N_x + i \\
 I_{-z} &= (k - 1)N_x N_y + jN_x + i \\
 I_{+z} &= (k + 1)N_x N_y + jN_x + i
 \end{aligned} \tag{3.3}$$

For periodic systems, this can easily be modified to take into account the connectivity at the system boundaries, and for non-periodic systems, each atom simply has fewer neighbours. To construct lattices which are non-cubic, it is useful to map onto a regular cubic grid, leaving vacancies. The neighbours in all cases can either be calculated heuristically on-the-fly or precomputed and stored in neighbour lists which determine between which spins the local fields such as the Exchange and Dzyaloshinskii-Moriya are computed in advance. The procedure for calculating a particular interaction then, is generally:

Algorithm 1 ComputeInteractionField()

```

 $N$  = number of particles
for Particle  $i \in [0, N - 1]$  do
    Get spin  $\mu_s^i$ 
     $\mathbf{H}_{\text{eff}}^i = 0$ 
    for all neighbouring particles  $j$  of particle  $i$  do
        Get coefficient of interaction (e.g.  $J_{ij}, D_{ij}$ )
         $\mathbf{H}_{\text{eff}}^i +$  = Interaction between particles  $\mu_i$  and  $\mu_j$ 
    end for
end for

```

For interactions in which the next-nearest neighbours are needed, such as higher-order exchange interactions, the neighbour list for each spin is simply extended to take this into account.

The dipolar field is handled somewhat differently. In a naive implementation, the dipolar field is a $\mathcal{O}(n^2)$ calculation with the number of particles, because each spin must interact with all other spins in the system, and so the cost grows enormously as larger problem sizes are tackled. In many atomistic simulations, the dipolar field is neglected entirely, despite it's relevance in many physical systems and this having been shown to affect even small systems. In others [58], approximate techniques are utilised in order to decrease the computational power. In some atomistic simulation software, including Fidimag [95, 96], the dipolar field calculation is implemented using a convolution. This technique was applied to spin dynamics in two dimensional systems by Mansuripur and Giles [97], and this was later extended to three-dimensional systems by Yuan and Bertram [65, 98]. The technique appears to have been developed independently, but it is in effect similar in form to a simple version of the Particle-Mesh algorithm, which was already in widespread use in electrostatic calculations in the 1970s [99], only without an interpolation step due to particles sitting on exact points on the lattice. The technique is applicable both to non-periodic and periodic systems where the system is mappable to a regular grid, and when the system is non-periodic, the technique is exact up to machine precision. One downside of this approach is that the restriction to regular grids is not appropriate for all atomistic-like systems, such as in coupling chemical molecular dynamics with spin dynamics [100].

We start generally, so as to show how the technique can be applied to different types of system. We denote the position of a source on a lattice of dimensions $n_x \times n_y \times n_z$ as r_a^{ijk} with the superscript ijk representing the real space position $(i a_x, j a_y, k a_z)$, where a_x, a_y, a_z are the lattice spacings. The interaction between a source and all other sources on the lattice can be written as:

$$G(i, j, k) = \sum_{i \neq i'} \sum_{j \neq j'} \sum_{k \neq k'} K((i - i')a_x, (j - j')a_y, (k - k')a_z) S(i', j', k') \quad (3.4)$$

Where $i - i' \in [-n_x + 1, n_x - 1]$, $j - j' \in [-n_y + 1, n_y - 1]$, and $k - k' \in [-n_z + 1, n_z - 1]$, G is a scalar field, $K(x, y, z)$ is a kernel describing the interaction, and S is the source strength. This is clearly the form of a 3-D discrete convolution. It does not seem obvious that we have thus far gained anything from this mathematics, and if we evaluate the relation directly on a computer then would we expect the same performance as a simple summation technique. However, fast algorithms such as the Cooley-Tukey algorithm [101] exist for computing the discrete Fourier transform, and this can be applied to discrete convolution. These fast algorithms

for computing the discrete Fourier transform are known collectively as Fast Fourier Transforms (FFTs), and they scale with the number of array elements N as $\mathcal{O}(N \log N)$, though performance is greater when N is a product of low value prime integers (i.e. $2^a 3^b 5^c$).

A convolution of two functions f and g is given as:

$$f * g = \int_{-\infty}^{\infty} f(\tau)g(t - \tau)d\tau \quad (3.5)$$

However, the convolution theorem relates that:

$$f * g = \mathcal{F}^{-1}(\mathcal{F}(f)\mathcal{F}(g)) \quad (3.6)$$

This property is strictly true only for data which is continuous. However, it can be applied to discrete data, using the discrete Fourier Transform so long as one of the functions in the convolution is a periodic summation of another function, such that:

$$f^i = \sum_{k=-N}^{k=N} p^{i-kN} \quad (3.7)$$

These relations generalise straightforwardly to convolutions in two and three dimensions.

For a non-periodic function, the data must be ‘zero-padded’ so that aliasing effects from the FFT do not occur; one does this by padding the size of the array to at least $2N - 1$ elements. One can pad further, usually to a power of two in order to see the performance benefits from more efficient FFT algorithms, and sometimes see the computation time decrease despite the increase in the number of elements.

Writing the potential and field for dipoles in such a system is more involved than for the general case, because the dipole source is a vector, and each dipole component contributes to the each component of the field. We can write out the calculation

of the potential and field as a matrix multiplication to show this explicitly:

$$\begin{bmatrix} \Phi \\ F_x \\ F_y \\ F_z \end{bmatrix} = \begin{bmatrix} \frac{x}{r^3} & \frac{y}{r^3} & \frac{z}{r^3} \\ \frac{3x^2}{r^5} - \frac{1}{r^3} & \frac{3xy}{r^5} & \frac{3xz}{r^5} \\ \frac{3yx}{r^5} & \frac{3y^2}{r^5} - \frac{1}{r^3} & \frac{3yz}{r^5} \\ \frac{3xz}{r^5} & \frac{3yz}{r^5} & \frac{3z^2}{r^5} - \frac{1}{r^3} \end{bmatrix} \begin{bmatrix} \mu_x \\ \mu_y \\ \mu_z \end{bmatrix} \quad (3.8)$$

$$= \begin{bmatrix} S_x & S_y & S_z \\ N_{xx} & N_{xy} & N_{xz} \\ N_{yx} & N_{yy} & N_{yz} \\ N_{zx} & N_{zy} & N_{zz} \end{bmatrix} \begin{bmatrix} \mu_x \\ \mu_y \\ \mu_z \end{bmatrix} \quad (3.9)$$

Notably, some of the matrix elements are redundant; of the matrix components N which contribute to the field there are only 6 which are independent. The application of the FFT accelerated technique requires pre-computation of the FFT of S_x , S_y , S_z , N_{xx} , N_{xy} , N_{xz} , N_{yy} , N_{yz} and N_{zz} . At each computation step, the FFT of the padded μ_x , μ_y and μ_z arrays can then be computed, and the potential and field computed in Fourier space as:

$$\begin{aligned} \tilde{\Phi} &= \tilde{S}_x \tilde{\mu}_x + \tilde{S}_y \tilde{\mu}_y + \tilde{S}_z \tilde{\mu}_z \\ \tilde{F}_x &= \tilde{N}_{xx} \tilde{\mu}_x + \tilde{N}_{xy} \tilde{\mu}_y + \tilde{N}_{xz} \tilde{\mu}_z \\ \tilde{F}_y &= \tilde{N}_{xy} \tilde{\mu}_x + \tilde{N}_{yy} \tilde{\mu}_y + \tilde{N}_{yz} \tilde{\mu}_z \\ \tilde{F}_z &= \tilde{N}_{xz} \tilde{\mu}_x + \tilde{N}_{yz} \tilde{\mu}_y + \tilde{N}_{zz} \tilde{\mu}_z \end{aligned} \quad (3.10)$$

These quantities are then back-transformed into real space using the inverse FFT algorithm, in order to determine the real-space potential and field.

In a direct computation of the field, if a position in the grid does not hold a spin, we can simply neglect it from the computation completely. Because of the strict constraint on one of the functions in the discrete convolution to be a periodic summation, this is not possible using the FFT. Many different lattice types are mappable onto a regular grid to satisfy the constraint, but at the cost of large increases in the memory necessary; the full dipolar tensor must be computed, and any ‘ghost’ spins on the lattice have are set to zero in the array. The field is subsequently calculated at the positions on the grid where no particle sits, and this can be seen as an inefficiency of the method. Consider, as an example, a BCC cubic lattice. We can construct a cubic unit cell of $3 \times 3 \times 3$ spins which fits naturally onto a regular grid, setting to zero 18 spins at the centre of each face and edge. In storing a single unit cell of this lattice, $\frac{2}{3}$ of the storage space

describes non-existent spins, which requires an increase in memory in storing the spin arrays of 200% per unit cell when compared to the naive method. Similarly, a triangular lattice can be imposed on a regular grid, and in this case it requires a 100% increase in memory. While this is a downside for non-cubic lattices, it can sometimes be useful to compute the potential and field of an ensemble of sources accurately at some distance by adding ghost sources around the object.

3.3 Micromagnetics

With Micromagnetics being a continuum field theory, the fields must be re-discretised in order to solve the system of differential equations computationally. There are two main approaches to doing so; using either a finite-element or a finite difference approximation.

More care needs to be taken in setting up micromagnetic simulations than in atomistic simulations because of the approximations made in discretising the fields. While in atomistic spin dynamics we can identify each atom as the smallest relevant particle, in micromagnetics each cell contains thousands of atoms, and so the discretisation needs to be both small enough that the approximations in the derivation of the Energy functionals are good, and large enough that the computational cost does not approach that of an atomistic simulation. In general, a good rule of thumb is that the discretisation length must be smaller than the exchange length, as otherwise independent domains cannot form in simulation [88, 102]. In order to settle on an appropriate value, one should perform a convergence study on the quantity of interest, be that a static property such as the energy or a dynamic property of the system such as the reversal or depinning field [103].

3.3.1 Finite Difference Method

The finite difference method is by far the simpler of the two methods used to treat micromagnetics computationally. The magnetisation field of the magnetic material is discretised onto a structured mesh of cubic cells. In this regard, much of the infrastructure is the same as for the atomistic case - the same indexing schemes can be used, and the connectivity between cells defined in the same way as for neighbouring atoms.

For the micromagnetic energy formulations, approximate derivatives of the magnetisation are straightforwardly calculated in Cartesian coordinates using central difference approaches, such that for derivatives with respect to x :

$$\frac{\partial \mathbf{m}(\mathbf{r})}{\partial x} \approx \frac{\mathbf{m}(r + h) - \mathbf{m}(r - h)}{2h} \quad (3.11)$$

$$\frac{\partial^2 \mathbf{m}(\mathbf{r})}{\partial x^2} \approx \frac{\mathbf{m}(r + h) + \mathbf{m}(r - h) - 2\mathbf{m}(r)}{h^2} \quad (3.12)$$

When a cubic grid is used, with $n = (n_x, n_y, n_z)$ cells, and cell spacings of (a_x, a_y, a_z) derivatives of \mathbf{m} at a cell (i, j, k) with position in space $r_{i,j,k} = (ia_x, ja_y, ka_z)$ can be evaluated on the mesh as:

$$\frac{\partial \mathbf{m}(\mathbf{r}_{i,j,k})}{\partial x} = \frac{\mathbf{m}(\mathbf{r}_{i+1,j,k}) - \mathbf{m}(\mathbf{r}_{i-1,j,k})}{2a_x} \quad (3.13)$$

$$\frac{\partial^2 \mathbf{m}(\mathbf{r}_{i,j,k})}{\partial x^2} = \frac{\mathbf{m}(\mathbf{r}_{i+1,j,k}) + \mathbf{m}(\mathbf{r}_{i-1,j,k}) - 2\mathbf{m}(\mathbf{r}_{i,j,k})}{a_x^2} \quad (3.14)$$

From this, one needs to simply iterate over all of the cells on a mesh in order to compute the derivative of the magnetisation field, and by simple substitution of the derivatives into the micromagnetic expressions given in the previous chapter, computational methods for evaluating the energy and/or the effective field are straightforwardly built up. Higher order finite-difference schemes can be utilised for increased accuracy, at the expense of computational time [103].

In general, one can construct a matrix formulation for the finite-difference method in micromagnetics, such that:

$$\mathbf{H}_{\text{eff}} = \mathbf{A}\mathbf{M} \quad (3.15)$$

When most of the relevant energy terms are included, this leads to a sparse matrix, because the Exchange and Dzyaloshinskii-Moriya interactions are built up from first and second derivatives and so for a given cell only it and its nearest neighbours need to be considered in the finite-difference scheme, and the Anisotropy and . However, including the non-local demagnetising interaction leads to a dense matrix, so in practice the problem is not usually formulated in this way. Fidimag and OOMMF take the approach of computing the effective field of each interaction term independently, which is probably the most flexible approach. For performance reasons, however, it is advantageous to iterate through the magnetisation array as few times as possible, and the software MuMax3 [104] instead takes this approach, by combining the Exchange, Zeeman and Dzyaloshinskii-Moriya interactions into a single loop iteration, with the downside that the various energy fields are not separable.

In all micromagnetic calculations, the computation of the non-local demagnetising interaction is, like in atomistic simulations, the dominant part of the calculation time - a naive calculation of the demagnetising field is an $\mathcal{O}(n^2)$. Calculation of the demagnetising interaction is generally implemented using the same convolution approach as in the atomistic case, leading again to an $\mathcal{O}(n \log n)$ scaling. Here, rather than describing the interaction between dipoles, the matrix components describe the interactions between discrete cuboids of magnetic material which are uniformly magnetised. These formulae were derived by considering the relative interactions between faces on the cuboids by Newell, et. al. [102, 105], with an explicit form of these interactions given as:

$$N_{xx}(\mathbf{r}) = \frac{1}{4\pi a_x a_y a_z} [2F(x, y, z) - F(x + a_x, y, z) - F(x - a_x, y, z)] \quad (3.16)$$

$$N_{xy}(\mathbf{r}) = \frac{1}{4\pi a_x a_y a_z} [G(x, y, z) - G(x - a_x, y, z) - G(x, y + a_y, z) - G(x - a_x, y + a_y, z)] \quad (3.17)$$

with the auxiliary functions $F(x, y, z)$ and $G(x, y, z)$ defined as:

$$\begin{aligned} F(x, y, z) = & 4f(x, y, z) - 2f(x, y, z - a_z) \\ & - 2f(x, y, z + a_z) - 2f(x, y - a_y, z) \\ & + f(x, y - a_y, z - a_z) + f(x, y - a_y, z + a_z) \\ & - 2f(x, y + a_y, z) + f(x, y + a_y, z - a_z) \\ & + f(x, y + a_y, z + a_z) \end{aligned} \quad (3.18)$$

$$\begin{aligned} G(x, y, z) = & 2g(x, y, z) - g(x, y, z - a_z) \\ & - g(x, y, z + a_z) - 2g(x, y - a_y, z) \\ & + g(x, y - a_y, z - a_z) + g(x, y - a_y, z + a_z) \\ & - 2g(x + a_x, y, z) + g(x + a_x, y, z - a_z) \\ & + g(x + a_x, y, z + a_z) + 2g(x + a_x, y - a_y, z) \\ & - g(x + a_x, y - a_y, z - a_z) \\ & - g(x + a_x, y - a_y + y, z + a_z) \end{aligned} \quad (3.19) \quad (3.20)$$

where

$$\begin{aligned}
 f(x, y, z) = & \frac{yz^2 - x^2y}{2} \operatorname{arcsinh}\left(\frac{y}{\sqrt{x^2 + z^2}}\right) \\
 & + \frac{y^2z - x^2z}{2} \operatorname{arcsinh}\left(\frac{z}{\sqrt{x^2 + y^2}}\right) \\
 & - xyz \operatorname{arctan}\frac{yz}{x\sqrt{x^2 + y^2 + z^2}} \\
 & + \frac{2x^2 - y^2 - z^2\sqrt{x^2 + y^2 + z^2}}{6}
 \end{aligned} \tag{3.21}$$

$$\begin{aligned}
 g(x, y, z) = & (xyz) \operatorname{arcsinh}\frac{z}{x^2 + y^2} \\
 & + \frac{3yz^2 - y^3}{6} \operatorname{arcsinh}\frac{x}{y^2 + z^2} \\
 & + \frac{3xz^2 - x^3}{6} \operatorname{arcsinh}\frac{y}{x^2 + z^2} \\
 & - \frac{z^3}{6} \operatorname{arctan}\frac{xy}{z\sqrt{x^2 + y^2 + z^2}} \\
 & - \frac{zy^2}{2} \operatorname{arctan}\frac{xz}{y\sqrt{x^2 + y^2 + z^2}} \\
 & - \frac{zy^2}{2} \operatorname{arctan}\frac{yz}{x\sqrt{x^2 + y^2 + z^2}} \\
 & - \frac{xy\sqrt{x^2 + y^2 + z^2}}{3}
 \end{aligned} \tag{3.22}$$

and with the expressions for other tensor components (for e.g. N_{yy}, N_{xz}) computed through cyclic permutation (e.g. $(x + a_x, y, z) \rightarrow (y + a_y, z, x)$).

For intermediate to long distances, calculation of the tensor with the exact formulae leads to large loss of precision errors because of catastrophic cancellation. It is therefore usual to use the asymptotic formulae to compute the demagnetising tensor [106] at distances greater than around 20 cell lengths. This has increasing relevance with graphical processing units (GPUs) starting to see widespread adoption in the simulation community, as the consumer-priced GPUs adopted in many academic clusters offer much lower performance for double-precision computations (usually around 20% of peak). With this in mind, it has been noted that it is important to calculate at least the demagnetising tensor in double-precision [107], and this is indeed the approach taken by the GPU-accelerated micromagnetic software MuMax3 [104], which otherwise uses a single-precision floating point representation.

3.3.2 Finite Element Method

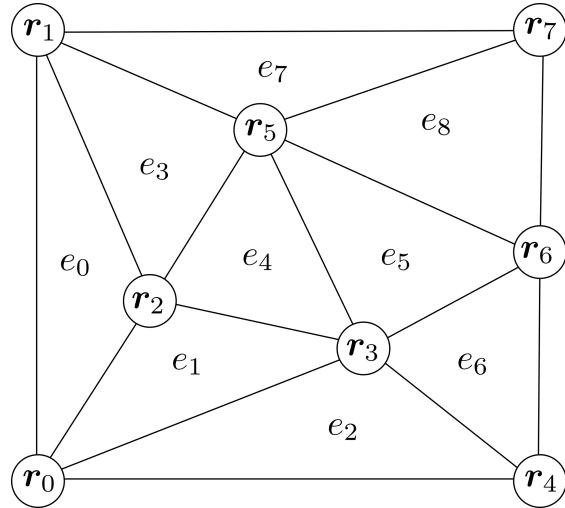


FIGURE 3.1: Figure showing a 2-D example triangular finite element mesh of a rectangle with $N = 7$ elements, constructed from 8 nodes.

In the finite element method (FEM), an unstructured mesh is used to discretise the domain, and elements can be of any shape, though in practice triangles or quadrilaterals (in 2D) and tetrahedra (in 3D) are often chosen. By subdividing the solution domain into a finite number N of sub-domains known as ‘elements’ (Figure 3.1), constructed from a set of nodes, the aim of the finite-element method is to construct a solution over the whole domain by finding values of the solution at the nodes. The number of ‘free’ nodes which gives the total number of unknowns in the linear system is fewer than the total number of nodes in the system, because some are constrained by the choice of boundary conditions - for Dirichlet boundary conditions all edge nodes on a finite element mesh have specified values.

3.3.2.1 Generalised Weighted Average Method and the Weak Form

To solve differential equations with appropriately specified boundary conditions using the finite-element method, the equation must be recast into the so called ‘weak’ form [108], which holds across the finite elements that make up the total domain. Taking as an example the Poisson equation in two dimensions:

$$\int_{\Omega} \nabla^2 u(x, y) \, dV = \int_{\Omega} f(x, y) \, dV \quad (3.23)$$

on the domain $\Omega \in \mathbf{R}^2$, where $u_{\partial\Omega} = g(x, y)$. To solve this computationally, we want to find a solution such that the residual R :

$$R = \int_{\Omega} \nabla^2 u(x, y) - f(x, y) \, dV = 0 \quad (3.24)$$

up to some tolerance, where $u = \sum_i u_i \phi_i(x, y)$ is an approximation of u formed of an expansion in basis functions, which is known as a trial function. The space of basis functions is known as the trial space, and the most general set of functions which can be used lie in the Sobolev space, the space of functions which have a bounded L2 norm over the domain (sometimes known as L2-integrable). Finding the minima of this equation by varying the coefficients a_i of the basis functions therefore leads to a good approximation of the true solution. One choice to do so is to minimise R everywhere in space by finding a weighted average of R across the total domain by weighting with a ‘test’ function v as:

$$\int_{\Omega} v R \, dV = \int_{\Omega} v(\nabla^2 u) \, dV - \int_{\Omega} v f \, dV = 0 \quad (3.25)$$

This test function is similarly constructed from a set of basis functions as $v = \sum_i v_i \psi_i(x, y)$, and a constraint is imposed that it vanishes on the boundary. Often, the set of basis functions for both the trial and test function are chosen to be of the same form ($\phi_i = \psi_i$), and if this is the case then the method known as a Galerkin method. We follow this convention here.

Simplifying the right hand side of Eq. 3.25 by integrating by parts:

$$\int_{\Omega} (\nabla^2 u) v \, dV = - \int_{\Omega} \nabla u \cdot \nabla v \, dV + \int_{\Omega} f v \, dV + \int_{\partial\Omega} (\nabla u) \cdot \hat{\mathbf{n}} \, dS \quad (3.26)$$

Note that the boundary terms vanish because of the constraint that $v = 0$ on the boundary. Then, the weak form is given as:

$$- \int_{\Omega} \nabla u \cdot \nabla v \, dV = \int_{\Omega} f v \, dV \quad (3.27)$$

Here, the u term only needs to be differentiable once; the requirement of differentiability on u is weakened from the initial form of the differential equation, and so this is the reason it is called the weak form [109].

Imposing the requirement that the value inside each element is uniquely specified by the values on the nodes making up that element leads to a globally continuous solution. For a particular node i with basis function ϕ_i , the value is 1 on the node,

and 0 on all other nodes, with the basis function spanning only the elements which contain that node - the basis function is compact over those elements (Figure 3.2). The value of the function on the node is therefore $u(x_i, y_i) = u_i \phi_i(x_i, y_i) = c_i$. Because the basis functions have value only within the elements on which the nodes they sit on are part of, the inner product on two element basis functions is only non-zero if they are part of the same element.

Using this knowledge, the integral over the whole domain is broken up - each side of the equation in the weak form can be transformed into a form which depends on integrals over the basis functions for each node i and the volume they span which we label e_i :

$$-\int_{\Omega} \nabla u \cdot \nabla v \, dV \approx \sum_i u_i v_i \int_{e_i} (\nabla \phi_i) (\nabla \phi_j) \, dV \quad (3.28)$$

$$= \sum_i u_i v_i A_{ij} \quad (3.29)$$

where the matrix A (also alternatively labelled K) is often referred to as the stiffness matrix. For the right hand side:

$$\int_{\Omega} f v \, dV = \sum_i f_i v_i \int_{e_i} \phi_i \, dV \quad (3.30)$$

$$= \sum_i f_i v_i b_i \quad (3.31)$$

where \mathbf{b} is sometimes known as the load vector or Because v_i appears on both sides, it can be eliminated, giving a matrix equation:

$$\sum_i u_i A_i = \sum_i f_i b_i \quad (3.32)$$

in which u_i are unknown, and the other terms can be directly computed. The Dirichlet boundary condition is incorporated through modification of the linear system. The integrals over basis functions in A and b are often computed in the finite-element method by using approximate quadrature formulae.

In order to solve the constructed linear systems, iterative methods such as the generalised minimal residual method (GMRES) are normally utilised for systems with more than a few hundred unknowns; for smaller systems Gaussian elimination can be used.

For most finite element solver codes, each element is mapped onto a reference element in a simplified coordinate space. By doing so, the basis functions for a

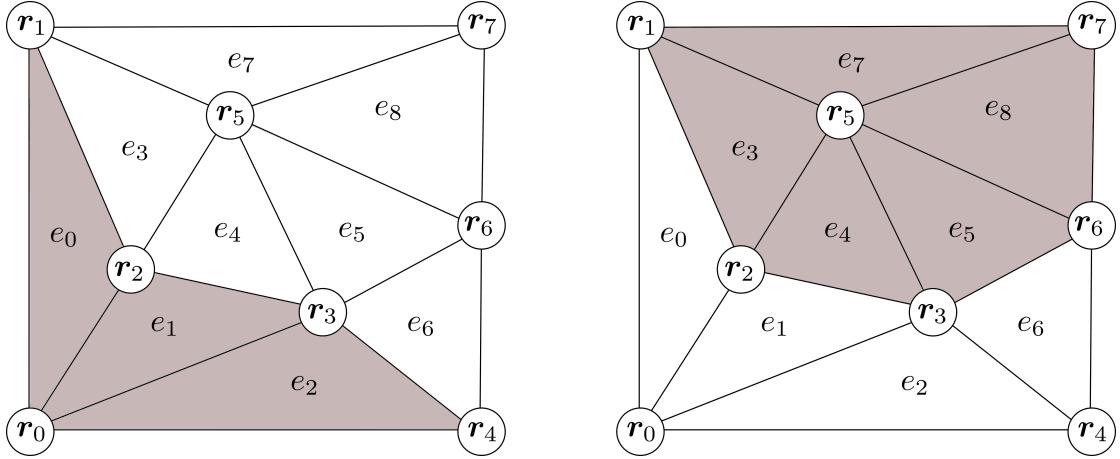


FIGURE 3.2: Figure showing the span of the basis function on node 0 (left) and node 5 (right).

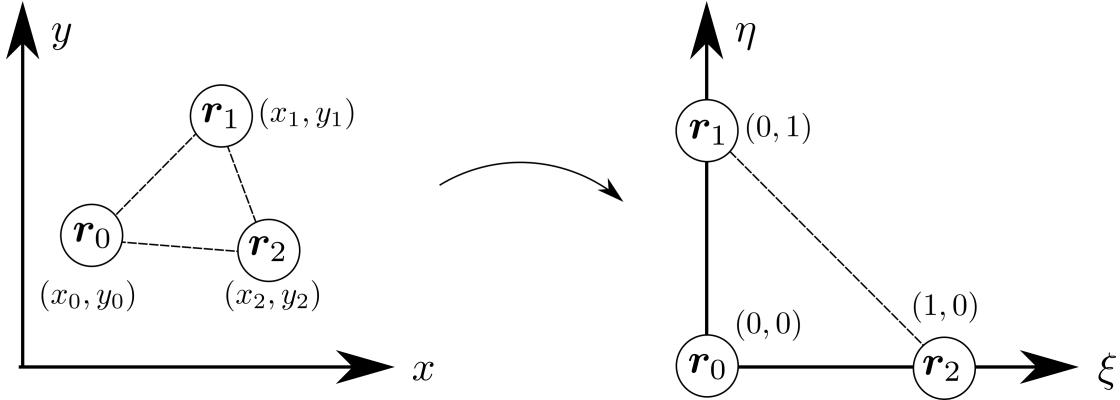


FIGURE 3.3: Figure showing a choice of mapping of a triangular element onto a reference element.

particular element type are constructed only once which simplifies the implementation of quadrature functions used to compute the integrals, at the expense of coordinate transforms. Taking for example linear basis functions on a triangular mesh, in the reference element shown in Figure 3.3, the form of the basis functions can be given as:

$$\begin{aligned}\phi_0(\eta, \xi) &= 1 - \eta - \xi \\ \phi_1(\eta, \xi) &= \eta \\ \phi_2(\eta, \xi) &= \xi\end{aligned}\tag{3.33}$$

Derivatives of the a function constructed from basis functions in the ordinary coordinate space, needed in the linear system, can straightforwardly be found by

evaluating the Jacobian, which can be precomputed and stored:

$$\begin{aligned}\frac{\partial f}{\partial x} &= \frac{\partial f}{\partial \eta} \frac{\partial \eta}{\partial x} + \frac{\partial f}{\partial \xi} \frac{\partial \xi}{\partial x} \\ \frac{\partial f}{\partial y} &= \frac{\partial f}{\partial \eta} \frac{\partial \eta}{\partial y} + \frac{\partial f}{\partial \xi} \frac{\partial \xi}{\partial y}\end{aligned}\tag{3.34}$$

3.3.2.2 Application to Micromagnetics

The finite-element approach to micromagnetics was introduced by Fredkin and Koehler in the late 1980s [110]. The most common use of the finite element method in micromagnetics does not so much use the finite element method itself (i.e. by constructing and solving a linear system subject to boundary conditions), but makes use of the finite element approximation to compute the effective field from the magnetisation, represented on a finite element mesh, in a convenient manner. The effective field is then used either to compute the right hand side of the Landau-Lifshitz-Gilbert equation on each node, or in one of various energy minimisation algorithms to find metastable states.

In the original approach by Fredkin and Koehler, the magnetisation was assumed to be constant throughout each finite element. Most implementations of the finite element equations for micromagnetism, including Finmag, the software used by the present author for the work in this thesis, instead follow a formulation more similar to that of Chen [111] whereby the set of 1st order linear polynomials were used as the basis set for each component, allowing the magnetisation to vary within cells. In Finmag, the magnetisation is represented by 1st order Lagrange basis functions, and coefficients such M_s and A are represented by 1st order discontinuous Lagrange basis functions in order that the boundaries between materials with different parameters can be resolved appropriately.

In order to compute the non-local demagnetising field, several approaches can be taken. In common to all techniques, the magnetostatic potential is generally found, and the numerical gradient is then taken. Sticking strictly to finite elements, however, it is difficult to enforce the boundary condition for the magnetostatic potential, which approaches zero at infinity, and early studies made use of a large, coarse "air-box" with no magnetic material, surrounding a magnetic region, though errors from this approach are large unless the air-box region is sufficiently large, which clearly increases the calculation time. Most publicly available finite-element software [112–116] have instead therefore followed the method of Fredkin and Koehler, [117] by coupling the boundary element method (BEM) and finite

element method. Rather than starting from the explicit solution, the magnetic scalar potential ϕ can be more simply calculated as the solution of the differential equation [118]:

$$\nabla^2 \phi(\mathbf{r}) = \begin{cases} 4\pi \nabla \cdot \mathbf{M}(\mathbf{r}) & \mathbf{r} \in \Omega \\ 0 & \mathbf{r} \notin \Omega \end{cases} \quad (3.35)$$

$$(3.36)$$

with boundary conditions on $\partial\Omega$

$$[\phi(\mathbf{r})] = 0 \quad (3.37)$$

$$\left[\frac{\partial \phi}{\partial \hat{n}}(\mathbf{r}) \right] = -4\pi \hat{n} \cdot \mathbf{M}(\mathbf{r}) \quad (3.38)$$

where \hat{n} is the unit vector normal to the surface and here, the scalar potential is split into two components $\phi = \phi_1 + \phi_2$ where:

$$\nabla^2 \phi_1(\mathbf{r}) = 4\pi \nabla \cdot \mathbf{M}(\mathbf{r}) \quad x \in \Omega \quad (3.39)$$

with boundary condition on $\partial\Omega$

$$\left[\frac{\partial \phi_1(\mathbf{r})}{\partial \hat{n}} \right] = 4\pi \hat{n} \cdot \mathbf{M}(\mathbf{r}) \quad (3.40)$$

and

$$\nabla^2 \phi_2(\mathbf{r}) = 0 \quad \mathbf{r} \in \mathbb{R}^3 \quad (3.41)$$

$$(3.42)$$

with boundary condition on $\partial\Omega$:

$$\left[\frac{\partial \phi_2(\mathbf{r})}{\partial \hat{n}} \right] = 0 \quad (3.43)$$

with the additional constraint that

$$\lim_{|\mathbf{r}| \rightarrow \infty} u_2(\mathbf{r}) = 0 \quad (3.44)$$

Then, ϕ_1 is directly calculated by solving the above Poisson-like equation using the finite-element method in the interior part of the domain. The boundary values

of ϕ_2 are then calculated using the boundary element method as:

$$\phi_2(\mathbf{r}) = \frac{1}{4\pi} \int_{\partial\Omega} \phi_1(\mathbf{r}') \frac{\partial}{\partial \hat{n}(\mathbf{r}')} \frac{1}{|\mathbf{r} - \mathbf{r}'|} dS \quad (3.45)$$

These boundary values are then used as the boundary condition while solving for ϕ_2 , again by the finite element method in the interior domain. The numerical gradient is then computed from the sum of the two computed potentials to find each component of the field.

Alternative techniques for finding the magnetostatic potential from a magnetisation distribution are the shell-transformation technique, in which an exterior space is meshed but the position of the outer shell scaled to match some far-distance on which the boundary condition can be set [119, 120], and both tree code and fast-multipole method codes, which have been implemented to solve for the magnetostatic potential by several authors [121–123]. The use of these techniques have in some ways been motivated by the difficulty of using the FEM/BEM approach on distributed systems due to the dense BEM matrices which require large communication bandwidth for iterative solution.

3.4 Evolution of the magnetisation

With the energy formulae, one can calculate the energy of a given magnetic state. However, this is not usually that interesting - one wants to find how this might evolve in space, or want to find the state which minimises the energy. We describe several of these techniques here. We show a general procedure one can use to update the magnetisation in Figure 3.4.

3.4.1 Time Integration

In both atomistic and micromagnetics (finite-element and finite-difference), the techniques used to solve the Landau-Lifshitz-Gilbert equation of motion are very similar. The time-dependent LLG equation must be discretised into a finite number of steps, and so the finite-difference approach can be used in the time domain to a first approximation:

$$\mathbf{m}(t + dt) \approx dt \frac{\partial \mathbf{m}(t)}{\partial t} \quad (3.46)$$

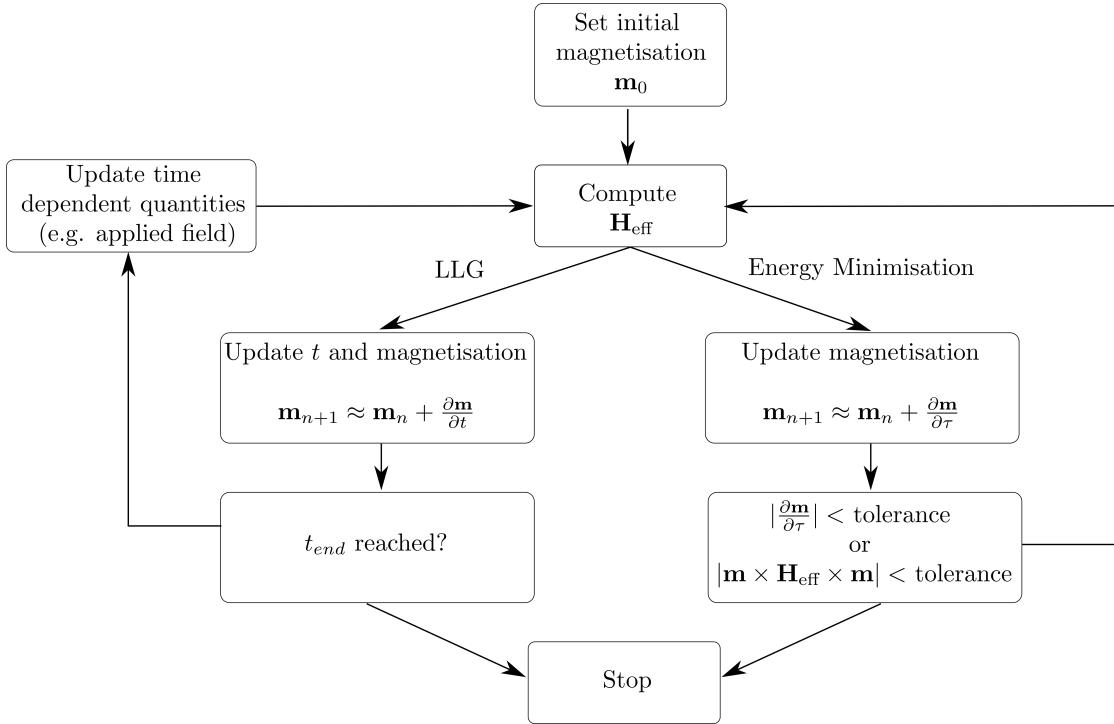


FIGURE 3.4: Algorithm procedure for both micromagnetic and atomistic simulations.

To facilitate this kind of updating formulae, the LLG is usually used in the explicit Landau-Lifshitz form:

$$\frac{\partial \mathbf{m}}{\partial t} = \frac{\gamma_0^*}{1 + \alpha^2} \mathbf{m} \times \mathbf{H}_{\text{eff}} - \frac{\gamma_0 \alpha}{1 + \alpha^2} \mathbf{m} \times \mathbf{m} \times \mathbf{H}_{\text{eff}} \quad (3.47)$$

The LLG equation is well known for being stiff; that is, very small time steps are required in order to maintain numerical stability [124, 125]. For this reason, a simple updating formula like that above is not practical for simulations. Runge-Kutta methods or implicit solvers such as the are ordinarily used to solve the LLG equation in the absence of thermal fields. In the presence of thermal fields, it is usual to use predictor-corrector type Heun's method [58].

One issue that arises in the numerical solution of the LLG in both atomistic and micromagnetic modelling is enforcing the constraint that the spin length remains constant. Some codes [104, 126] use an explicit re-normalisation after each time step. One issue with this when using an adaptive time stepping methods is that the spin length may drift within a time step, leading to inaccurate calculations of the effective field. In Fidimag and Finmag, a drift term is added to the calculation

of the LLG, which acts only in the direction parallel to the magnetic vectors:

$$\frac{\partial \mathbf{m}}{\partial t} \rightarrow \frac{\partial \mathbf{m}}{\partial t} + c \frac{\partial \mathbf{m}}{\partial t} (1 - \mathbf{m})^2 \mathbf{m} \quad (3.48)$$

where c is an empirically determined constant.

Alternative geometric integration schemes such as the implicit midpoint method [127] exist, which preserve the length of the magnetisation as a consequence of their construction. These have not been widely applied in the field, in part due to the lack of adaptive time-stepping schemes which make the computational cost prohibitive to their application. Recently, however, several such adaptations of the method have been developed [127–129].

In order to find realistic, metastable states, it is often useful to relax a system under Landau-Lifshitz-Gilbert dynamics, usually with a value of the Gilbert damping α much greater than realistic values for a given material in order to suppress precessional dynamics.

3.4.2 Energetics

As an alternative to studying the dynamics or using relaxation to find metastable states, it is sometimes faster to use direct energy minimisation. Two general families for doing so are the steepest descent method and the conjugate gradient method. In this thesis we use the steepest descent method popularised by Exl et. al. [130] in which the evolution of the magnetisation is described by the iterative scheme:

$$\mathbf{m}^{n+1} = \mathbf{m}^n - \tau \frac{\mathbf{m}_n + \mathbf{m}_{n+1}}{2} \times \mathbf{H}_{\text{eff}}(\mathbf{m}_n) \quad (3.49)$$

where τ is a fictitious time step.

It is often interesting to find the minimum energy paths between magnetic states, in order to determine whether states found through direct energy minimisation are physically realisable, or if they are likely to be inaccessible - states in which a large energy barrier for formation exists may be rare in practice. The Nudged Elastic Band Method (NEBM) is one such technique for doing so. In the NEBM, a set of images $\mathbf{Y}_i, i \in 0, 1, \dots, n$ of a system are created - the so called ‘band’ of the method. Each image in the band shares identical magnetic properties (for e.g. spin magnitude or saturation magnetisation magnitude, exchange interaction strength) with the other images. However, the magnetisation of each band is allowed to vary

between each image and for a given spin p in band i the magnetisation is labelled $\mathbf{m}_p^{(i)}$. The images \mathbf{Y}_0 and \mathbf{Y}_n are fixed, and are set to be the two states for which a transition is to be determined. These two images must be energy minima down to a high numerical precision, in order for the algorithm to converge, found either through LLG dynamics or direct energy minimisation. The other states must be set such that they interpolate between the two equilibrium states. This can be performed in several ways, such as:

- Using a linear/quadratic/cubic interpolation of the angle between each spin of the equilibrium states for each image.
- Creating images which interpolate between the two states in such a way as to imitate a particular transition.

The NEBM itself is described by a dynamical equation:

$$\frac{\partial \mathbf{Y}_i}{\partial \tau} = -\mathbf{Y}_i \times \mathbf{Y}_i \times \tilde{\mathbf{G}}_i + c(1 - \mathbf{Y}_i^2)\mathbf{Y}_i \sqrt{\left(\frac{\partial \mathbf{Y}_i}{\partial t}\right)^2} \quad (3.50)$$

which describes the evolution of the band under a fictitious time τ , with the right hand term acting to keep the spin/magnetisation length normalised (due to drift from numerical errors).

The vector \mathbf{G}_i describes the by-node ‘force’ acting on each image in the band, and is comprised of two components:

$$\mathbf{G}_i = \mathbf{H}_{\text{eff}} + k\mathbf{F}_i \quad (3.51)$$

where \mathbf{F}_i is defined as:

$$\mathbf{F}_i = \|\mathbf{Y}_{i+1} - \mathbf{Y}_i\| - \|\mathbf{Y}_i - \mathbf{Y}_{i-1}\| \quad (3.52)$$

The first component, which we know to be the effective field, simply acts on each image as to move it towards a local energy minima (remembering that the energy gradient is exactly the effective field). The right hand term, which acts as though each image is one of a chain of particles on a chain, keeps the images distributed along the band, to avoid them from clustering together in the energy landscape around local energy minima or the equilibrium states. The spring constant k is a parameter of the algorithm which controls the strength of this force. The total force \mathbf{G}_i must be projected onto the tangent plane of the magnetisation; this avoids

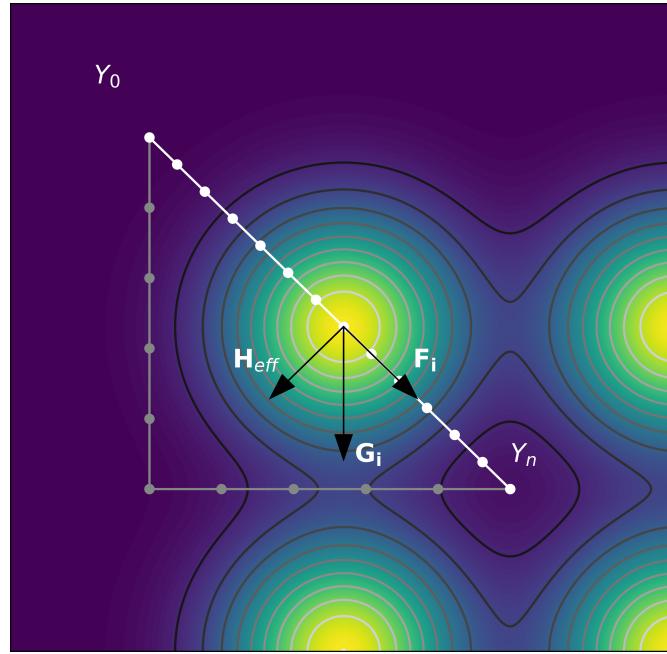


FIGURE 3.5: Example Energy Landscape - Here, the white band is the initial interpolated band between two states, which traverses a maxima in the landscape, while the grey band represents a minimal energy path which the band has evolved towards.

the force from acting in such a way as to break the constraints of fixed spin length, and can be carried out as:

$$\tilde{\mathbf{G}}_p^{(i)} = \mathbf{G}_p^{(i)} - (\mathbf{G}_p^{(i)} \cdot \mathbf{m}_p) \mathbf{m}_p \quad (3.53)$$

We show the direction of the forces in Figure 3.5. The norm $\|\mathbf{Y}_{i+1} - \mathbf{Y}_i\|$ defines a ‘distance’ in the phase space of the bands, determining how far apart spins are in each image, which must be considered carefully. We follow the formulation of the NEBM method as applied by Bessarab, et. al. [131, 132], and use the geodesic distance, defined in Cartesian form as:

$$\|\mathbf{Y}_i - \mathbf{Y}_j\| = \sqrt{\sum_p \arctan \left(\frac{\mathbf{m}_p^{(i)} \cdot \mathbf{m}_p^{(j)}}{|\mathbf{m}_p^{(i)} \times \mathbf{m}_p^{(j)}|} \right)} \quad (3.54)$$

though we note that other definitions of the distance have been used in previous works.

Initially, the NEBM needs to be set up with images distributed roughly evenly along the energy path; this can be achieved in a number of ways but in general, some sort of interpolation of the spin vectors between the initial and final states is used [131–135]. The algorithm can be sensitive to the choice of the initial states of

the band, with convergence strongly affected by the appropriateness of the initial guess. The NEBM algorithm is considered to have converged once changes of all of the varying images in the band are below a certain tolerance. Care needs to be taken that a large enough set of images is created so that the method converges to a minimum energy path, that is, the path over which the energy barrier is minimal; with too few images, the path is not minimal. This can be checked by adding additional images and rerunning the procedure - if sufficiently many images are used then the barrier should not change as the number of images n increases.

3.5 Approach Choice: Micromagnetics or Atomistic?

The type of problem under study necessitates different treatment. The micromagnetic approach can clearly not describe ferrimagnetism or antiferromagnetism, because the assumption that the field varies slowly required in the derivation of the continuum form is broken in these cases. Extensions to micromagnetics in order to treat these types of systems have seen some success [136], but the techniques are not widely adopted in the field, in part due to a lack of open source implementations of these methods in software. In some materials, exchange between next-nearest neighbours can also be important to the energetics of the system, and again micromagnetics is not suitable for such materials. Nonetheless, it is worth pointing out that computational micromagnetics has in general been a major success story in the last 30 years, with accurate predictions made for many real-world systems.

Even within micromagnetic problems, the computational approach needed to solve certain problems must be carefully considered. The two most commonly used micromagnetic software packages, OOMMF and MuMax are both finite difference based, and it is well known that for curved boundaries, the various energy terms have large errors [104, 126]. While there are methods to mitigate this, such as through the use of an immersed boundary techniques [137], these are not widely implemented or used. For these types of problems, a finite element discretisation is instead usually chosen. The finite-element approach is, however, much less widely used than finite differences in general, due in part to the difficulty of implementation when compared to the finite-difference method. In addition, from a user perspective, most of the freely available and formerly popular finite element software packages for computing magnetisation dynamics are no longer maintained, such

as Magpar [112] and Nmag [114], with the remaining codes such as Magnum.Fe [116] and FastMag [138] available through commercial licensing, and the Finmag package used in this thesis, developed at the University of Southampton and the European XFEL, only recently made public. Other finite-element micromagnetic codes such as MERRILL [139] and Micromag [120, 140, 141] are targeted at the geosciences community, and therefore focus on the energy of and energy barriers between magnetic states, rather than Landau-Lifshitz-Gilbert dynamics.

The modelling work carried out in Chapters 5 and 6 of this thesis was conducted using the finite-element micromagnetic software package Finmag [115], which was developed between 2011 and 2018 Hans Fangohr's group at the University of Southampton, and to which the present author contributed mainly meshing and geometry creation functionality, and Singularity images for running on HPC clusters. Chapters 4 and 7 of this thesis utilised Fidimag [95, 96], a micromagnetic and atomistic modelling software which started in 2012 and which is aimed at being flexible for new development and easy to use, and which is still under active development. Only Fidimag's atomistic modelling component was used in this thesis. The authors main contributions to Fidimag include the refactoring and porting of the software to the C++ language, improvements to the build system to allow parallel builds, the fast-multipole method routines described in Chapter 4, and the addition of higher order anisotropy calculations. A small amount of use was made of the Object Oriented Micromagnetic Framework (OOMMF) developed by Mike Donahue and Donald Porter at NIST [126], likely the most widely used and robust software package in micromagnetics and which pioneered many of the approaches taken by others; this was mostly for verification of the finite-element results such as in Chapter 6.

It is notable that three of the aforementioned finite-element micromagnetic packages (Magnum.fe, Finmag and Micromag) were built around the FEniCS software package [142–146], which allows the user to simply specify the weak form of a partial differential equation, which is then translated into a computational representation of the mathematics by the UFL package [147]. In this framework, the finite element family can be chosen from a wide variety of element types, and the order of the basis functions used can be adjusted straightforwardly [148, 149]. In general, very little code needs to be written by the user to solve a problem with FEniCS, which makes it straightforward to adopt. Internally, FEniCS makes use of high performance sparse matrix solvers from the PETSc library [150], which allows it to scale to many thousands of processors without input from the user,

with matrix assembly handled by the package UFC [147, 151]. Despite the flexibility and ease of use at a high level, one downside of the FEniCS approach is that extending the code to work in some cases (such as in implementing the Boundary Element Method, which is not available in FEniCS, for the demagnetising field) is difficult and requires the use of low level APIs, which are subject to change without documentation. With Finmag, for example, most development time over the last four or so years was in trying to maintain compatibility with new FEniCS versions rather than in improving the algorithms or implementing new features, and a decision was taken on the grounds of time to fix at an older version of FEniCS to reduce the maintenance burdens. Nonetheless, it was recently argued [152] that high-level packages like this are serving to make finite-element methods accessible to a wider audience in computational science, and so it is likely that in the future the finite-element approach will see further adoption in the micromagnetic community. In addition future improvements to FEniCs and other packages such as BemPP [153] which couples to FEniCs for BEM routines are likely to reduce the amount of work needed to implement a comprehensive micromagnetic solver using the finite-element method.

Chapter 4

Hierarchical Methods for Atomistic Long Range Field Calculations

In both micromagnetics and in spin dynamics simulations, the dominant part in computational time is the calculation of the non-local dipolar (or demagnetising) field. The broadest interest in accelerating long-range particle-particle interactions has, in the past, been in two fields; the astrophysical community where simulations have used the Barnes-Hut method (sometimes known as a ‘treecode’ method) to simulate interacting gravitational bodies, [154], and in chemical molecular dynamics where the Fast Multipole Method method is used to calculate Coulombic repulsion between charged ions. Both of these algorithms can be applied to calculate fields or potential, and are more general than these two domains, but have not seen widespread adoption. Both of the techniques rely on approximations in order to accelerate the calculation, but how justifiable these approximations are depends on the problem domain.

Implementations of the Fast Multipole Method (FMM) for the Poisson equation usually utilise a spherical harmonic expansion of the Green’s function, which leads to an irreducible representation. Doing this in a computationally efficient way is non-trivial because calculating the spherical harmonic functions necessary for such an expansion effectively requires making use of recursion relations in order to avoid numerical issues, due to catastrophic cancellation when numbers of different orders of magnitude are summed. Cartesian Taylor expansions also form a straightforward basis in which to expand the potential, though the multipole expansions are reducible and so are less computationally efficient than spherical harmonic expansions for the high accuracy regime. The expansion of the kernel

used in different software packages varies widely across different areas of study; in astrophysics a Cartesian basis is primarily used, especially in the Barnes-Hut (BH) method [155], but in chemical molecular dynamics studies it is more usual to make use of the Spherical Harmonic expansion . Similarly, the method utilised varies enormously, and application specific performance optimisations can be made for, which preclude against code-reuse when attempts are made to apply the methods to other problems.

In the past, an approach based on template metaprogramming in C++ was utilised by Visscher and Apalkov [67] to provide efficient recursive implementations of multipolar Cartesian operator functions for micromagnetic cell sources. In this work, the surface charges through each of the six faces of a micromagnetic cell were used as point monopoles, from which the dipolar field could be calculated, allowing the use of standard Cartesian fast multipole operators. In this paper, however, no comparison of performance with the standard method of computing the dipolar field was performed. A subsequent paper followed a similar template metaprogramming approach for applying the FMM to point dipoles [66], and in this study, a performance increase was shown over the FFT convolution method in some cases. In addition, some recent studies have began to utilise treecodes or the FMM to accelerate the demagnetising field computation in finite-element micromagnetic calculations, and shown either good performance improvements or worse performance but significant memory reductions over the more typically used Boundary Element Method [107, 122, 156]. A templating approach was also used by Wang. et. al. in order to implement Cartesian operators for the FMM as applied to the Boundary Element Method [156]. This templating approach, however, is not straightforwardly generalisable to languages which do not support code generation at compile time, such as C and Fortran, and it makes it difficult to apply broad optimisations across the templated code. In most FMM and BH implementations based on the Cartesian expansion, hand-written code is therefore used to implement the various FMM operators; in Dehnen’s FalcON code this goes as far as hand implementing vectorisation of the various operators [157]. This is difficult, because beyond several expansion orders it is tedious to ensure code correctness, and the number of terms in the multipole expansion grows with n as $n(n + 1)/2$. Implementing operators by hand in this way also makes it difficult to enable code reuse; for example, implementing both the multipolar BH and FMM in the same code base requires a large amount of duplication of code. Generalising to provide efficient operators for point sources of different orders (i.e. point monopoles, point dipoles, point quadrupoles) is also important for

adoption of the method. It is important to note that for many problems which are computationally intractable, reduced order models of systems can be constructed using multipoles, by treating objects with complex internal structures as points with a multipole expansion up to some given order, and these sources can often have non-negligible quadrupole terms [158].

Symbolic code generation is a technique which has, in recent years, been applied to the generation of functions for the computational solution of ordinary and partial differential equations. The FEniCs Form Compiler (FFC) library [159] constructs functions for the evaluation of variational forms for assembling finite-element matrices, and is used as part of both the FEniCs and Firedrake projects [160, 161]. The OpenSBLI [162–164] project generates finite-difference stencils in the language of the high-performance OPS library [165] from symbolic representations of differential equations, while the Devito [166] project achieves similar goals using symbolic code generation functionality in the SymPy library to generate efficient finite-difference kernels written in C. In the context of the fast multipole method, code generation has previously been utilised by Coles and Masella in order to provide an implementation of the Cartesian basis Fast Multipole Method in the closed source PolarisMD code, for the calculation of the electric potential and field from polarisable atoms in molecular dynamics, [167] with this work then being extended to support the use of more efficient operators through detracing techniques introduced by Applequist, which reduce multipole moments in one form to an irreducible form which cannot be then stored any more efficiently [157, 168–170].

In this work, we describe the implementation and details of an open-source code generation library, fmmgen, [68] which produces a set of operators for the Cartesian BH and FMM methods, and provides OpenMP parallelised example implementations of the methods, for the solution of the dipolar field in atomistic simulations (i.e. point dipole sources). We draw attention to how optimisations and simplifications can be enabled at different stages in the code generation to improve performance, and comment on the effectiveness of optimisation strategies.

A preprint of this work was made available on arXiV (2020) [171], and the data from this study and scripts to reproduce the figures are also available [172]. The present author and Prof. Hans Fangohr conceived of this study, and the present author wrote the fmmgen library, performed the simulations and interpreted the results and wrote the manuscript.

4.1 Mathematical Basis

We begin by showing the mathematical details necessary to construct the FMM and Multipolar BH methods to compute the potential and field the Laplace Equation in a Cartesian basis for source points of arbitrary order (i.e. Monopoles, Dipoles, Quadrupoles, ..., 2^n -poles. We denote the minimum ‘order’ of point sources in a system source as s , such that a monopole has order $s = 0$, a dipole $s = 1$, etc.

We here use the mathematical notation of monomials, which is widely used in the Fast Multipole literature. Here:

$$\begin{aligned}\mathbf{n} &= (n_x, n_y, n_z) \\ \mathbf{n} + \mathbf{m} &= (n_x + m_x, n_y + m_y, n_z + m_z) \\ \mathbf{n}! &= n_x! n_y! n_z! \\ \mathbf{r}^{\mathbf{n}} &= x^{n_x} y^{n_y} z^{n_z} \\ |\mathbf{n}| &= n_x + n_y + n_z \\ \binom{\mathbf{n}}{\mathbf{k}} &= \binom{n_x}{k_x} \binom{n_y}{k_y} \binom{n_z}{k_z}\end{aligned}$$

Consider the expansion of the Coulomb Potential from two well-separated cells A and B , with centres \mathbf{z}_a and \mathbf{z}_b , which containing points \mathbf{x}_a and \mathbf{x}_b respectively. We define vectors $\mathbf{r}_a = \mathbf{x}_a - \mathbf{z}_A$ and $\mathbf{r}_b = \mathbf{x}_b - \mathbf{z}_B$. When a charge q_a is located at \mathbf{x}_a , the potential at \mathbf{x}_b can be evaluated as:

$$\phi(\mathbf{r}) = \frac{q_a}{|\mathbf{x}_b - \mathbf{x}_a|} \quad (4.1)$$

Taylor expanding this around the point x_a and truncating at order p gives an approximate function for the evaluation of the potential:

$$\phi(\mathbf{x}_b - \mathbf{x}_a) \approx q_a \sum_{|\mathbf{n}|=0}^p \frac{(-1)^{\mathbf{n}}}{\mathbf{n}!} (\mathbf{x}_a - \mathbf{z}_a)^{\mathbf{n}} \nabla^{\mathbf{n}} \phi(\mathbf{x}_b - \mathbf{z}_a) \quad (4.2)$$

By grouping terms, a multipole term defined around the centre \mathbf{z}_A can be written:¹

$$\mathcal{M}_{\mathbf{n}}(\mathbf{z}_A) = \frac{(-1)^{|\mathbf{n}|}}{\mathbf{n}!} q_a (\mathbf{x}_a - \mathbf{z}_a)^{\mathbf{n}} \quad (4.3)$$

¹This definition varies between fields and authors. Notably, the factor of $(-1)^{\mathbf{n}}/\mathbf{n}!$ is often absorbed into the local expansion definition. It is also worth noting that the definition of the dipole and quadrupole moments can vary; for e.g. in Chemistry the dipole moment vector for

Because we have translated the particle's moment into a multipole expansion, this is known as the 'Particle-to-Multipole' or P2M operator. For a given \mathcal{M}_n term centred at \mathbf{z}_a , the shifted multipole expansion at a centre \mathbf{z}'_a can be derived through the substitution of $(\mathbf{x}_a - \mathbf{z}_a) = ((\mathbf{x}_a - \mathbf{z}'_a) + (\mathbf{z}'_a - \mathbf{z}_a))$, expanding out in powers and substituting multipole terms where recognised.

$$\mathcal{M}_n(\mathbf{z}'_a) = \sum_{|\mathbf{k}|=0}^{p-|n|} \frac{(\mathbf{z}_a - \mathbf{z}'_a)^\mathbf{k}}{\mathbf{k}!} \mathcal{M}_{n-\mathbf{k}}(\mathbf{z}_a) \quad (4.4)$$

This is usually referred to in the literature as the 'Multipole-to-multipole' or M2M operator, because it translates the multipole moments from one position to another. Using (4.4), expressions for calculating the multipole expansion of arbitrary order source particles can be written by considering a 'source' multipole \mathcal{S}_n . For a Coulomb charge, such that $\mathcal{S}_{(0,0,0)} = q$, and all other terms would be zero. For a dipole, $\mathcal{S}_{(1,0,0)} = \mu_x$, $\mathcal{S}_{(0,1,0)} = \mu_y$ and $\mathcal{S}_{(0,0,1)} = \mu_z$, with all other terms zero. Mixed systems can also be considered. Thus, in an arbitrary system where the lowest order of source is s , the expansion can be written:

$$\mathcal{M}_n(\mathbf{z}) = \sum_{|\mathbf{k}|=0}^{p-|n|} \frac{(\mathbf{z}_a - \mathbf{x}_a)^\mathbf{k}}{\mathbf{k}!} \mathcal{S}_{n-\mathbf{k}} \quad (4.5)$$

For the charge only case, we see that we can straightforwardly recover through Eq 4.3 through the knowledge that all terms except $\mathcal{S}_{(0,0,0)}$ are zero.

The potential can then be rewritten in terms of these Multipole terms - we can calculate the potential at a second particle, and so this is known as the 'Multipole-to-Particle' (M2P) operator.

$$\phi(\mathbf{x}_B - \mathbf{x}_A) \approx \sum_{\mathbf{n}=s}^p \frac{(-1)^\mathbf{n}}{\mathbf{n}!} \mathbf{r}_a^\mathbf{n} \mathcal{M}_n \nabla^\mathbf{n} \phi(\mathbf{x}_B - \mathbf{z}_A) \quad (4.6)$$

This and the other preceding expressions forms the mathematical basis of the multipolar Barnes-Hut method. However, one can go further and take another expansion, this time around \mathbf{z}_B , and truncating such that the maximum order of terms is the same, giving:

$$\phi(\mathbf{x}_B - \mathbf{x}_A) \approx \sum_{\mathbf{n}=s}^p \sum_{\mathbf{m}=0}^{p-|\mathbf{n}|-s} \frac{(-1)^\mathbf{n}}{\mathbf{n}! \mathbf{m}!} \mathbf{r}_b^\mathbf{m} \mathcal{M}_n \nabla^{\mathbf{n}+\mathbf{m}} \phi(\mathbf{x}_B - \mathbf{z}_A) \quad (4.7)$$

a two charge system is normally given as directed from positive to negative charge; in Physics this is reversed. It is important to check then which convention is used in any software package.

Grouping terms again in Eq. 4.7, a new expansion, ‘local’ to the target particle can be evaluated centred around \mathbf{z}_B using what is known as the ‘Multipole-to-Local’ (M2L) operator.

$$\mathcal{L}_{\mathbf{n}}(\mathbf{z}_B) = \sum_{|\mathbf{m}|=0}^{p-|\mathbf{n}|-s} \frac{(-1)^{\mathbf{n}}}{\mathbf{m}!} \mathcal{M}_{\mathbf{m}}(\mathbf{z}_A) \nabla^{\mathbf{n}+\mathbf{m}} \phi(\mathbf{z}_B - \mathbf{z}_A) \quad (4.8)$$

Then, the potential can be evaluated in terms of the local expansion with the ‘Local-to-Particle’ (L2P) operator.

$$\phi(\mathbf{x}_B) \approx \sum_{|\mathbf{n}|=s}^p \frac{1}{\mathbf{n}!} (\mathbf{x}_b - \mathbf{z}_b)^{\mathbf{n}} \mathcal{L}_{\mathbf{n}}(\mathbf{z}_B) \quad (4.9)$$

Derivatives of the potential then be calculated by differentiating this expression with respect to the component axis:

$$\frac{\partial^{\mathbf{k}} \phi}{\partial \mathbf{r}^{\mathbf{k}}} \approx \sum_{|\mathbf{n}|=s+|\mathbf{k}|}^p \frac{1}{(\mathbf{n}-\mathbf{k})!} (\mathbf{x}_b - \mathbf{z}_b)^{\mathbf{n}-\mathbf{k}} \mathcal{L}_{\mathbf{n}}(\mathbf{z}_B) \quad (4.10)$$

If the order of the derivative is greater than $p - s$, this expression is not sufficient. In this case, a finite-difference approximation must be used.

4.2 Implementation

4.2.1 Operator Generation

Here, we attempt to give a description of the open source code generation framework, fmmgen, [68] is implemented and how it can be used. The framework is built in Python, using the symbolic algebra package SymPy [173], and generates source code output in C and C++, with the reasoning that code generated in these languages by the framework can be straightforwardly incorporated into other projects without great difficulty or the requirement of large dependencies.

The code generation of each of the multipole operator equations can be broken up into different stages, each of which can be used independently. The user must specify the minimum source order s , the maximum expansion order p , and the output they desire (potential, field, or both). From these parameters, a mapping between \mathbf{n} values and one-dimensional array indices is created. By default this

mapping is lexicographic:

$$\begin{aligned}
 0 &: (0, 0, 0) \\
 1 &: (1, 0, 0) \\
 2 &: (0, 1, 0) \\
 3 &: (0, 0, 1) \\
 4 &: (2, 0, 0) \\
 5 &: (1, 1, 0)
 \end{aligned} \tag{4.11}$$

such that the total monomial order of a given term is strictly increasing. Nonetheless, if another ordering is preferred (for e.g. in some fields the quadrupole moments are ordered differently), it is possible to use change this by simply using a different array mapping. If it is known in advance that certain terms will always be zero, terms can be removed from the mapping in order to create simpler symbolic representations of the multipole and local expansion operators. We also make use of the source order parameter given by the user to reduce the memory needed to store the multipole and local expansions; this is possible because it is not possible to construct a multipole with a net n^{th} -moment from sources of order $s > n$.

A set of expansion functions are implemented, which are used to construct symbolic representations of \mathcal{M}_n , the particle-to-multipole operator, and \mathcal{L}_n , the multipole-to-local operator, at a given \mathbf{n} , as well as the shifting operators for these (the Multipole-to-Multipole and Local-to-Local operators). These must make reference to the mapping, in order to return the correct array indices. Generator functions use the set of expansion functions and iterate through the full list of n values needed for a particular problem, and an array representation of each operator is formed. This is repeated for each expansion order, and a least-recently-used (LRU) cache is used to store results which come up repeatedly in the generation stage to reduce the code generation time. We finally generate a symbolic representation of the operator functions for both the Barnes-Hut and Fast Multipole Method which can calculate the required quantities from a multipole (M2P) or local expansion (L2P), or from another source (P2P).

Once the full set of symbolic operators is generated, a code writing class is used to turn the mathematical representation of the operators into C or C++ code. While the SymPy library can provide some basic code-generation functionality, by default it generates unoptimised code which leaves much room for improvement in performance terms. To this end, we implemented a set of optimisations which can be enabled and disabled at the code generation stage by the user of the library. We

leave these as options rather than enabling by default, because it is then easy to test that the optimisations affect only the performance, and because the optimised code is often more difficult to read and hence debug.

Coles et. al. previously discussed how in code generation of multipole operators, [167] they reduce the number of mathematical operations in the code through an automated process called common sub-expression elimination (CSE), which analyses the code for repeated calculations across multiple lines, and pulls these out as factors. We utilise this technique, which is implemented in the SymPy library, to pull out expressions over the entire computational kernel, which can stretch to thousands of lines of code, enabling optimisations that would be impossible to perform by hand. Prior to using CSE, we pre-process the operators to increase the chance of finding common subexpressions. These preprocessing stages rationalise powers (for e.g. replacing $(x^2)^2$ with x^4 , factor terms, and remove extraneous multiplications which sometimes appear in the code generation stage (e.g. $(1.0)x$). In Figure 4.1, we show the effect that CSE has on the Multipole-to-Local operator. This has the greatest effect on performance in the calculation of the Multipole-to-Local operator for the FMM and Multipole-to-Particle operators for the Barnes-Hut method, which make use of the calculation of the derivatives of $1/r$ up to a given order.

In traditional codes, the computations of derivatives of these derivatives up to an order p are usually performed incrementally, such as by using the $\mathcal{O}(p^6)$ formula of Cipriani and Silvi [174] or using an $\mathcal{O}(p^4)$ recursive formula as described by Challacombe et. al. [175]. With the code-generation, we were able to implement straightforwardly an optimisation noted by Dehnen [157] by making use of the harmonicity of the Poisson Green's function, which allows us to calculate derivatives as:

$$\nabla^{\mathbf{n}+(0,0,2)}\phi = -\nabla^{\mathbf{n}+(2,0,0)}\phi - \nabla^{\mathbf{n}+(0,2,0)}\phi \quad (4.12)$$

This reduces the number of mathematical operations for higher order calculations. We note that while the SymPy library provides some metrics for the number of mathematical operations in given expressions, these are not an effective way of deducing the computational cost of generated code. The choice of compiler and the enablement of compiler optimisations drastically affects the FLOP count, and some operations take more clock cycles than others. This is especially true for Reduced Instruction Set Computer (RISC) processor architectures, on which all

```

1  void M2L_1(double x, double y, double z, double * M, double * L) {
2      double R = sqrt(x*x + y*y + z*z);
3      double D[4];
4      D[0] = (1 / (R));
5      D[1] = -1.0*x/(R*R*R);
6      D[2] = -1.0*y/(R*R*R);
7      D[3] = -1.0*z/(R*R*R);
8      L[0] += D[0]*M[0] + D[1]*M[1] + D[2]*M[2] + D[3]*M[3];
9      L[1] += D[1]*M[0];
10     L[2] += D[2]*M[0];
11     L[3] += D[3]*M[0];
12 }
13

```

(A) M2L Operator without CSE

```

1  void M2L_1(double x, double y, double z, double * M, double * L) {
2      double Rinv = pow(x*x + y*y + z*z, -0.5);
3      double D[4];
4      double Dtmp0 = (Rinv*Rinv*Rinv);
5      D[0] = Rinv;
6      D[1] = -Dtmp0*x;
7      D[2] = -Dtmp0*y;
8      D[3] = -Dtmp0*z;
9      L[0] += D[0]*M[0] + D[1]*M[1] + D[2]*M[2] + D[3]*M[3];
10     L[1] += D[1]*M[0];
11     L[2] += D[2]*M[0];
12     L[3] += D[3]*M[0];
13 }
14

```

(B) M2L Operator with CSE

FIGURE 4.1: Here, we see the affect of enabling common-subexpression elimination for the 1st Order Multipole-to-Local operator in the FMM method when $s = 0$. Prior to enabling this subexpressions such as $1/R^3$ are repeated multiple times across multiple lines of code (a), while with it enabled, these are factored out into temporary stack variables (b).

complex instructions are built up from simple instructions, rather than having dedicated complex instruction hardware. This means that for accurate FLOP counts, only sampling with runtime tools gives an effective approach for measuring performance improvements through code simplification.

The code also supports the replacement of evaluations of the power function `std::pow(x, n)`, where n is a positive or negative integer value, with repeated multiplications. It is well known that this can be an effective optimisation in numerical codes, but in practice it can be tedious to implement, and beyond a certain point round off errors begin to accumulate from intermediate results [176]. In the code generation stage, these operations can be replaced up to some maximum n_{\max} , the optimum which can be determined through profiling for a given architecture, precision and compiler combination.

4.3 Methods

We implemented both the Barnes-Hut and Fast Multipole Method in a single code base using the generated operator functions. We here provide some detail on the implementations.

4.3.1 Tree Decomposition

The key construct of the BHM and FMM is the choice of an appropriate data structure for containing information about the positions of particles - typically this is a tree. In two dimensions, a hierarchical decomposition of some bounded region of space gives a quadtree and in three dimensions an octree, so named because on splitting a cell, it divides into four and eight child cells respectively. A tree is constructed by taking a set of positions, and applying criteria for subdividing space recursively. This criteria is a limit either on the number of particles in a single spatial cell, or on the ‘depth’ of the tree, restricting how far cells can subdivide. Figure 4.2 shows an example of how this can look for a two-dimensional tree. Alternative types of tree, such as k-dimensional (k-d) trees have been utilised in multipole codes in the past [177, 178]. Constructing the tree can be done in many different ways, and the layout of the tree data structure is crucial to the performance of a fast multipole code, and for taking advantage of parallel architectures.

We use class based space decomposition, whereby a `Cell` data structure is used to hold references to the particles held within itself. Each cell holds a with a reference to it’s parent (the cell which contains it) and its children (the cells it contains). In our implementation, the decomposition of space is performed with a class based method. The code for this is straightforward to implement, because the code to split a cell can be written as a recursive function.

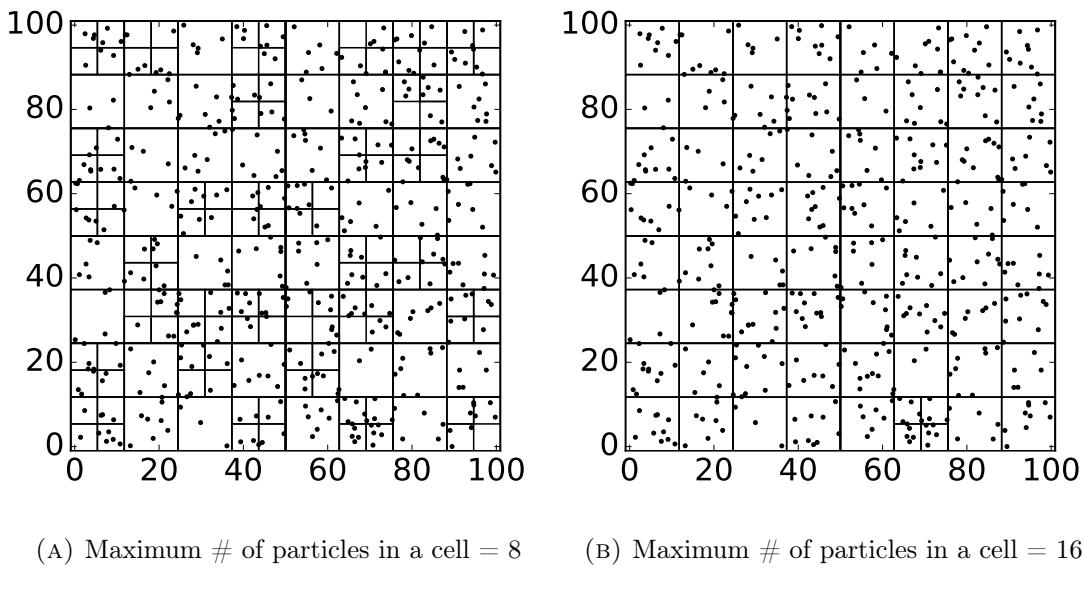


FIGURE 4.2: This figure shows spatial decomposition in two dimensions at different levels, for the same set of 500 particles. The decomposition of space here depends on the criteria for the maximum number of particles in a cell. When the maximum number of particles in a cell is reached, the cell is split into four more cells. This process is carried out recursively until the criteria is reached in every cell. We can see that when the maximum number of particles is 8, a finer discretisation of space is achieved. This type of spatial decomposition is known as a multi-level quadtree.

Algorithm 2 BuildTree()

n_c = Number of particles in cell before splitting
 Add root cell to list of cells
for all particles **do**
while current cell has more particles than n_c **do**
 current cell \leftarrow root cell
 Calculate octant of current cell where particle is
 if octant cell doesn't exist **then**
 create octant cell and push to list of cells
 end if
 current cell \leftarrow octant cell
end while
 Add particle to current cell
if current cell holds more particles than n_c **then**
 SplitCell(current)
end if
end for

Algorithm 3 SplitCell(cell)

```

for particles in cell do
    if cell octant does not exist then
        create octant cell and push to list of cells
    end if
    add particle to octant cell
    if number of particles in octant  $\geq n_c$  then
        SplitCell(octant)
    end if
end for

```

Because of the recursive nature of this class based tree splitting, it is difficult to parallelised the implementation of a class based tree building algorithm. However, this is only of interest in the case where the tree needs to be rebuilt many times, and it becomes a significant performance cost. In the case of spin dynamics, however, spins or micromagnetic cells are usually placed at fixed positions, which reduces the cost, and thus the whole tree does not need to be reconstructed. This does not mean that the tree data structure layout can be neglected, however, as traversing the data structure forms a significant part of the calculation. Alternative approaches for laying out the tree data include using space filling curves such as the Morton or Hilbert curves to map the position of each cell in some grid coordinate space to an index in memory for it's multipole and local expansion data [179, 180]. The order of cells in that data structure is defined such that each cell can find the cell indices of it's parent and child cells by simple calculation, rather than needing to store their indices as in a class based tree structure. Ordering in these ways also reduces how far in memory fetch operations must go, and so reduce cache invalidation, increasing performance. In order to mitigate this issue in the class based scheme developed here, multipole and local expansions are stored sequentially in a single array, with each cell holding a pointer to it's expansion data. This memory is allocated only once the tree has been constructed.

4.3.2 Tree Traversal

For the BH method, we evaluate the multipole expansion on cells at the lowest level of the tree, and then pass this upwards using the multipole-to-multipole (M2M) operators. Then, for each particle, located at \mathbf{x}_p the tree is traversed from the top level downwards. A cell is considered to be near to a particle if it meets the

Barnes-Hut multipole acceptance criterion:

$$\frac{r_{cell}}{|x_p - x_c|} < \theta_{BH} \quad (4.13)$$

which relates the cell size to the distance, and an opening angle parameter θ , which is a user supplied parameter which controls the accuracy.

If a cell has no child cells, and the cell does not meet the acceptance criterion, then the cell's particles are looped through, and the interaction is calculated directly using the Particle-to-Particle (P2P) operator. If the criterion is met, then the interaction between the cell and the particle is instead computed using the Multipole-to-Particle (M2P) operator. Finally, if the cell has child cells, then the procedure is repeated on these. For the FMM, we implemented the dual-tree traversal algorithm which has seen widespread adoption, rather than the classic FMM introduced by Greengard and Rokhlin in which cell-cell interactions only occur between neighbouring cells and their children, [181] as this has much in common with the Barnes-Hut approach. The initial procedure here is the same as the Barnes-Hut method; multipoles are computed for cells on the lowest level of the tree and then shifted upwards. Then, the tree is traversed from top to bottom. Cells which fulfil the multipole acceptance criterion:

$$\frac{r_{c_A} + r_{c_B}}{R} < \theta_{FMM} \quad (4.14)$$

interact via the Multipole-to-Local (M2L) operator, while cells which do not are recursed into until either their children fulfil the criteria, or a leaf cell is reached, at which point the cells interact directly. For more straightforward parallelisation, as opposed to the task-based parallelism favoured by some authors, we traverse the tree at initiation in our test implementation, and store the sorted interaction lists which can then be iterated through with loop-based parallelism.

4.4 Testing

We provide a test application with the library which can be configured to allow the evaluation of the potential and/or field from a set of source particles of arbitrary order, using either the Barnes-Hut or FMM approach, which allows for a straightforward comparison between the two methods and their performance. We ran tests with this text executable on a machine with a 4-core 3.4GHz Intel i7 6700 machine. We note that this processor is affected by the Spectre and Meltdown

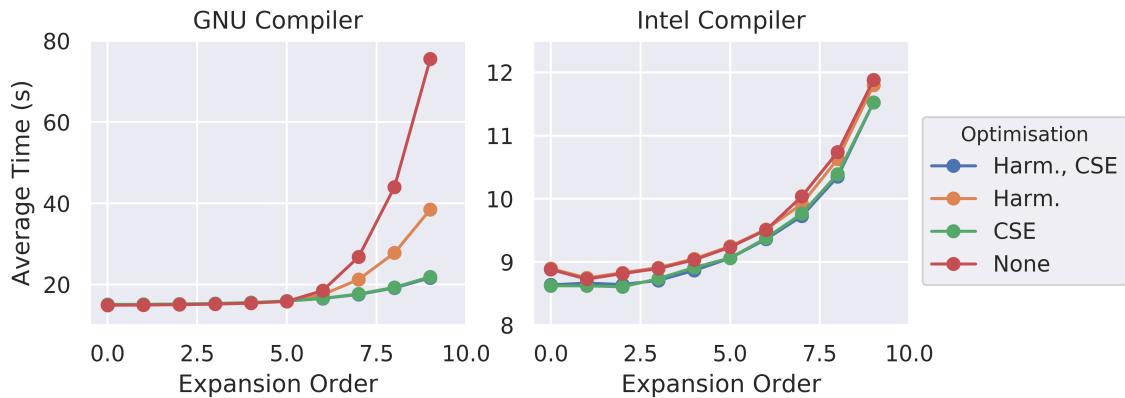


FIGURE 4.3: Performance comparison between compilers and with CSE and Harmonic Derivative evaluation disabled and enabled for a system of 10^5 charged particles.

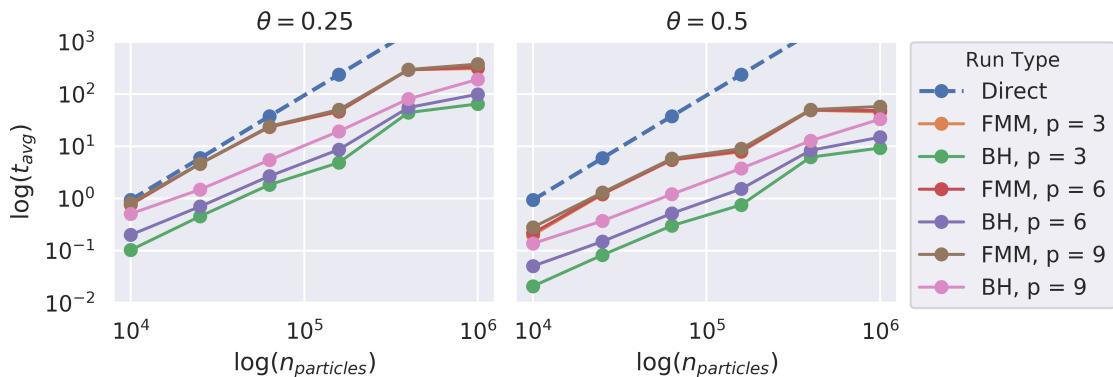


FIGURE 4.4: Performance between compilers and with CSE and Harmonic Derivative evaluation disabled and enabled.

vulnerabilities, and testing was performed with the Linux kernel version 4.15.0-55-generic, which includes mitigations for this, which have been reported to affect the performance of some HPC applications [182]. The executables were compiled with the Intel Compiler v.19.0.3.199.² Tests were performed with OpenMP enabled, and with options set to prevent thread migration between cores and idle threads from sleeping, and with hyperthreading disabled. All of the timing results shown below are averaged over three runs in order to reduce the effect of system calls and background processes on the runtime measurement.

Initially, we tested how the performance optimisations described in the previous section affected the performance of the potential and field calculation via the Fast Multipole Method, for a system of 10^5 randomly distributed charged particles

²The executables were compiled with the Intel Compiler v.19.0.3.199 and g++ v.7.4.0. The compilation flags in both cases were `-O3 -ffast-math -march=native -fopenmp`

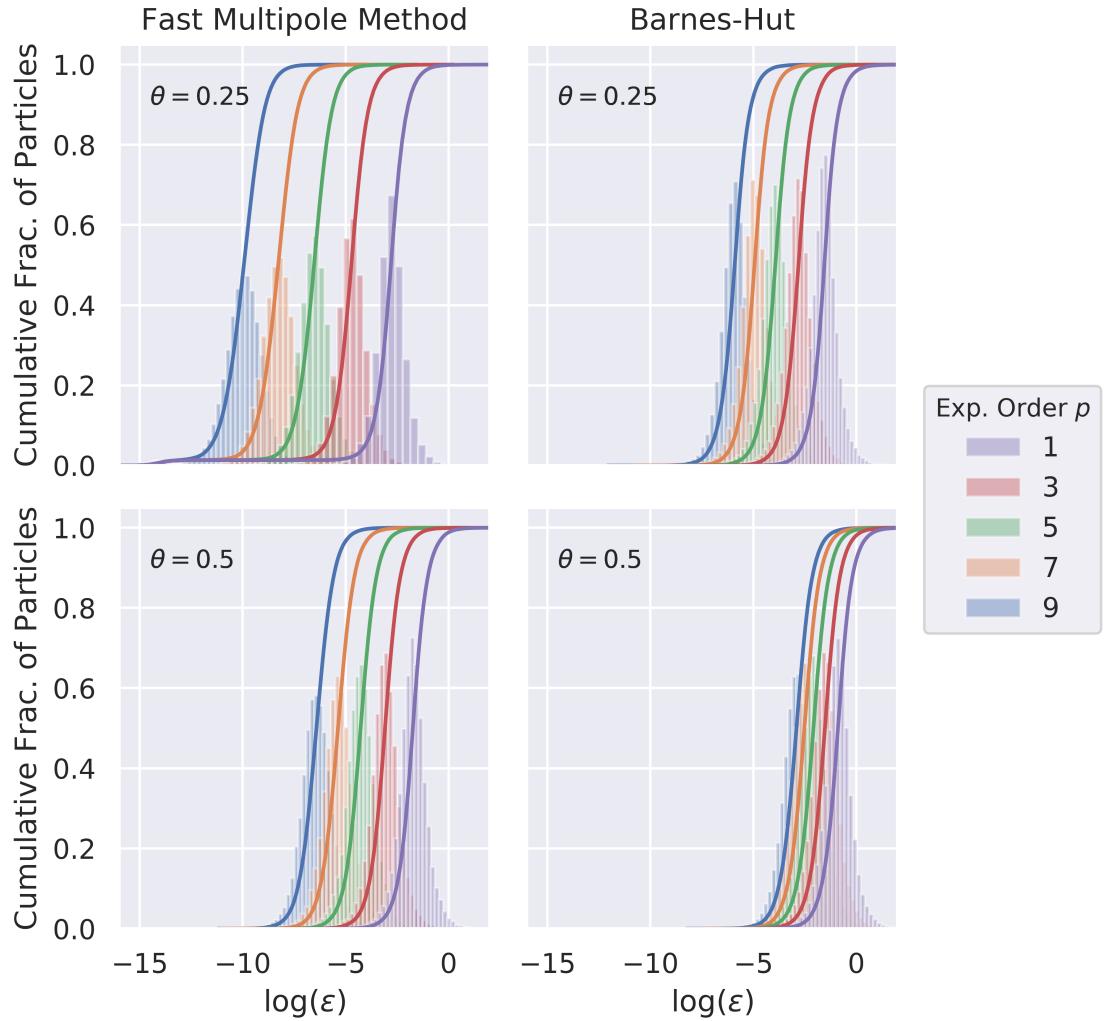


FIGURE 4.5: Here we show a comparison between the FMM and BH methods for two different values of θ . We can see that the BH method outperforms the FMM method in both cases, but that the BH expansion order affects the runtime of that method more substantially than for the FMM.

in $[-10^{-9}, 10^{-9}]^3$, with $\theta = 0.3$ and $n_{\text{crit}} = 128$. We compiled executables with generated operators with CSE and the computation of derivatives through the reuse of results and the harmonicity property enabled and disabled, the results of which are shown in Figure 4.3. We found that in general, the timing results were relatively consistent, with the runtime increasing progressively with the expansion order.

With the GNU compiler, enabling the harmonic derivatives optimisation led to a decrease in performance at 9th of around 50% while enabling CSE led to a

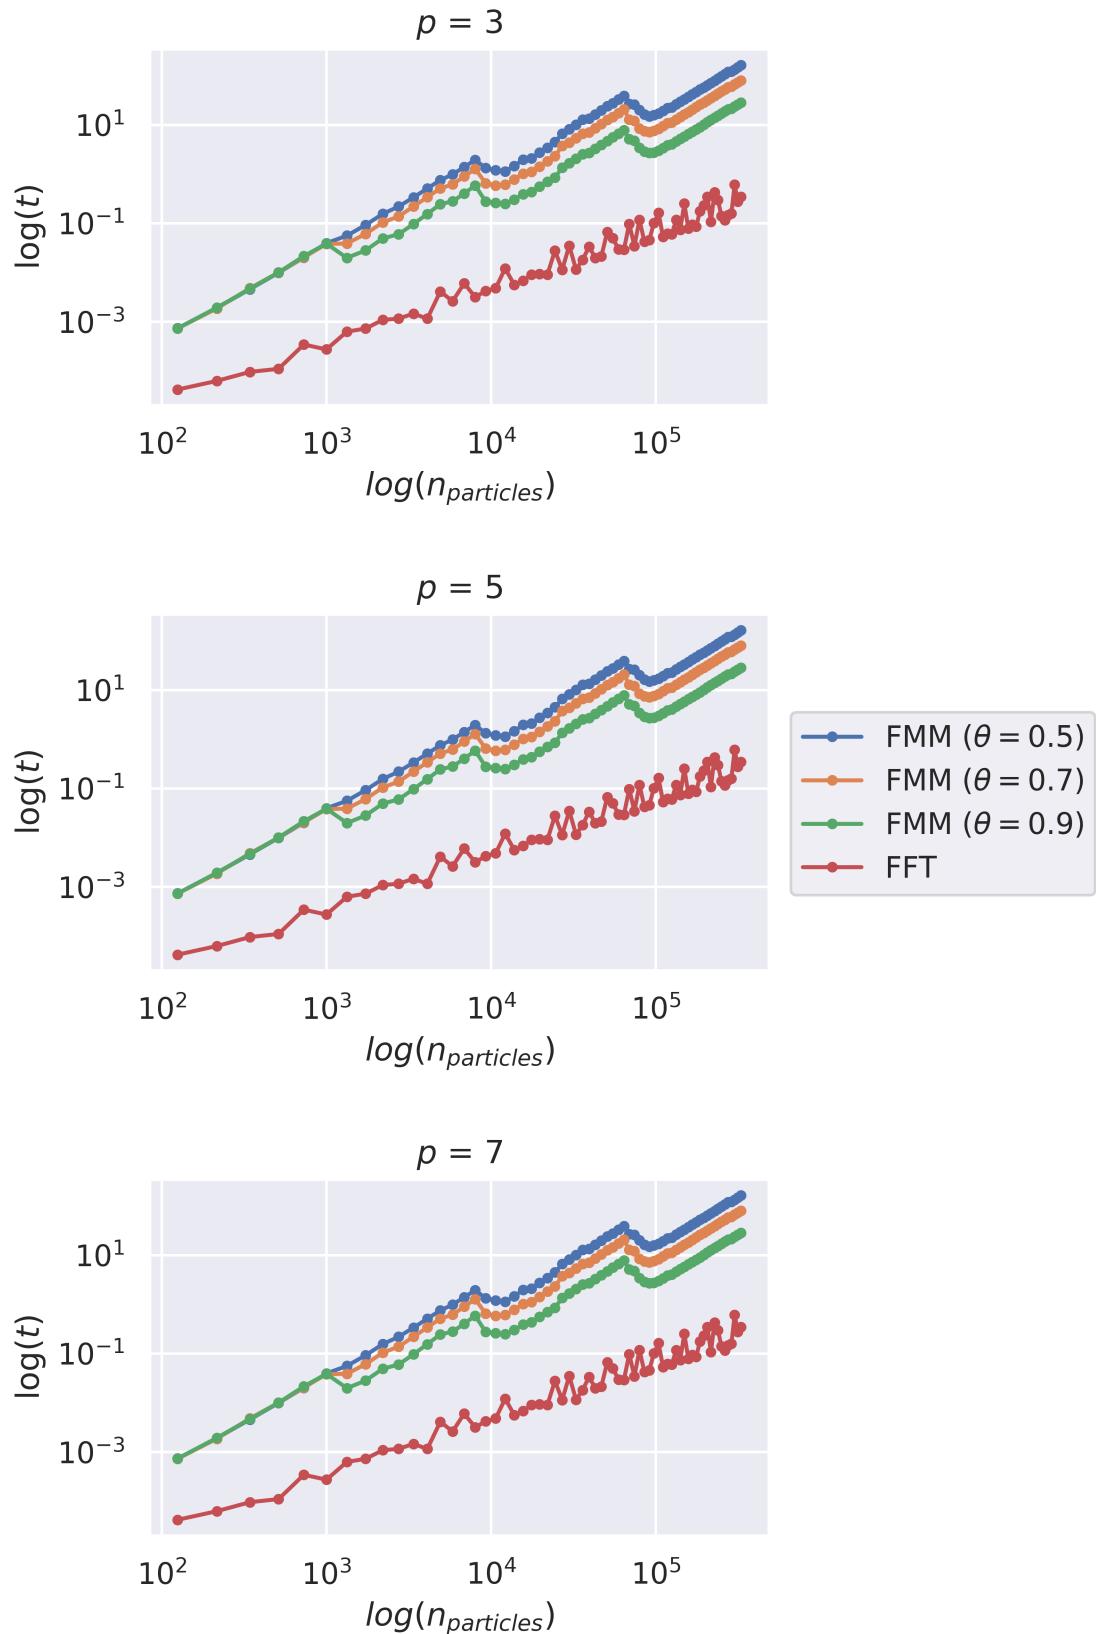


FIGURE 4.6: Here we show the performance of the FMM against the FFT for dipolar field calculation, by comparing runtime as the number of spins in a cube is increased at different expansion orders. We note that between expansion orders, we see a negligibly small difference in performance with variation in θ having much more of an impact on runtime.

75% decrease. For the Intel compiler, the corresponding decreases were around 1% and 2%. At lower expansion orders, we see very little performance increase, and this is because there are fewer opportunities for eliminating common factors. The difference between the GNU and Intel compilers was investigated. Analysis with Intel VTune showed that substantial numbers of the repeated operations at high optimisation levels were cached in compilation with the Intel Compiler. In both cases, there was not a significant difference in performance when both optimisations were enabled. All subsequent tests were run with the Intel compiled executable with both CSE and Harmonic derivatives enabled.

In Figure 4.4, we show the scaling of the FMM and BH methods with regards to the number of particles, where the number of particles is chosen so that points are equidistant in log-space. We can see that for numbers of particles up to 10^6 , the BH method outperforms the FMM. In both cases, the exact break even point over the direct method depends on the expansion order, but is less than 10^4 charges. We can see that increasing the expansion order gives a clear delineation of the runtime of the Barnes-Hut method while in the FMM, there is less of an impact; this is because the M2P kernel is evaluated many more times in the BH method than the equivalent M2L kernel is in the FMM method, and it is why the method scales more poorly ($\mathcal{O}(n \log n)$ for BH vs $\mathcal{O}(n)$ for the FMM) at large numbers of particles.

We note that the two multipole acceptance criterion are not directly equivalent, despite having a similar controlling effect on accuracy, because in the BH it directly relates the particle-cell distance and the cell size, while in the FMM it is a cell-cell distance and size parameter. As a result of this, to achieve similar error characteristics with the two methods, θ_{FMM} should be around twice θ_{BH} . This can be seen in Figure 4.5, where we show the error distributions for the two methods at different expansion order at $\theta = 0.25, 0.5$ for a system of 50000 particles with $n_{\text{crit}} = 128$.

4.5 Performance comparison with Fast Fourier Transform Approach for Dipolar Field

As a real-world test case for atomistic spin dynamics, we integrated the dipole field FMM calculation generated by the library into the computational nanomagnetism software Fidimag [95]. To check the implementation, we compared it against the

FFT convolution based dipole field solver for a system of atomic dipoles in a cubic arrangement with n spins on each axis resulting in n^3 spins. We compute the field with the convolution method using the Fast Fourier Transforms computed from the library FFTW (version 3.3.7) with OpenMP parallelism enabled. In order that the comparison is fair, we neglected any start up time which is one-off, and so do not include the pre-computation of the demagnetising tensor for the FFT technique or the tree construction for the FMM technique.

In all tests, it was found that the FMM method was over an order of magnitude worse in performance terms compared to the FFT convolution technique. We also note that for some θ values ($\theta > 0.7$), in realistic test simulations in which the Landau-Lifshitz-Gilbert equation was used to relax the system, we found that simulations either failed to converge using the FMM, or took more integration steps to do so, as a result of the loss of accuracy in the method. This suggests that, at least on parallel shared memory architectures, it is difficult to achieve competitiveness in performance against the Fast Fourier Transform convolution technique with the Fast Multipole Method.

Our results contradict prior performance studies on the fast multipole method in atomistic lattice systems, where the method showed speed-ups over the FFT convolution method for the numbers of particles commonly used in atomistic simulations. We note that the method shown in one paper promising speed-ups from the Cartesian FMM used the scalar non-parallelised FFT routine from Numerical Recipes [183], which was likely to be significantly slower than the FFT methods in FFTW (originally released in 1999, and with the much improved version 3 released in 2003 which added vectorised forms of the FFT) even at the time of publication [66]. Though we have chosen here to show the results by way of comparison with the FFT in FFTW due to this being freely available across architectures and operating systems, we note that performance of the FFT through the FFT interface supplied in Intel's commercial Math Kernel Library was found to be around 2.5x that of FFTW on the same hardware used in this study, and so the FMM fares even worse by comparison.

We did not attempt to introduce a micromagnetic version of the algorithm into Fidimag. This was because, using the technique of Visscher [67] due to needing to use the surface charges of each cell, we would expect the performance of the P2M part of the calculation to be six times worse than in the atomistic case. Given that the algorithm was not competitive with the FFT convolution technique even without this extra cost, and because it is more usual to consider

Despite this, the inclusion of the FMM method into our code Fidimag is designed such that it enables the study of systems where particles do not lie on a lattice, enabling the computation of the dipolar field in problems where it was not previously possible other than by direct computation in this software, such as in clusters of nanoparticles [184]. We note that the atomistic spin dynamics codes Vampire [58] and Vinamax [185] make use of approximations for computing the dipole field that appear similar to the Barnes-Hut method with $p = 1$ and $s = 1$. From our own tests, we found that approximation at this level of expansion order is not sufficient to maintain an acceptable level of accuracy in simulations in general, because it can lead to an error on the dipolar field of over 100% on individual particles. This may or may not manifest itself in simulations, and is strongly dependent on other parameters and the relative strength of the dipolar field against other energy terms.

4.6 Suggestions for Performance Improvements of Hierarchical Calculations

In order to achieve highly parallel performance, modern processors make great use of core and vector parallelism to improve performance. With core parallelism, each CPU has multiple cores which can act independently from one another and calculate different things, thereby reducing the computation time when all cores are applied to a given problem. Vector parallelism, on the other hand, allows parallel operation on data within a core. We show an example of this for multiplication of two arrays in Fig. 4.7.

Each modern processor architecture has a set of vector instructions can be utilised. Current flagship Intel processors make use of the AVX2 or AVX512 instruction set, which allows operations on 256-bit and 512-bit registers respectively. A floating point number is usually stored in 4 bytes (32 bits) as single precision, or 8 bytes in double precision, so up to 8 (16) floating point numbers or 4 (8) double precision numbers can be operated on in parallel once on current architectures.

Exploiting vector registers however, is difficult in practice. Compilers can vectorise loops to some extent but if branching occurs in a loop then vectorisation will fail. Processors also have cache memory to improve performance, and fetching from the cache is done in chunks (of 64 bytes - 8 floats or 4 doubles). The fastest FMM codes, ExaFMM makes use of explicit vectorisation in the P2P kernels,

while FalcON uses vector registers in the M2L kernel [157, 181] They do this by calling vector intrinsic functions [186] made available in language extensions, and which are available in some compiler families.

Implementing explicit vectorisation into code is not straightforward in general, however. Different vector instruction sets must be used for each processor architecture, and in general they are not backwards compatible, or compatible between processor vendors (though ARM’s implementation of vectorisation, Scalable Vector Extension, is more general and does not require knowledge of the vector register size in advance). The calling conventions for vector intrinsic functions can also differ between compilers, and so in a HPC code which must generally be compilable by at least GCC and the Intel compilers, and preferably also Clang, Cray and Portland Group, it can require a large amount of work to write performance-portable code. Tools such as the Vector Class Library [187] and xsimd [188] exist which aim to allow software developers to write portable vectorised code by making use of their data structures rather than raw arrays. In general, using a library such as this is the easiest way to implement vectorisation in a cross-platform manner.

At present, no explicit (i.e. hand-implemented) vectorisation is done in the operator functions generated by fmmgen. This is likely a major source of the difference between the performance of the convolution approach and the FMM, as in general, FFT libraries are highly optimised. However, it is straightforward to see how vectorisation can be implemented. Considering the various kernels, if particles are sorted by their parent cell, precomputed sparse matrices multiplying column vectors holding the updates spin values and operated on in parallel order to compute the operators. We suggest that to increase performance, this could be incorporated into the code generation in an explicit manner much more straightforwardly in code generation than by hand.

```
1 #include <immintrin.h>
2
3 void scalar_multiply(double *a, double *b, double *c, int N) {
4     for(int i = 0; i < N; i++) {
5         c[i] = a[i] * b[i];
6     }
7 }
8
9 // Byte alignment - hardware specific.
10 constexpr int align = 32;
11 // Number of values each vector intrinsic function can act on:
12 constexpr int k = align / sizeof(double);
13
14 void vector_multiply(double *a, double *b, double *c, int N) {
15     __m256 reg1;
16     __m256 reg2;
17     for(int i = 0; i < N/k; i++) {
18         int o = i*k;
19         reg1 = _mm256_load_ps(&a[o]);
20         reg2 = _mm256_load_ps(&b[o]);
21         reg1 = _mm256_mul_ps(reg1, reg2);
22         _mm256_store_ps(&c[o], reg1);
23     }
24 }
25 }
```

FIGURE 4.7: Scalar and vector multiplication functions. The vector multiply function here makes use of Intel AVX2 instructions, and breaks the loop into chunks which are processed on vector registers.

4.7 Discussion

In this work, we have implemented and shown the efficacy of code generation for the multipolar Barnes-Hut and Fast Multipole Methods, and have described the implementation of this into an publicly available framework. While we have achieved substantial increases in performance over the direct method, there are several areas in which further progress can be made. Notably, the use of an irreducible representation of the operator functions through the use of a detracing operator can reduce the storage space needed for the Cartesian FMM [168]. We also note that performance was disappointing compared to the FFT-convolution method which is commonly used in dipolar atomistic simulation. It is worth stating, however, that, the code developed still has utility for more general spin dynamics simulations, whereby atoms are not placed on lattices. The study of Stoner-Wohlfarth nanoparticles which are not distributed on a lattice has been of much recent interest due to their applications in medical technology [185]. In these types of simulations, the methods developed and made public in this work have considerably better performance for small numbers of particles.

It is important to note that while the algebraic complexity of the spherical harmonic expansion is lower (at $\mathcal{O}(p^4)$ for a naive implementation or $\mathcal{O}(p^3)$ when rotations are used to reduce the cost of the local expansion translation), at low orders it has been shown by various authors that the computational cost of using the Cartesian method is often still lower, despite scaling as $\mathcal{O}(p^6)$. It has, however, been shown by the proliferation of consumer-grade GPU hardware in computational research that in many cases, accuracy of less than 10^{-7} is sufficient in many numerical applications. It is with this in mind that there is still much to recommend about the Cartesian approach over the Spherical Harmonics technique.

It is of our opinion that the specialised nature of many libraries towards specific problems means that hierarchical methods have not been as successfully adopted as other numerical techniques, and indeed, part of our own motivation for this work was in the difficulty of applying existing packages to our own problems of interest, namely nanomagnetic dipoles. We note that, for example, in gravitational systems where the domain origin is chosen as the centre of mass, the dipole moment will always vanish [168, 189]. A specific and widely used optimisation for the fast multipole method in this case, therefore, is to neglect entirely the calculation of the dipole moments in a system, which precludes the reuse of a hand-written gravitational FMM code for other applications where the dipole moments are non-vanishing, without some modification.

After the course of work was completed, it was brought to the attention of the author that an updated version of the work by Coles and Bieri [190] had been published [168], in which the previously closed source code for generating multipole operators had been released openly as the software package ‘Mosaic’, available freely on GitHub [191].

Chapter 5

Magnetic Skyrmions

5.1 Introduction

Skyrmions are a type of topological defect first predicted by Tony Skyrme in the context of field theories in Nuclear Physics [192]. The skyrmion's nature is defined by the property that it can not be continuously deformed from the lowest energy state (sometimes known as the vacuum state) to the defect state. This concept can be thought of simply in the case of a Möbius strip; the twist in a continuous piece of paper cannot be removed without ripping the paper. This affords the skyrmion a degree of 'topological protection' which leads to an increased stability when compared to other states. In the context of magnetism, the skyrmion is a quasi-two-dimensional structure which originates from the competition between different magnetic interactions. The skyrmion in micromagnetic theory is characterised by the topological charge [193]:

$$S = \frac{1}{4\pi} \int \int dx dy \cdot (\partial_x \mathbf{m} \times \partial_y \mathbf{m}) \quad (5.1)$$

The presence of a skyrmion in an infinite continuous field gives a winding number of 1. Along the radius of the skyrmion, the magnetisation in the out of plane direction varies from $-M_s$ to $+M_s$ continuously. The two types of skyrmion which form, the Néel skyrmion and Bloch skyrmion (shown in Fig. 5.1) differ in how the magnetisation rotates along this radius, and they are named for the respective types of domain wall which also rotate in the same way.

The existence of the skyrmion state in magnetism was first predicted by Bogdanov and Yablonskii [194]. In this work, using the mean-field Ginzburg-Landau (GL)

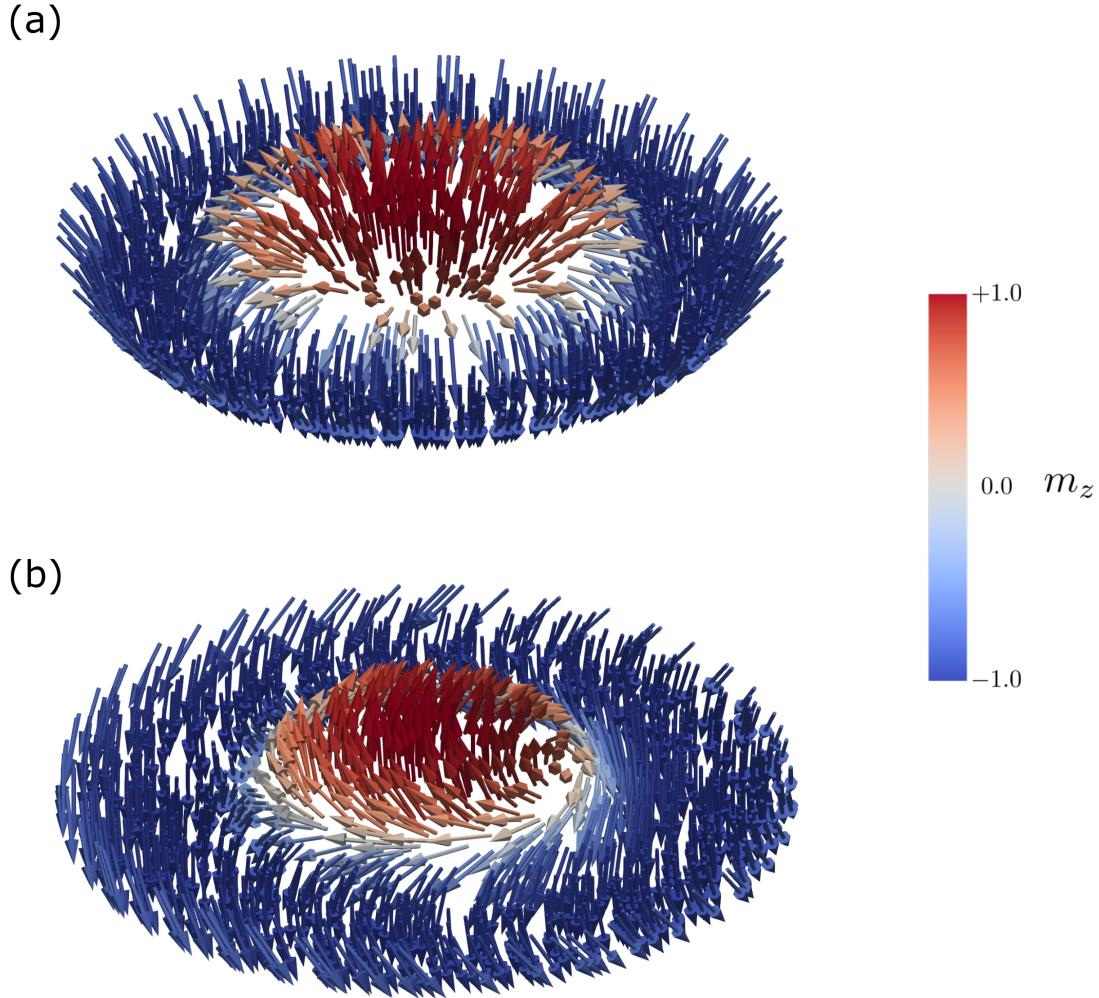


FIGURE 5.1: This figure shows (a) a Bloch Skyrmion and (b) a Néel Skyrmion. Bloch Skyrmions are found in materials such as FeGe and MnSi, which have a bulk Dzyaloshinsky-Moriya interaction, while Néel Skyrmions are found in materials with interfacial Dzyaloshinsky-Moriya interaction. While the m_z component of both types of skyrmions are the same, the magnetisation rotates differently along the skyrmion radius.

theory in an analogous way to the work done by Abrikosov in predicting the vortex lattice [195], the authors added a term corresponding to the Dzyaloshinskii-Moriya interaction (DMI) in bulk crystals of various crystallographic classes to the free energy functional, and showed that it admitted an inhomogeneous mixed ‘vortex’ lattice state which could be formed in a particular range of high applied magnetic fields, as the free energy of this state was lower than that of the spiral (or helical) and ferromagnetic states under these conditions. This was later extended in a work by Bogdanov and Hubert where a phase diagram between the combined effect of the applied field and anisotropy and the other material parameters was calculated, and showed a triple point between the ferromagnetic state, the skyrmion lattice

and the spiral state [28]. The thermal effects were later taken into account, and the prediction of a spontaneous skyrmion in two dimensional samples in a finite temperature range was made [29]. The skyrmions predicted by these papers are general; different crystallographic classes admit energetically favourable skyrmions of different types.

In materials where symmetry breaking occurs at the boundary, a different type of DMI occurs, and the strength of this interaction can be significant [196, 197]. A later paper by Bogdanov and Rößler added a different DMI term phenomenologically to the GL energy functional to model the effects of this DMI, and showed that skyrmions could be seen outside of the bulk types of crystals in which they had previously been predicted [198]. The skyrmions predicted to form in these types of systems have a specific form of the DMI, and this leads to the Neél skyrmion.

The first experimental study in which skyrmions were indirectly observed was published by Mühlbauer and collaborators in 2009. In this work, samples of the bulk B20 material MnSi were subject to an incident neutron beam applied parallel to a magnetic field after field cooling, and the scattering pattern showed sixfold symmetry, in an ‘A-Phase’ characteristic of a skyrmion lattice. The skyrmion lattice was observed in samples cooled to between 25 and 35 Kelvin, and subject to an applied field of between 0.1 and 0.2T [199]. A subsequent work by the same group showed that an additional contribution to the Hall resistivity which could be explained only by the Topological Hall Effect, giving further evidence that the ‘A-Phase’ was indeed caused by Skyrmions [200].

Observations of skyrmions in a broader range of materials has since been carried out. FeGe has been particularly well studied, and it was in this system which the first real space images of skyrmions were obtained using Lorentz Transmission Electron Microscopy (LTEM) [201, 202]. It was determined that the presence of the skyrmion lattice in FeGe was critically dependent on the thickness of the sample, with the extent of skyrmion stability in the $\mathbf{B} - T$ phase diagram smaller in thicker films. Theoretical work has predicted that chiral modulations occur along the core of a skyrmion when extended through a sample, which significantly changes the energy landscape of skyrmionic states. In particular, for thicker films, this is energetically unfavourable, which explains these experimental observations [203]. Later works extended real-space imaging to other B20 materials, with a particular achievement being the imaging of the skyrmion lattice in MnSi, with the much shorter helical period of 19 nm [204]. Skyrmions have been observed in the insulating B20 multiferroic material Cu_2OSeO_3 [205, 206].

In addition to the observations of Skyrmions in bulk crystals, much experimental focus has been on thin-film and ultra-thin film systems [34, 207–209]. The motivation behind this is because by careful engineering, the Dzyaloshinskii-Moriya constant can be tuned. The first experimental observation of a skyrmion lattice in such a structure was made in a hexagonal monoatomic layer of Fe on a surface of Ir(111). In this particular system, the skyrmion lattice was observed to have twofold symmetry, and was imaged using spin-polarised scanning tunnelling microscopy (SP-STM) [210]. Since then, the number of experimental studies in this area has seen enormous growth. By now, the tuning of the material parameters through the use of multi-layers has been tried by many experimental groups, with notable achievements the realisation of room temperature skyrmion stability [211, 212] and skyrmions as small as 10 nm [213].

Experimental evidence of skyrmion motion was obtained first in MnSi samples, where the application of an in-plane current through a sample placed under a magnetic field or temperature gradient caused rotation of the skyrmion lattice phase, was seen by rotation of the diffraction pattern [214]. Direct imaging of skyrmion motion was obtained soon after, with Lorentz TEM imaging of FeGe under the application of a current showing a low critical current for skyrmion motion of less than 10^2 A cm^{-2} [215]. This study showed that the induced skyrmion motion had both translational and rotational components. Another key experiment which showed the use of spin transfer torques in skyrmion manipulation demonstrated the creation and destruction of individual skyrmions. In a three layered Pd/Fe/Ir(111) system, in which the PdFe acts as a single magnetic layer, the application of a voltage sample allowed selective injecting of a skyrmion into the magnetisation of the sample [33].

Much has been made of proposed technological applications of Skyrmions, and research in this direction has been a large focus of interest [36, 38]. In terms of technical applications, it is rare to find a scientific paper in this area without promises that skyrmions can revolutionise magnetic storage devices or electronics. The skyrmion is heralded as a successor to magnetic domain based storage in many papers due to the aforementioned degree of topological protection. Despite this, it is well known that any such protection is limited; while it is true that the destruction of a skyrmion in an infinite magnetic system in a continuum theory leads to an infinite energy barrier and as such the skyrmion cannot be destroyed, magnetic materials are not continuum but discrete. It has been shown now by many authors [132, 134, 216] that the energy barriers to the skyrmion destruction are not impassable, and this is discussed further in Chapter 6. This is somewhat

obvious - if the barrier to creating or destroying a skyrmion was infinitely high, we would never observe skyrmions in materials at all. Nonetheless, the works in which barriers have been calculated have suggested that the skyrmion energy barrier is still high enough to provide a level of thermal stability suitable for magnetic storage. It is perhaps relevant to mention that this increased interest in skyrmions comes at a time of reduced demand for magnetic storage devices; as of early 2020 sales of magnetic hard disk drives are dropping quarter by quarter because of the rise in semiconductor based solid state drives. This has left researchers working towards alternatives which can in the future supplant existing technologies.

Probably the most widely discussed idea for applying Skyrmions towards technological applications to has been that of using ‘skyrmion racetracks’ for data storage, proposed by Fert and collaborators in a commentary piece in 2013 [26]. The diameter of skyrmions in multilayered systems is to some extent tunable through the layering of different materials, but is generally around an order of magnitude lower than conventional magnetic domain storage bits, which motivates their use in storage devices. The proposed mechanism behind racetrack storage is that a current is applied in-plane, causing the skyrmions to move along a nanowire. This is a logical extension of previous work which used the idea of domain walls in a racetrack [22]. In research into domain wall based racetrack storage, it was found that the current-driven motion of domain walls was inhibited by local defects where there is either no magnetic material or an increased anisotropy, making the domain walls difficult to drive reliably, and causing data loss [217]. While several authors have shown the motion and creation of skyrmions in racetracks both theoretically [25, 37, 39, 218–220] and have shown skyrmion motion experimentally [212], recent studies have suggested the same problems with pinning and maintaining data integrity are likely to occur with skyrmions [44, 221]. While the threshold current for skyrmion motion is lower than that of domain walls, the skyrmion velocity is comparable to the domain wall velocity above the domain wall critical current, and any proposed storage device would rely on the current being above this value [212, 215]. An additional problem is that the skyrmion also moves with a deflection due to the topological hall effect, driving the motion towards the sample edge in a racetrack [222].

Detecting skyrmions reliably is still an ongoing issue. While imaging techniques can be used, for devices this is not sufficient, and the magnetostatic field measured at a distance from a skyrmion is low. The topological Hall effect component of the resistivity changes in the presence of a skyrmion, but the skyrmion also moves in the presence of a current, which influences data integrity, and there have

been many proposed theoretical electrical detection mechanisms [26, 223, 224]. The experimental study of such structures is challenging, though several groups have now fabricated samples in which individual skyrmions have been electrically detected, though at temperatures much lower than would be feasible for realistic devices [46, 224, 225].

The dynamic behaviour of skyrmions in the bulk lattice regime is now relatively well understood [226, 227]. However, in practice, to fabricate useful devices, the nanostructure hosting the skyrmion needs to be of a size not much larger than the magnitude in size as the skyrmion itself, as only by this will it be able to compete with other avenues of research. To this end, it is important to understand how skyrmions behave in isolation in confined nanostructures. In 2013 it was shown that Néel skyrmions could be admitted into cylindrical nanodisks as a lower energy state than the ferromagnetic state with an interfacial type DMI [31]. In FeGe disk systems, a comprehensive study similarly showed that the ground state energy was either a skyrmion or an ‘incomplete skyrmion’, also known as a radial vortex. This study also investigated hysteretic behaviour of these disk systems [32], with skyrmion state destruction mediated by a Bloch point shown. This work was expanded to investigate the dynamic behaviour of skyrmions in disk systems, by studying the resonance modes [228]. As previously mentioned, due to the discrete nature of magnetic systems, there is no ‘true’ topological protection afforded to the skyrmion in magnetic systems; a finite energy barrier always exists between the ferromagnetic and skyrmion state. In particular, work has shown skyrmion creation at the sample edge can be energetically preferable, and strategies for skyrmion creation here. [132, 221, 229, 230].

Experimental measurements show that isolated skyrmions can exist for long timescales well outside of the small pocket of the temperature-field ‘A-phase’ in which the skyrmion lattice is found in many materials. However, as discussed by several authors [32, 231], there has still been relatively little investigation as to the both the behaviour and stability of isolated skyrmion states, and how the shape of nanostructures affects this in bulk chiral materials. This is despite it being well known that boundary shapes play a key role by inducing shape anisotropies in the energy landscape in ordinary ferromagnets [232]. A particularly surprising paper showing that the understanding of isolated skyrmions in bulk ferromagnets was the recent work of Du [233] which showed both in experiment and simulation that the pairwise interaction between Skyrmions can change from attractive to repulsive at low temperatures.

In this chapter, the energy landscape of FeGe nanostructures of triangular and square shapes of different sizes is investigated, and a ground state phase diagram for each state is constructed for each type of sample by varying the size and applied field. It is shown that the energetics are substantially modified in comparison with similarly sized nanodisks. An experimental procedure for observing skyrmions is proposed for FeGe nanostructures, and a range given over which one could expect to find Skyrmion states in such samples. The work shown here was published as "Skyrmion states in thin confined polygonal nanostructures", R. A. Pepper, et. al. Journal of Applied Physics 123, 093903 (2018) [234]. The data used in this study and scripts to generate the figures are available as 'Dataset for "Skyrmion states in thin confined polygonal nanostructures"', (Version 1.0.0), Zenodo (2017) [235]. The division of work in this was as follows. The present author, Dr. Marijan Beg and Prof. Hans Fangohr conceptualised the idea behind the work and interpreted the results. The present author wrote and executed the simulations, analysed the data and prepared the manuscript. All authors on the paper worked on the simulation software Finmag behind the study.

5.2 Skyrmions in Thin Polygonal Nanostructures

5.2.1 Method

Through micromagnetic simulations, film systems of FeGe of thickness 10 nm are studied using a fully three-dimensional model. This model is chosen as it has been rigorously shown both theoretically and experimentally that in films of cubic helimagnets, chiral modulations occur along all three spatial dimensions, which reduces the skyrmion state energy in 3D systems of thickness lower than the helical length - the length over which the magnetisation undergoes a full rotation [236–239]. The dynamics of the magnetisation field \mathbf{m} are modelled by the Landau-Lifshitz-Gilbert (LLG) equation

$$\frac{\partial \mathbf{m}}{\partial t} = \gamma_0^* \mathbf{m} \times \mathbf{H}_{\text{eff}} + \alpha \mathbf{m} \times \frac{\partial \mathbf{m}}{\partial t}. \quad (5.2)$$

Here, $\gamma_0^* = \gamma_0(1 + \alpha^2)$ where γ_0 is the gyromagnetic ratio, and $\gamma_0 < 0$. The constant α is the Gilbert damping coefficient. The effective magnetic field is calculated as $\mathbf{H}_{\text{eff}} = -(\delta w / \delta \mathbf{m}) / (\mu_0 M_s)$, where w is the total energy density given as:

$$w = w_{\text{Exchange}} + w_{\text{DM}} + w_{\text{Zeeman}} + w_{\text{Demag}} \quad (5.3)$$

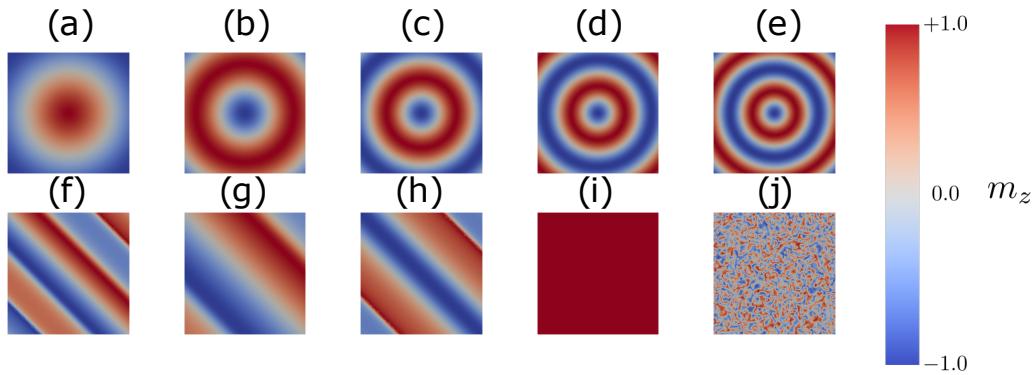


FIGURE 5.2: Initial magnetisation configurations from which each geometry is relaxed, shown here in a 140 nm side length film. The states are (a) incomplete skyrmion, (b) isolated skyrmion, (c) and (d) over-complete skyrmions (e) target state (f), (g) and (h) helical states of different helical lengths (i) uniform state, and (j) random state.

The symmetric exchange energy density is $w_{\text{Exchange}} = A(\nabla \mathbf{m})^2$ where A is the magnetic exchange constant. The bulk Dzyaloshinskii-Moriya interaction (or anti-symmetric exchange) in a material of crystallographic class T is given as $w_{\text{DM}} = D\mathbf{m} \cdot (\nabla \times \mathbf{m})$ where D is the DMI energy constant. The Zeeman energy is calculated from the applied field \mathbf{H} as $w_{\text{Zeeman}} = -\mu_0 M_s \mathbf{m} \cdot \mathbf{H}$. The demagnetising field is calculated using the Fredkin-Koehler hybrid FEM/BEM method.[117] For the simulations of FeGe, we use the parameters [32] $A = 8.78 \times 10^{-12} \text{ J m}^{-1}$, $D = 1.58 \times 10^{-3} \text{ J m}^{-2}$, $M_s = 3.84 \times 10^5 \text{ A m}^{-1}$. The finite-element discretisation was set such that the distance between mesh nodes was no greater than 3 nm, smaller than the relevant micromagnetic length scales for the given material, which has a helical length of 70 nm and exchange length $l_{\text{ex}} = \sqrt{\frac{2A}{\mu_0 M_s^2}} = 9.67 \text{ nm}$

The model chosen does not take into account temperature fluctuations. FeGe has an ordering temperature of 278.7 K [240] In large samples of FeGe studies showing skyrmion states have been performed within the temperature range of between 5 K and 278 K which show the skyrmion lattice state. In a study by Zheng et. al. [241], robust target states, where the magnetisation rotates over two periods over the width of the sample geometry, were observed in nano-pillars of FeGe with Lorentz TEM measurements taken at 95 K. With this in mind, the effect of temperature fluctuations is not expected to be substantial, so long as any corresponding experimental studies are done at temperatures below the ordering temperature. A ground state phase diagram is computed for two types of confined FeGe sample; square and triangular films of 10 nm thickness, through dynamic simulations. Dynamic simulations are used in order that all discovered states

are physically realisable, as direct energy minimisation can in some circumstances find states which are difficult to reach due to large energy barriers. The energy landscape of these structures is explored by changing the applied magnetic field, which is varied between 0 mT and 800 mT, and which is applied in the z direction, into the plane of the system. In squares samples, films which have a side length between 40 nm and 180 nm are simulated, while in triangles a larger range of side lengths between 40 nm and 220 nm so that coverage is given to samples of similar volume. Initially, the magnetisation of each point in the phase space is set to each configuration of a set of initial states; the definition of these states is the same as those used in the study of Beg et al [32]. The set of initial states, which includes uniform magnetisation, skyrmionic state profiles, helical profiles, and a random magnetisation state (which is repeated three times) are shown in Figure 5.2. This systematic exploration is done in order to capture as many equilibrium states as possible for each simulated system. In order to construct the ground state phase diagrams, we relax systems from these initial states under the LLG equation, until the system has settled into a local (or global) minima in the energy landscape. States are considered to be in equilibrium, and simulations are stopped, when the value of $|\partial\mathbf{m}/\partial t|$ is less than a tolerance of 0.01 degrees per nanosecond, at which point the magnetisation is no longer changing. A damping factor value of $\alpha = 1$ in order to achieve convergence to the final states quickly, by suppressing the precessional dynamics, which does not affect the final state. Once the dynamics have subsided according to the above criterion, the total energy of these relaxed states is computed. The lowest energy state that is found is then identified (from the total set of simulations starting from different initial configurations) as the ground state for the given geometry and applied field value, which allows the construction of a $d - B$ phase diagrams of the ground states. High energy relaxed states are considered to be metastable. To perform the simulations, the finite-element micromagnetic simulator Finmag, [115] is used. For time integration, a preconditioned BDF method is used from the SUNDIALS library [242].

5.2.2 Results

5.2.2.1 Equilibrium States

A wide variety of equilibrium states (formed of both the ground and meta-stable states) are obtained from the simulations in the systems, and in Figure 5.3, the regions in $d - B$ phase space where each state can form as an equilibrium state

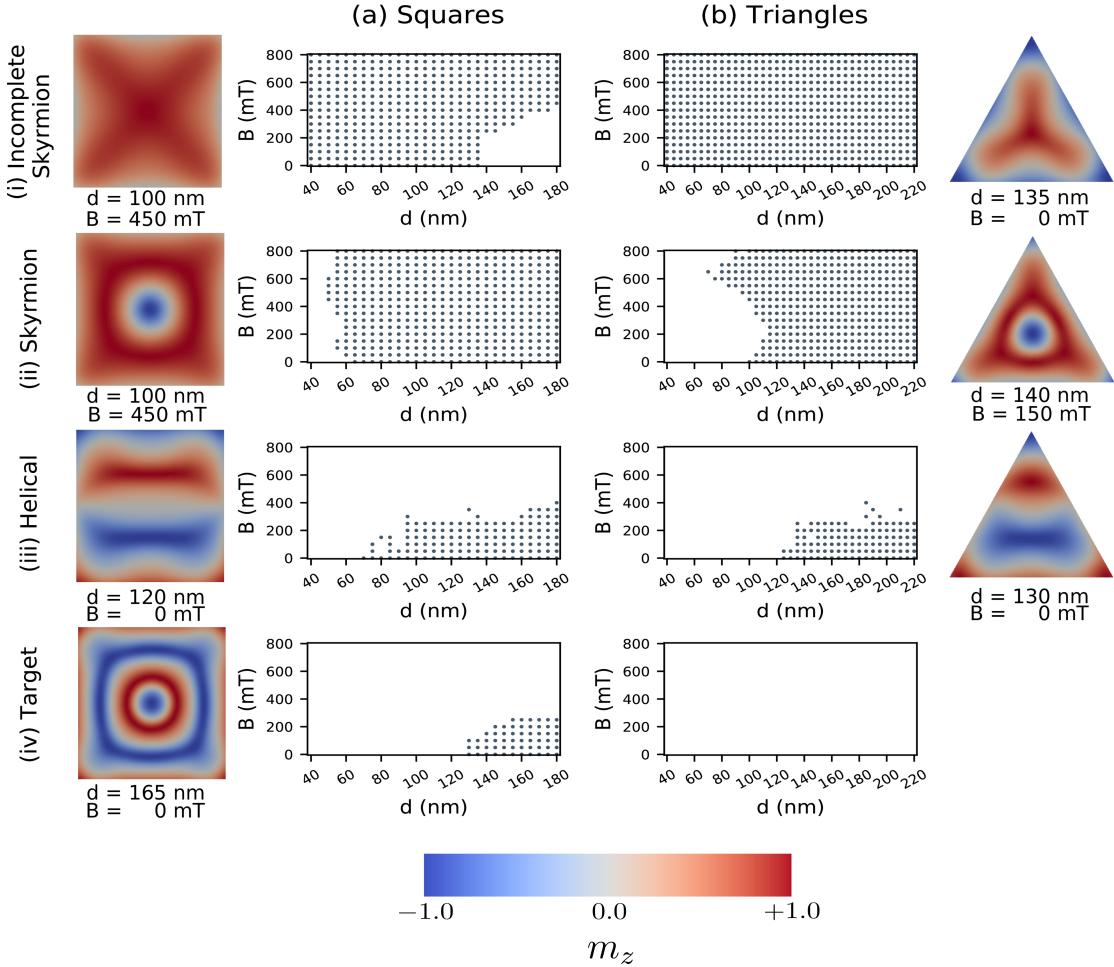


FIGURE 5.3: Equilibrium regions for states found in different geometries. In the first and last columns the z -component of the magnetisation for examples of states (i) to (iv) in the square and triangle systems is shown. The d - B graphs show dots when a meta-stable state of that type was found for that size and applied field. Incomplete skyrmion states (top row), are not stable for large square systems, with a field lower than around 400 mT. However, this is not seen in triangles. In contrast to squares, target states are not observed as meta-stable in any region of phase space studied in triangles.

are shown. The equilibrium states can be broadly classified into several groups.

1. *Incomplete Skyrmions* - These states are named [32] as such due to the presence of a quasi-uniform magnetisation across the system, and a skyrmion-like core. Due to the DMI, at the boundaries magnetisation is twisted. These states are known by many different names in the literature; other authors give the labels ‘radial vortex’ [243], ‘quasi-ferromagnetic’ [31] and ‘edged-vortex’ [244] in the literature to identify this type of state. (Figure 5.3 (i))
2. *Isolated Skyrmions* These states, normally axially symmetric in disks, are

distorted by the boundary of the confined geometry in both triangular and square systems. (Figure 5.3 (ii))

3. *Helical States* A large variety of rotational spin textures form metastable states in the studied systems. (Figure 5.3 (iii))
4. *Target States* Target states can be considered as an isolated skyrmion, with an additional radial half-helical rotation. (Figure 5.3 (iv))
5. *Skyrmion Clusters* Multiple clusters of skyrmions form metastable states in the geometries when strong fields are applied to the system, resulting in a smaller skyrmion radius. These states are found as high-energy metastable states for larger system sizes and for high applied fields. (Figure 5.4)

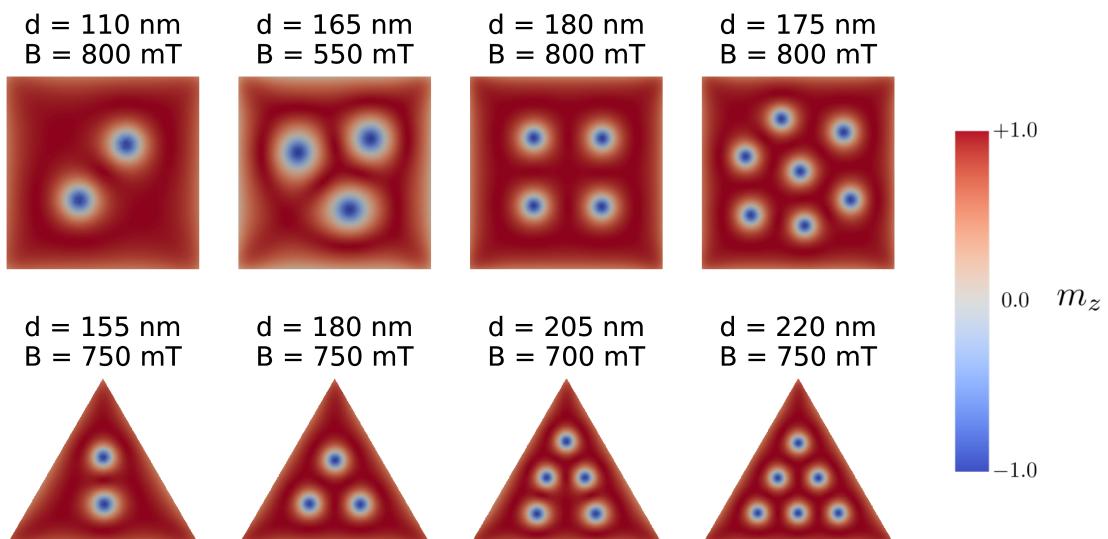


FIGURE 5.4: Examples of the z-component of the magnetisation for high energy states containing multiple skyrmions. For square systems, states containing up to 10 skyrmions were observed and in triangles, a maximum of 6 skyrmions were observed, in both cases with close packing of the Skyrmions.

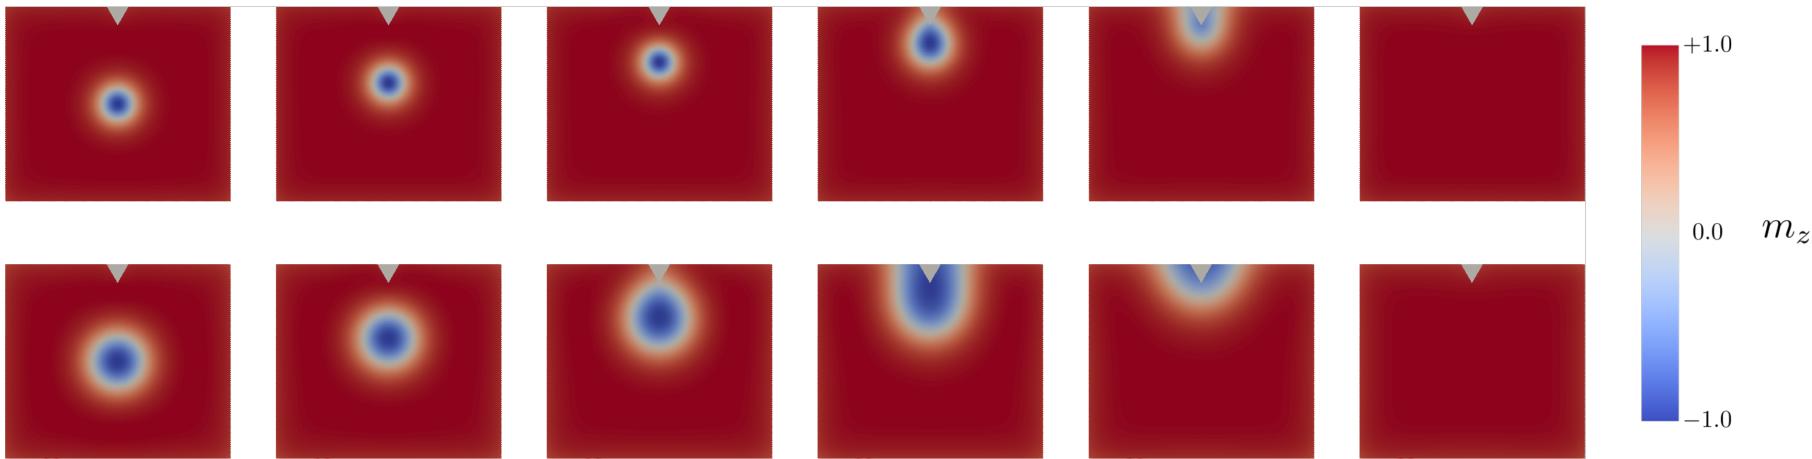


FIGURE 5.5: In this $d - B$ ground state phase diagram for the square geometry, the z-component of the magnetisation for the lowest energy state found for each sample size and applied field value is shown. There are three regions of interest. (a) The incomplete skyrmion state forms the bulk of the phase diagram. (b) For a narrow region and for low applied field values it can be seen that helical states form the ground state. (c) Isolated skyrmions form the ground state for large sample sizes. As the field is increased, the skyrmion shrinks such that more of the system aligns with the applied field, and for high field values the skyrmion no longer forms the ground state.

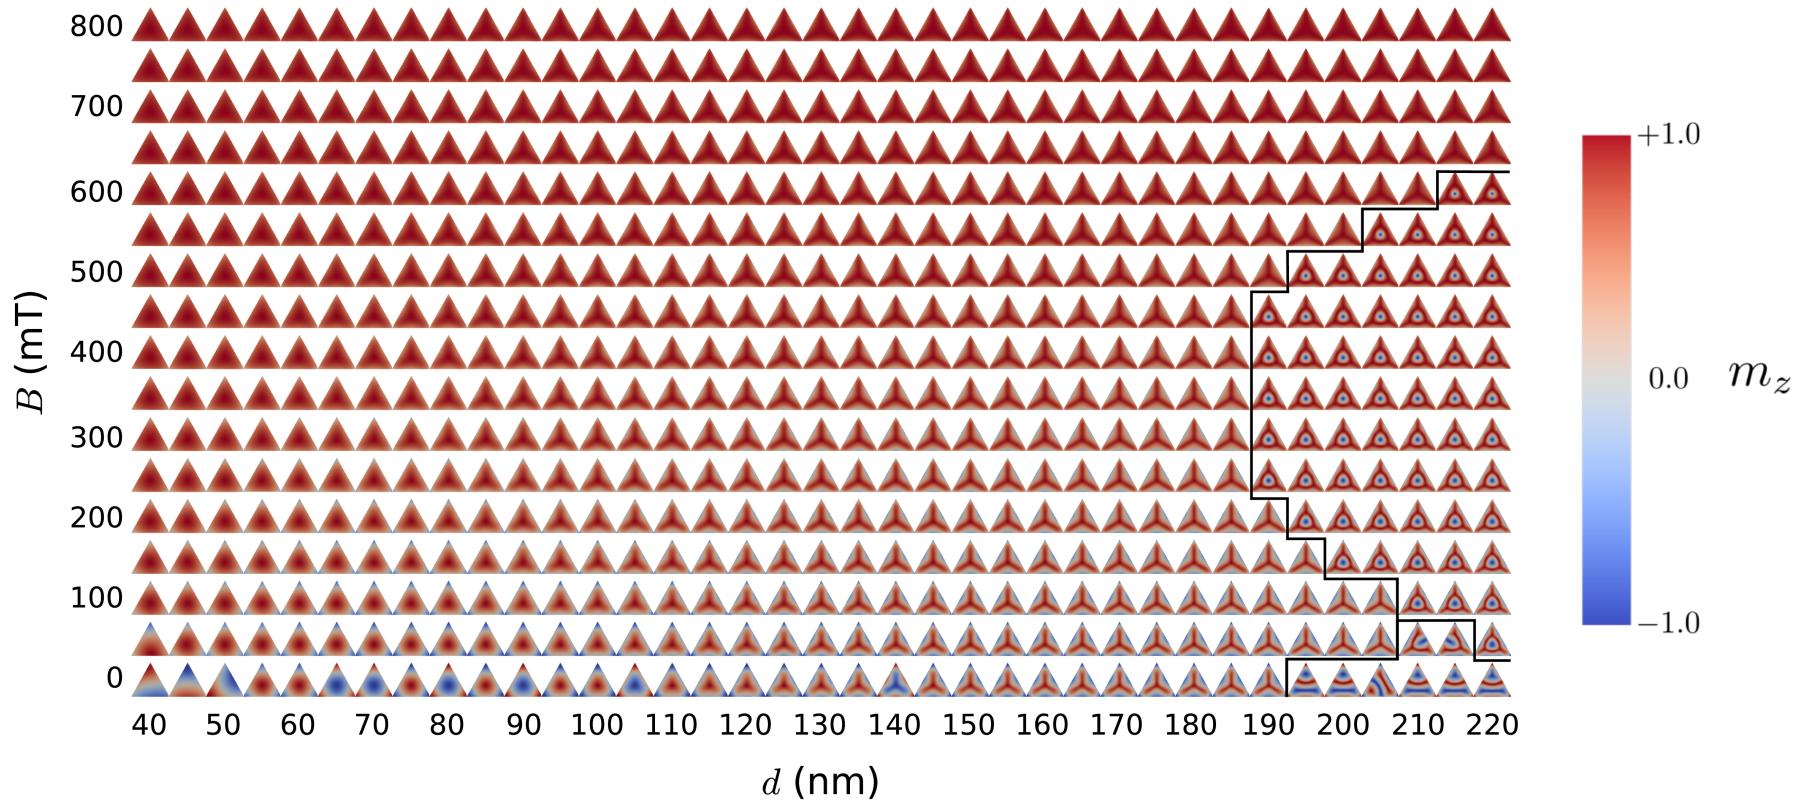


FIGURE 5.6: Here, the $d - B$ ground state phase diagram for the triangle geometry is shown, with the z -component of the magnetisation shown for each state obtained. As in Figure 5.5 there are three regions of interest. (a) The incomplete skyrmion state forms the bulk of the phase diagram. (b) For very large triangles, again at low applied field values, helical states form the ground state. (c) Skyrmions form the ground state for large sample sizes, but these are not observed without an applied field.

5.2.2.2 Ground States

The ground state phase diagrams (Figs. 5.5 and 5.6) show the lowest energy states identified for each geometry size for a given applied field. For the square systems, there is a large region where isolated skyrmions form the lowest energy state for sample sizes as low as 110 nm with an applied field of 350 mT. For larger sizes, the range of applied fields where skyrmions form the ground state increases, and at 155 nm, the skyrmion is the ground state with no applied field. For all sample sizes studied, fields of above 700 mT result in magnetisation saturation. With no applied field, from 100 nm to 150 nm there are several types of helical states form the ground state.

In Figure 5.6, the ground state phase diagram for the triangular systems is shown. In contrast to the square systems, skyrmions are not identified as the ground state when no applied field is applied for any sized sample investigated, which shows a strong indication that the shape of the boundary of the system plays a crucial role in the energetics of magnetic skyrmions in confined geometries. Skyrmion states do form the ground state for systems of side length $d > 185$ nm when an applied field between 50 and 600 mT is present. For systems of side length $d > 190$ nm, a number of helical states form the ground state with no field. Between 40 and 50 nm, quasi-helical type states are observed, though the lengths in these systems are below the helical length of FeGe.

These results are qualitatively similar to those seen in disk systems, in that the same types of state were found and the phase diagram is qualitatively similar. However in a previous wochiral bobberk, in nanodisks of FeGe of the same thickness, skyrmions formed the ground state with no applied field for smaller systems than in squares, with observation at disks of diameter greater than 135 nm [32], and no helices were seen as the ground state.

The incomplete skyrmion states identified in the triangular geometry vary significantly depending on the size of the systems. Notably, tilting of the magnetisation at the boundary of the sample due to the DM interaction causes the magnetisation to point most strongly along the axis of the applied applied field, with the strongest alignment along the axes of symmetry in both the square and triangular states, which can be seen in the incomplete skyrmion images shown in Figure 5.3 (i).

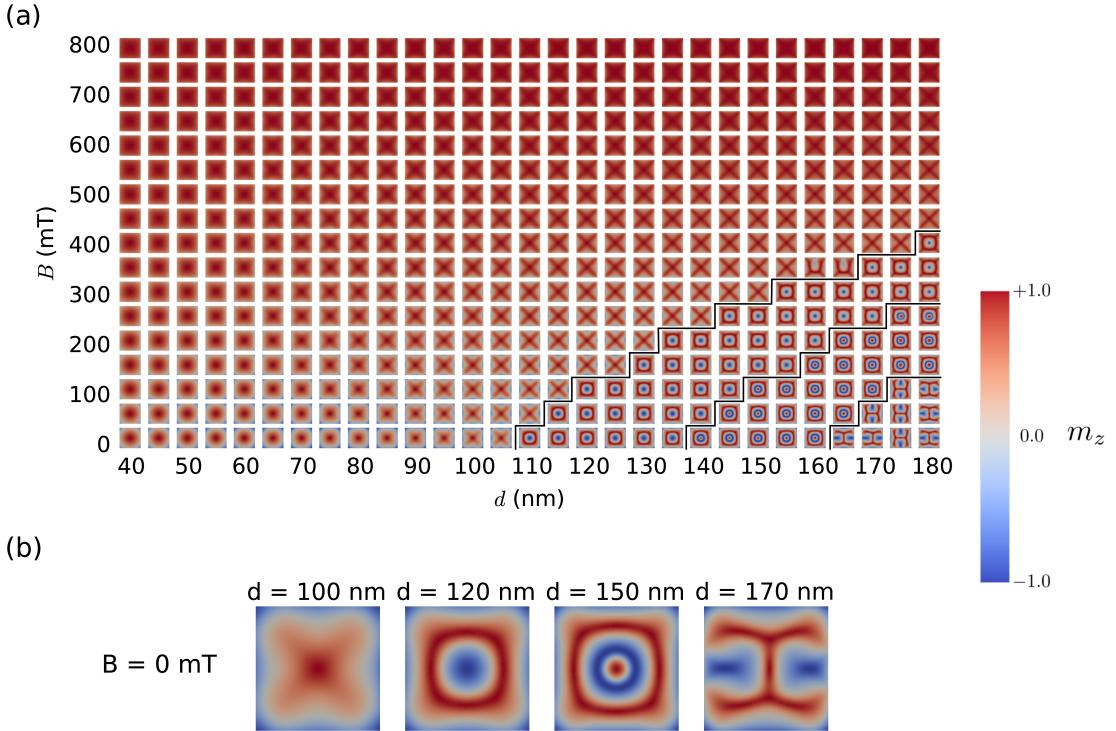


FIGURE 5.7: Obtained states from relaxing square systems from uniform magnetisation. In (a), there are four regions of different type of states - from left to right (i) incomplete skyrmions (ii) isolated skyrmions (iii) over-complete skyrmions, and (iv) helical type states. We predict that these configurations can be achieved in an experimental study where first a high saturation field is applied in the out-of-plane direction, and then the field is reduced to the value shown on the y -axis. In (b), the final state obtained for each of these configurations is shown, with no applied field.

5.2.2.3 Proposed experimental study

Of particular interest are states obtained from relaxing the systems from the uniform state. Experimentally, these states could be realised by initially applying a very strong applied field, to ensure that the magnetisation of a sample is saturated, and then rapidly reducing the applied magnetic field. The states obtained from doing this in the square sample are shown in Figure 5.7. There are four distinct bands of states, with incomplete skyrmions forming the bulk of the phase diagram. Skyrmion states are obtained in a narrow band, between 110 and 135 nm, with no applied field, and at larger sizes of system up to 180 nm with an applied field of 400 mT. For system sizes, from 140 to 160 nm with no applied field, target states are observed, and at 165 nm and above with no applied field, helical states form.

The corresponding uniform applied field results for triangular systems is shown in Figure 5.8. Here, similar results are obtained; in the bulk of the phase diagram incomplete skyrmion states are obtained. For large systems of between 190 and

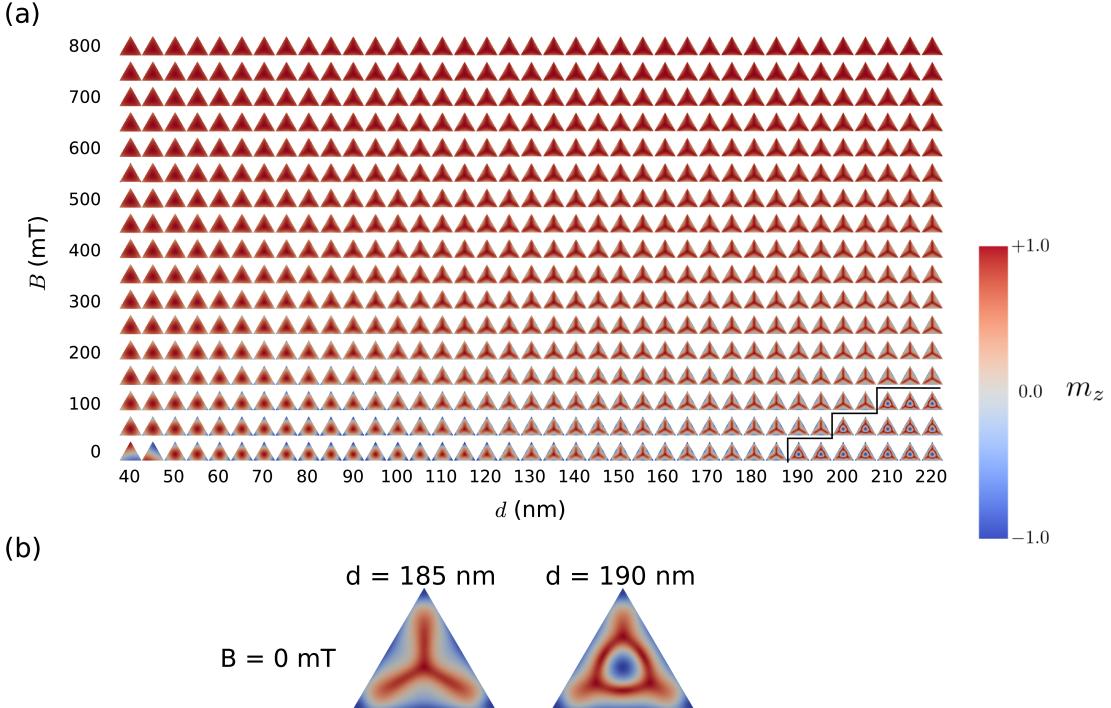


FIGURE 5.8: Here, the z-component of the magnetisation for the states obtained by relaxing triangular systems from uniform magnetisation is shown. In (a), there are two regions of different type of states - from left to right (i) incomplete skyrmions (ii) isolated skyrmions. In (b), the final state obtained for each of these configurations with no applied field is given.

220 nm, when the field is reduced to a value between 0 and 150 mT, a small band of skyrmion states is shown.

5.2.3 Conclusion

In this work, it is shown through micromagnetic simulations that in 10 nm thick confined geometries of FeGe, skyrmions can form the lowest energy state. When there is no applied field, there exists a lower bound of side length d between 150 and 155 nm, below which skyrmions do not form in square systems, and between 100 nm and 150 nm, a variety of helical type states form the ground state. In triangular systems, the incomplete skyrmion state forms the ground state in most of the phase space studied, and in large systems skyrmions form the ground state between fields of 50 and 600 mT.

Over the same range of sizes and fields studied, a wide variety of states are in equilibrium, and it is shown where these states can be obtained through metastability diagrams. The states obtained from relaxing uniformly magnetised states from the saturated state are shown in both the square and triangular systems in order

to motivate experimental work on FeGe confined geometries, and we predict that skyrmion states should be experimentally accessible in both square and triangular systems. In larger systems, we also predict that in large square systems, target states should be accessible using the same procedure.

To take work on confined FeGe nanosystems further, a potential option is to study elongated systems of disks and polygonal states, though these are much more computationally challenging. In one study, similar states to those found here were shown through hysteresis calculations in FeGe elongated nanocylinder systems up to around 50 nm via hysteresis. However, no phase diagram showing the ground states of such a system have been conducted to the authors knowledge. In such elongated systems, we may expect to see additional states than in this study due to the larger geoemtrical size, such as the exotic chiral bobber states [245] in large confined systems, which have recently been observed experimentally in thin film FeGe systems [246].

Chapter 6

Edge Defects and Energy Barriers in Skyrmion Monolayers

6.1 Introduction

Several techniques have been applied to investigate the energy barriers for transitioning between states in magnetic systems since the study became feasible in the late 1980s. The transition between states is most probable to occur over minimum energy paths - the paths in the energy landscape where the barrier between states is lowest, and so finding the barrier implicitly involves finding the minimum energy paths for a given system. The earliest computational technique to study the energy barriers in nanomagnetic systems was introduced by Enkin and Dunlop, [247] whereby a simple planar micromagnetic model of magnetite cubes was constrained at the surface to point in a given direction, with interpolating states minimised, in order to study single domain reversal mechanisms. Enkin and Williams then extended this reduced order formulation further to allow variation within a given plane [248]. Simple Stoner-Wohlfarth type models were later utilised to model the reversal properties of grains in magnetic recording media during the mid-to-late 1990s [249], while other studies attempted to estimate the upper limits of energy barriers by recording the energy during transitions when systems were evolved under Landau-Lifshitz-Gilbert dynamics [250].

The first attempt to directly compute the minimum energy paths, rather than assuming that the path followed in dynamic studies was minimal, was performed by Dittrich, et. al., which introduced to micromagnetics the Nudged Elastic Band Method (NEBM), a technique previously used to study structural transitions in

Chemistry, which can be utilised when the start and end states between a transition are known [133, 251, 252]. Numerical evaluation of energy barriers via a Lagrange multiplier method was later used by Paz, et. al. [253]. to study atomistic and micromagnetic systems, with the restriction that an initial rotation direction must be specified, favouring some annihilation mechanisms over others.

The study of energy barriers in systems containing skyrmions has been of recent interest, because of the various proposals for skyrmion based device technologies mentioned in the previous chapter. Crucial to any skyrmion-based device is the need for robustness; the skyrmion state should have a lifetime much longer than the information must be stored for. In general, for commercial storage applications, the lifetime of data must be at least 5 to 10 years, which is a challenging requirement. The skyrmion lifetime in a given sample is directly related to the energy barriers of different routes to destruction. For a particular minimal energy path, the lifetime can be estimated by an Arrhenius law:

$$\tau = \frac{1}{f} e^{-\frac{\Delta E}{k_B T}} \quad (6.1)$$

where ΔE is the barrier height of the minimal energy path, T is the temperature and f is the attempt frequency. Similarly, the same equation can be used to determine the spontaneous creation time with knowledge of the activation energy for the creation mechanism. The first study to consider minimum energy paths for skyrmion destruction and creation was performed by Bessarab, et. al. [131], who introduced a modified version of the NEBM on an atomistic system, utilising geodesic distances, and studied the energy barrier between the ferromagnetic and skyrmion state via skyrmion collapse due to a large applied field. Further studies have expanded on this by studying other destruction mechanisms, [51, 134, 254] namely mediated with a Bloch Point and through the edges of a sample, looking at the effect of destruction through the boundary [132]. Authors have also considered the system size, applied field, and confinement of skyrmions, as well as the destruction mechanisms for anti-skyrmions [255]. Notably, because of the Bloch Point energy destruction mechanism, it is important that that atomistic rather than micromagnetic modelling is used such that the energy barrier is computed accurately. In the above studies, the energy barrier for skyrmion destruction at the boundary has been found to be substantially lower than the other destruction mechanisms. Evidence that this is the case in reality is that skyrmions have been observed to preferentially form at the edges of systems in experiments. With knowledge of the energy barriers of different destruction and creation routes, one exciting recent study used the techniques as kinetic Monte Carlo and forward flux

sampling to estimate the lifetimes of a given state [256].

It is well known that defects aid in the nucleation and destruction of skyrmions, with simulation studies showing that nucleation and destruction can occur at notches in confined systems [37, 44], and in experimental studies the nucleation of skyrmions at defects has also been observed [257]. Thus far, however, only a few studies have looked at the effect of defects on the minimum energy path for skyrmion creation and destruction mediated by defects. In one study, [258] the energy barrier was computed for in-track destruction on a Pd/Fe/Ir(111) system at a defect site with either one or three atomic vacancies, and it was found that the energy barrier was reduced by more than 50% at the site of the three atom defect. A more comprehensive study looked at the effect of defects which vary the local Exchange, Dzyaloshinskii-Moriya and Anisotropy strengths across sites in Co/Pt(111) systems [259].

In realistic systems, we expect that defects are present not only within the sample but also at the surface edge, as a result of fabrication. In lamellae cut from crystals, edge defects can result from ion implantation, resulting in magnetically dead surface regions, while images of nanostructures grown on a substrate also clearly show rough surfaces. A recent study looked at energy barriers in confined nano-islands of Pd/Fe and Pd₂Fe in which clear edge defects can be seen, but did not model systematically how the edge defects affect the energy barriers of skyrmion to ferromagnetic transitions [135]. In another work, the energy barriers to skyrmion destruction past at semicircular notches along a nanotrack edge are considered micromagnetically when an applied current is present in the system, but the effect of defect size was not considered [44].

In this work we consider a simple, atomistic system comprising of a monolayer of atoms, with edge defects of triangular, square and Bezier curve shape. We aim here to try and fill the gap in the energy barrier literature by calculating the minimum energy paths to skyrmion nucleation and destruction via such edge defects, using the Nudged Elastic Band Method. We additionally study the contribution of the dipolar field to this barrier, and show that it can significantly reduce the energy barrier for large defects. The present author conceived of this study, wrote and performed the simulations and interpreted the results, and discussed the results with Dr. David Cortés-Ortuño. The data from this study has been made available [260].

6.2 System Definition

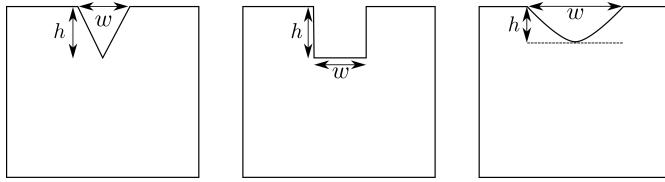


FIGURE 6.1: Definitions of the three types of defect - triangle and rectangular and bezier.

We study using the atomistic component of Fidimag [95] a confined rectangular monolayer of 201 x 200 Cobalt atoms using the Heisenberg energy model:

$$\mathcal{H} = - \sum_{i \neq j} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j + d \hat{\mathbf{r}}_{ij} \times \hat{z} - \sum_i k_u (\mathbf{S}_i \cdot \mathbf{u})^2 + \sum_i \mu_i \mathbf{H}_d \cdot \mathbf{S}_i \quad (6.2)$$

where the parameters used are those given in Table 6.1. For each system, we consider an edge defect cut from the top edge of the system. We introduce two types of edge defects in these simulations; triangular, rectangular and curved, (Figure 6.1), each with a maximum depth d and width w . For all defects, magnetic atoms are removed when they are centred within the region bounded by the defect region. The curved defects are introduced into the system by a bounded Bezier curve defined as:

$$f(t) = (1-t)^2 \mathbf{P}_0 + 2t(1-t) \mathbf{P}_1 + t^2 \mathbf{P}_2 \quad (6.3)$$

where \mathbf{P} are control points bounded such that $\mathbf{P}_0 = (\frac{L_x-w}{2}, L_y)$, $\mathbf{P}_1 = (\frac{L_x}{2}, L_y - 2d)$ and $\mathbf{P}_2 = (\frac{L_x+w}{2}, L_y)$ and $t = [0, 1]$.

Atomistic Cobalt Parameters	
a	0.25 Å
μ	0.846 meV
J_{ij}	27.026 meV
k_u	0.676 meV
d_{ij}	27.026 meV

TABLE 6.1: The parameter set used for all simulations in this chapter. The parameter set is the same as that used for a Cobalt monolayer in [132].

We apply the NEBM algorithm as described in section 3.4.2 of this thesis. We perform interpolation of the spin vectors to find the starting points for each image in the NEBM band. We create a skyrmion magnetisation texture using the ansatz of a skyrmion located at the centre of the sample, where within a radius R the

magnetisation is set to:

$$m_r = -\cos(\pi r/R) \quad (6.4)$$

$$m_\theta = \sin(\pi r/R) \quad (6.5)$$

for a radius $R = 10$ nm, and set this state to be the \mathbf{Y}_0 image. For the \mathbf{Y}_n image, we set the magnetisation to the $+z$. We then use the steepest descent energy minimisation algorithm implemented in Fidimag [130] to find local energy minima for the two initial states, which are then used as the starting points for the NEBM. In order to construct the other, non-fixed images in the band, we start with the equilibrium ferromagnetic state, and then select a circular region of radius R (such that R is greater than the skyrmion radius). The circular region for a given region is has its origin shifted by some distance $i(L_y + R)/n$ where i is the image number and n is the number of images. Then, within this region, the relaxed skyrmion state is interpolated onto the equilibrium state. This gives a band of images in which the Skyrmion is moving towards the edge of the system. We note that none of these states are relaxed at initialisation of the band (i.e. spins in each system would feel a torque due to the effective field). This type of interpolation was used to find the edge nucleation energy path of interest, rather than other potential paths described in other works [132].

In all systems studied here, the relaxed skyrmion state is of higher energy than the relaxed ferromagnetic initial state, and hence the skyrmion is only metastable, with the method considered to have converged to a metastable state once the magnetisation has stopped changing with respect to the fictitious time τ as $|\partial\boldsymbol{\mu}/\partial\tau| < 10^{-6}$. For these systems, we choose in all cases defects small enough that the initial skyrmion magnetisation is not affected.

The iterative NEBM algorithm is then applied. In all cases, a spring constant $k = 10$ is used, and a total of 26 images are used (including the two fixed images). We consider the band to have relaxed when the maximum rate of change for images in the band reaches a tolerance of $|\partial\mathbf{Y}| < 10^{-6}$. Initial simulations showed that the result did not show dependence on the spring constant k over the range of values chosen, and this value was used because it gave a good rate of convergence.

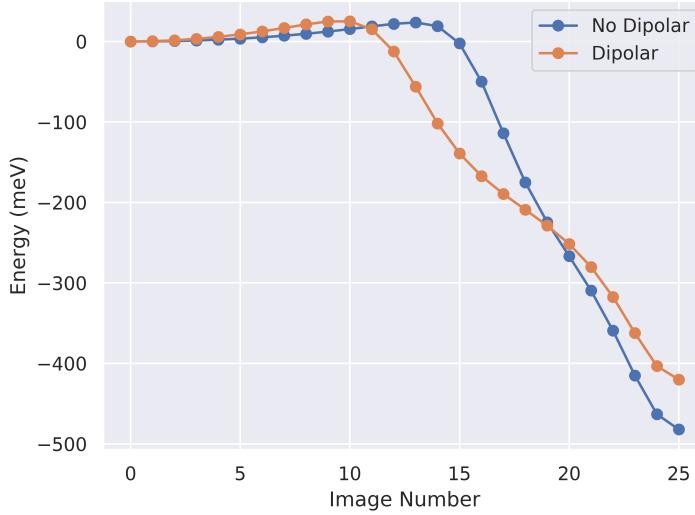


FIGURE 6.2: We show here the minimum energy path found for the system when no defect is present, both with and without the dipolar interaction.

From the relaxed band we can compute the skyrmion destruction and creation activation energy barriers heights as:

$$\Delta E_{\text{Destruction}} = \max(E_i) - E_0 \quad (6.6)$$

$$\Delta E_{\text{Creation}} = \max(E_i) - E_n \quad (6.7)$$

where E_i is the energy of the image i in the relaxed energy band, and E_0 and E_n are therefore the skyrmion state energy and the ferromagnetic state energy respectively.

6.3 Results

6.3.1 Triangular Defects

In a Cobalt monolayer, atoms sit in a regular triangular lattice. Because of this, triangular defects are a natural playground for considering the effect of the size of the defect on the skyrmion destruction energy barrier when the defect is smoothly cut from the lattice, because when $h = \sqrt{3}/2w$ no rough edges are present, and the ‘sharpness’ of the defect is fixed; any effect of the defect on the energy barrier is therefore only an effect of the defect size and not the shape, which is essentially fixed. Initially, a reference calculation is performed for a pristine system containing

no defect, both with and without the dipolar interaction included. For this system the energy bands are shown in Figure 6.2. We can see that the path followed by the system with the dipolar interaction is slightly different; the skyrmion is destroyed in an earlier image in the dipolar system. We find that this is due to the increase in the skyrmion size, which is consistent with the results of other publications [132]; because of this the skyrmion makes contact with the boundary in an earlier image in the series. For this system, the energy barrier to destruction of the skyrmion is found to be 23.57 meV with no dipolar interaction, and 24.93 meV, while the creation activation energies are found to be 505.67 meV and 445.12 meV. The activation energy barrier for this destruction mechanism of skyrmions in Cobalt is therefore substantially lower than that for the corresponding creation mechanism.

In the Cobalt system, the skyrmion state is highly localised, and the total energy of the state is not significantly affected by the boundary of the sample. Thus, for sufficiently small defects in the system, the energy barrier is not affected by the displacement of the skyrmion towards or away from the defect. This was verified by checking that the energy difference for states in the band was negligibly small as the Skyrmion moved towards the defect for images in the band.

We now compare this result to systems containing a triangular edge defect. The defect width w is varied between 0.5 and 7.5 nm, and the defect height correspondingly varied as described above. We find that regardless of the presence of the defect and its size, the skyrmion essentially follows the same annihilation mechanism through the boundary as in the pristine system. The energy barrier heights found for this system are shown in Figure 6.3. We find that the energy barrier to skyrmion destruction is substantially reduced in both the dipolar and non-dipolar cases with increasing defect size. While in the pristine case, the energy barrier is higher with the dipolar interaction, we see a cross over, whereby for large defects ($w > 6.5$ nm), the barrier is lower in the dipolar case than in the non-dipolar case. We show in Figure 6.4 the overall skyrmion destruction process for the defect of width 7.0 nm, for the non-dipolar and dipolar cases.

In order to determine why exactly the energy barrier falls when a defect is introduced, we must consider the changes in the contributions of the individual energy terms across the band of images. In Figures 6.5 and 6.6 we show the individual spin orientations close to the boundary, and in Figure 6.7 we consider the energy changes between the initial skyrmion state, and the ferromagnetic state, in the presence of the Exchange, Dzyaloshinskii-Moriya and Anisotropy energies only,

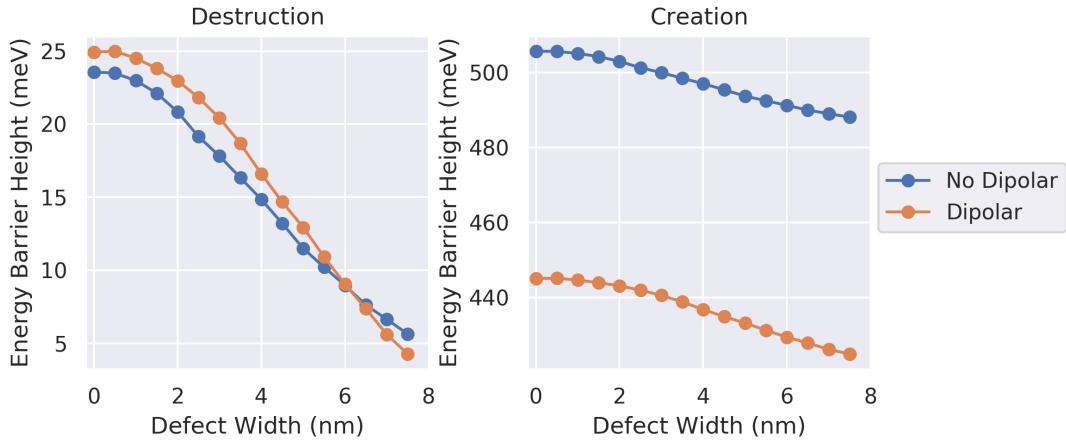


FIGURE 6.3: Energy barrier height for systems containing triangle defects. We include the no-defect case for reference. We find that the activation energy to create a Skyrmion is substantially larger than the activation energy for destruction at a defect.

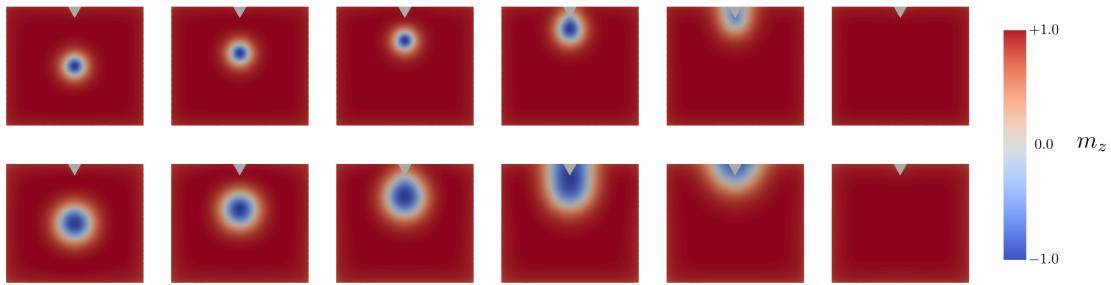


FIGURE 6.4: Path to destruction for a triangular defect of width $w = 7.0$ nm and height $h = \sqrt{3}/2w$, without (above) and with (below) the dipolar interaction present. We can clearly see that the skyrmion is larger in the dipolar system, but in both cases we see that the destruction/creation mechanism is similar. A video of this transition is available [261].

and with no defect present in Figure 6.7 (a)). In the skyrmion state at the beginning of the band, clearly, the anisotropy energy is not minimised because there are spins rotated relative to the z axis, forming the Neél wall which makes up the Skyrmion boundary. In both the ferromagnetic background and the skyrmion core, the magnetisation is varying slowly; the Exchange energy is therefore low in these collinear regions, and the Dzyaloshinskii-Moriya energy high. In the ferromagnetic state, the both the Exchange and Anisotropy energies are minimised because the magnetisation is aligned almost uniformly, but the Dzyaloshinskii-Moriya energy is high for the same reason. The barrier height between the skyrmion and ferromagnetic state then, is to a first approximation a result of the competition between the increase in Dzyaloshinskii-Moriya energy, and the decrease in the Exchange and Anisotropy energies.

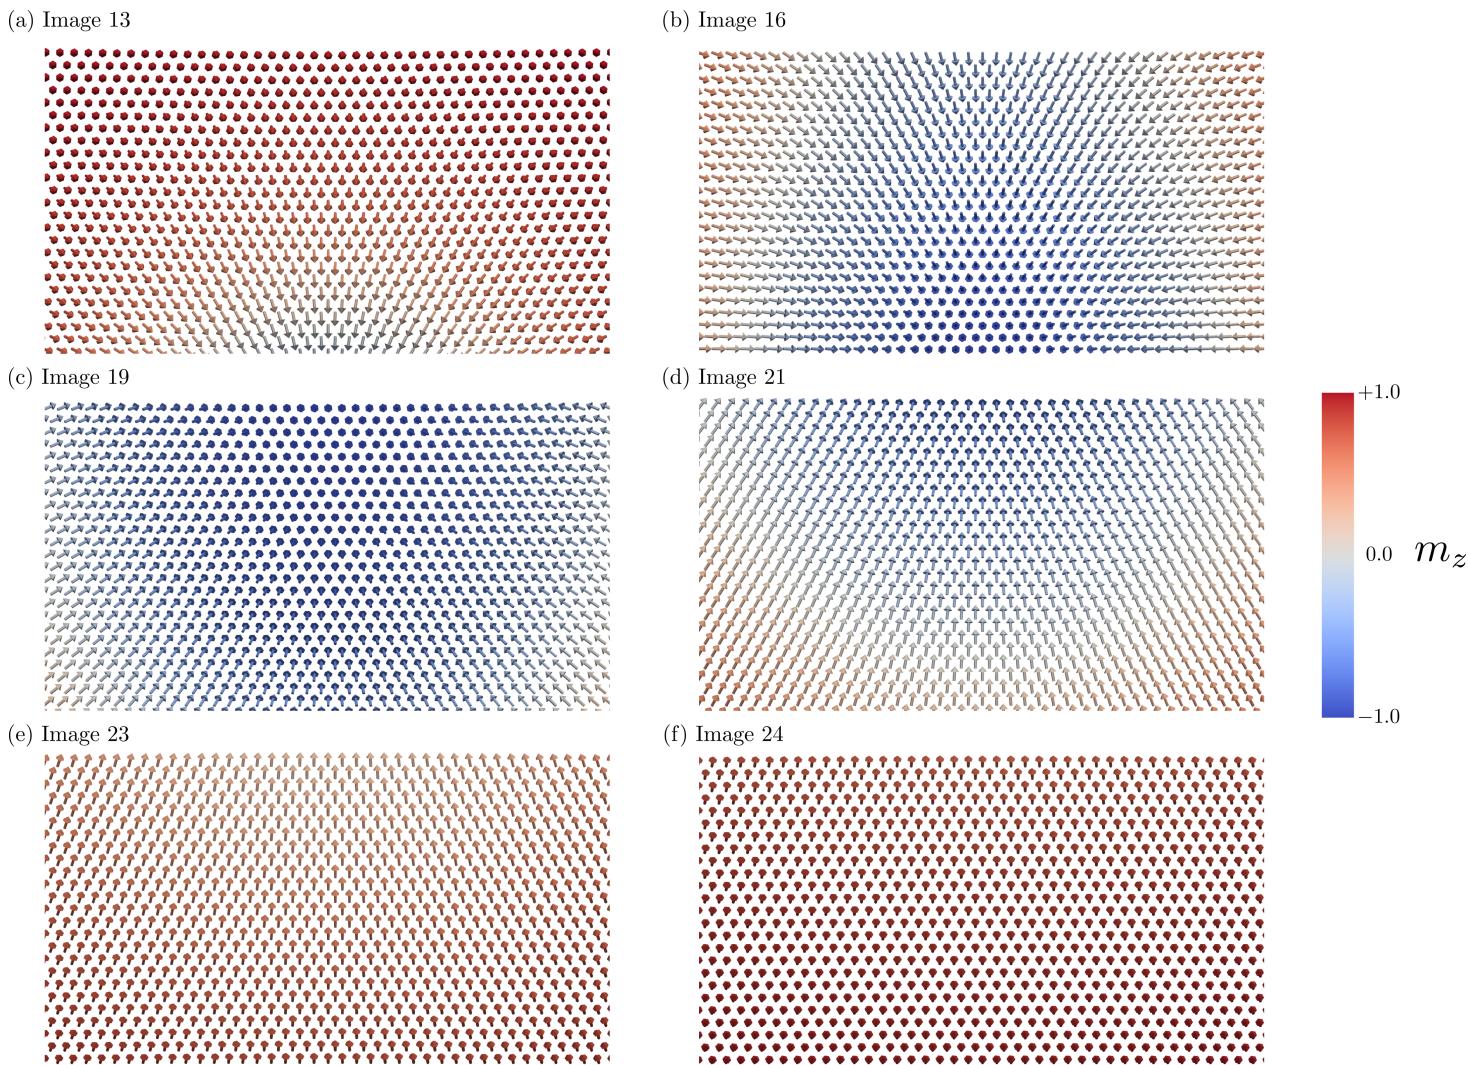


FIGURE 6.5: Close-up view of the no-defect case of a Skyrmion annihilation.

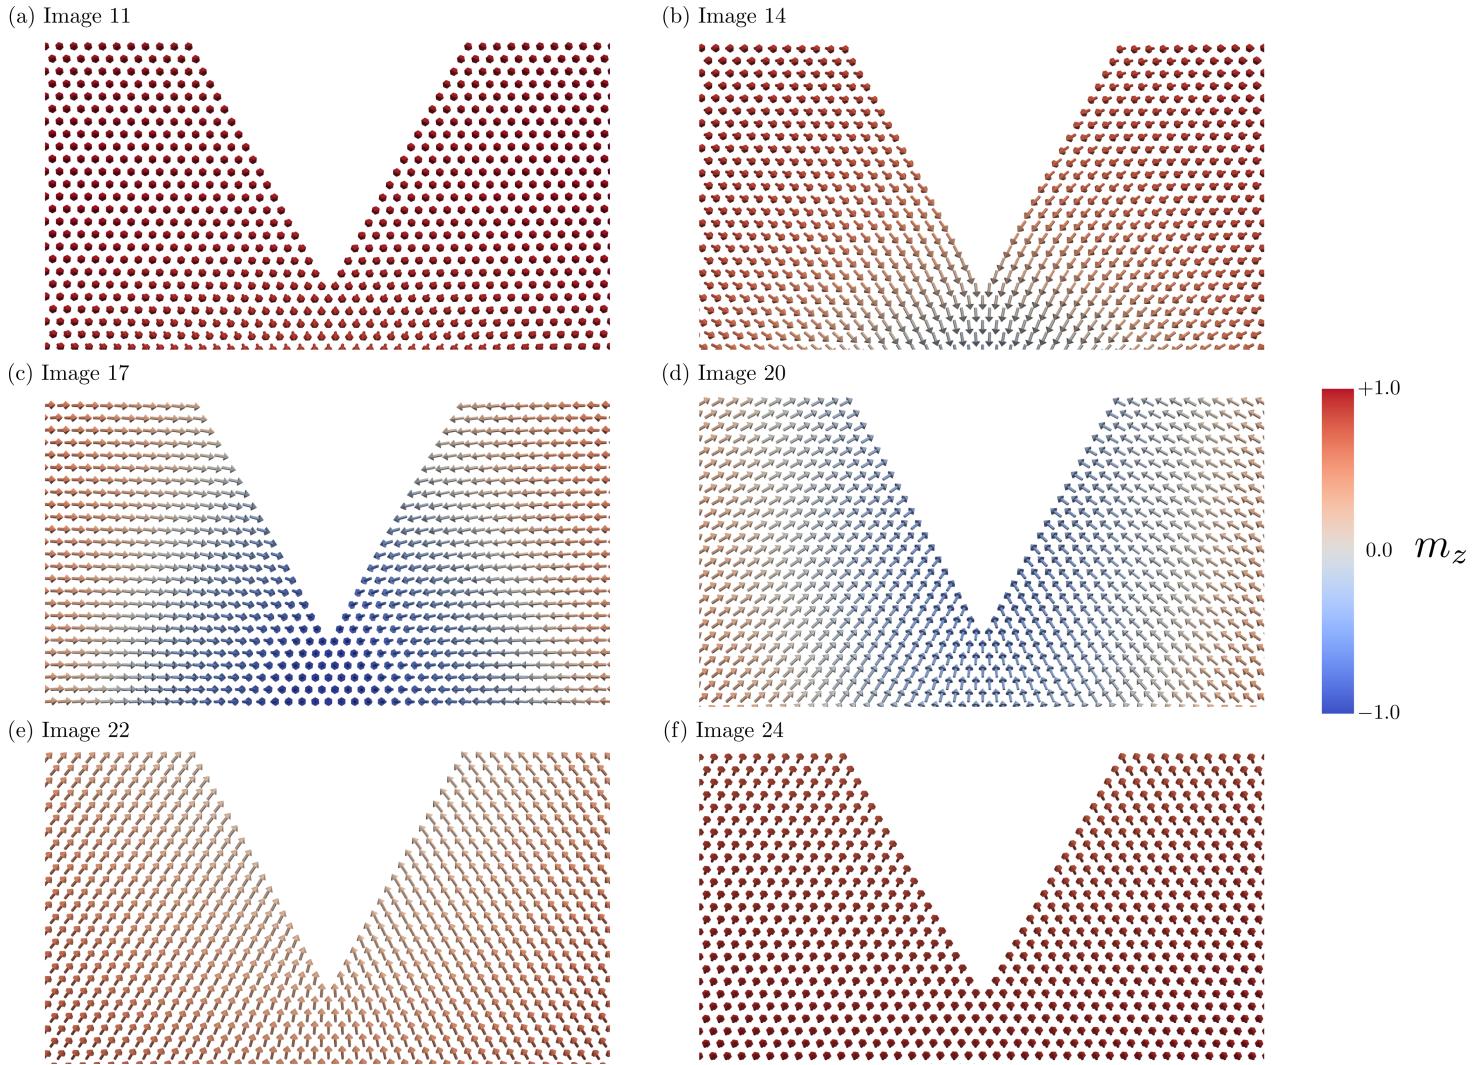


FIGURE 6.6: Close-up view of the Skyrmion annihilation at the edge of a system with a $w = 7.0$ nm triangular defect.

Consider now then, what must occur for the skyrmion to annihilate. When the skyrmion magnetisation texture approaches the boundary of the sample, the spins along the skyrmion path must undergo a full rotation, forcing them from an initial alignment in the $+z$ direction (Figure 6.5 (a)) to pointing into the sample (Figure 6.5 (b)), and then into the $-z$ direction (Figure 6.5 (c)). This process happens again in reverse as what remains of the skyrmion texture moves out of the domain; first the spins point rotate to point along the y -axis (Figure 6.5 (d)), In terms of the uniaxial anisotropy, rotating away from preferred axis is costly, and in particular, rotating through a full circle is the most costly possible evolution. This coherent rotation of nearby spins observed through the evolution of the system is both energetically expensive for the Dzyaloshinskii-Moriya interaction and cheap for the Exchange energy.

What happens then, on the introduction of a defect? We can see the evolution of the energy bands in Figure 6.7(b), for the triangular system with width 7.0 nm and it is clear that as the Skyrmion approaches the defect, the sharp triangular defect substantially reduces the energy barrier contribution from the Dzyaloshinskii-Moriya interaction while affecting little the energy of the Exchange interaction. We can attribute this to the reduction in part in the number of spins which must coherently rotate in any given point in the transition, due to the physical size of the defect. In addition, for all spins along the boundary of the system (both when a defect is and is not present), there are a reduced number of nearest neighbours, and the effective field acting on such spins from the Dzyaloshinskii-Moriya and Exchange interactions is therefore weaker than those located away from the edges of the system. The magnetisation inherently wants to align along the effective field so that the energy is reduced, but because the effective field is generally weaker at the boundary due to each spin having fewer neighbours, each spin moment is less strongly held in place. For the Exchange and Dzyaloshinskii-Moriya interactions, which are the strongest contributors to the effective field, the effective fields are a product only of the local environment of each spin, and so a reduction in the number of nearest neighbours for a particular spin substantially reduces the effective field acting on it. Introducing an edge defect then, increases the number of weakly held atoms which the skyrmion magnetisation texture comes into contact with in the annihilation process. Because of this, the energy barrier to push the skyrmion out of the system is reduced.

This still leaves the effect of the dipolar field to explain - why exactly does the energy barrier increase with no defect, and fall below that of the corresponding system with no dipolar field for larger defects? We know that the dipolar field

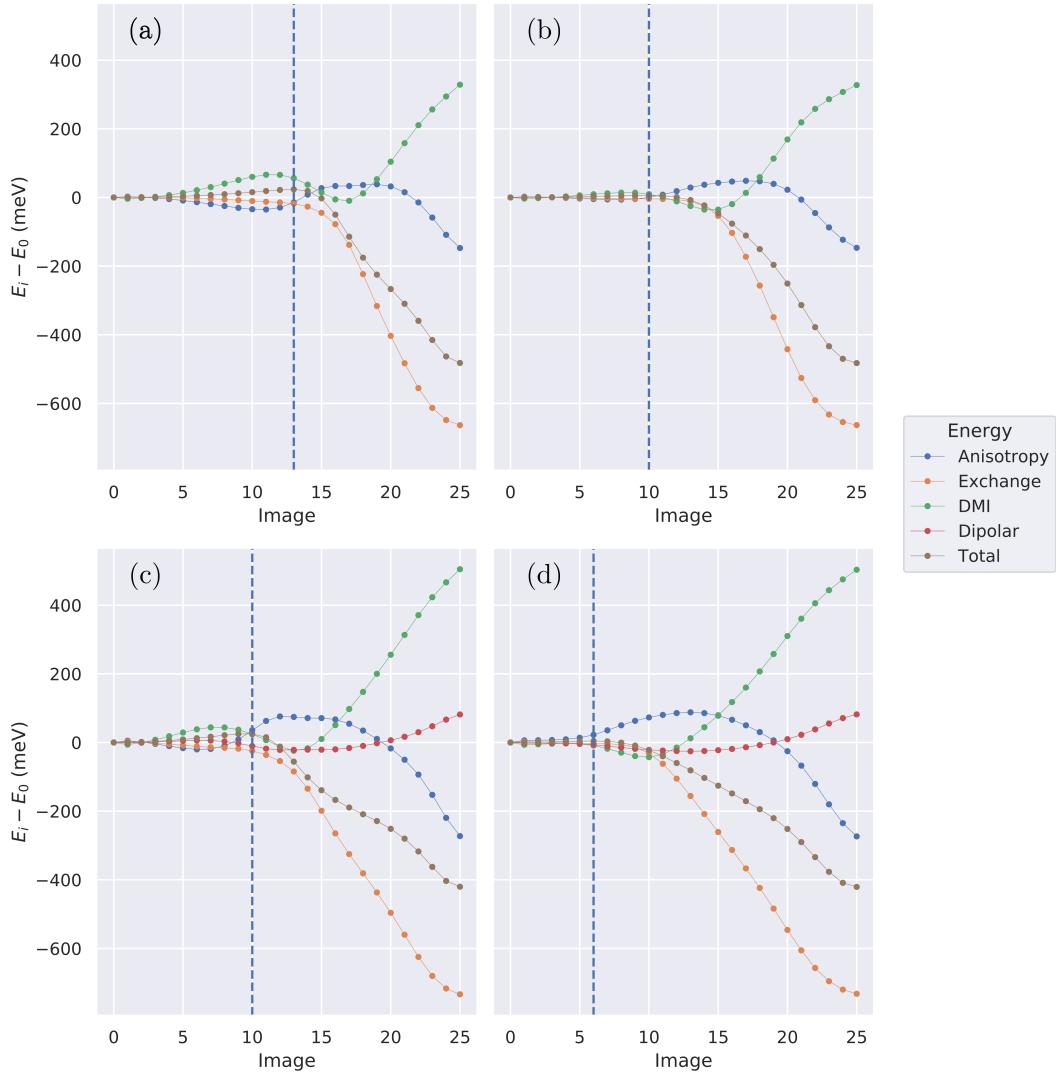


FIGURE 6.7: Change of each energy term from the initial energy along the images in the NEBM band. In (a), no defect is present and in (b) a triangular defect of 7.0 nm width is introduced. In (c) and (d), we show the change in energies for the same systems, but with the addition of the dipolar field. We show the highest energy point along the band with a vertical line.

is substantially weaker than the Exchange and Dzyaloshinskii-Moriya interactions over short length scales. When a defect is small, the reduction of the dipolar field due to the defect is very small indeed in comparison with the other effective field contributions acting on each spin. We can see from Figure 6.7(c) and (d) that indeed, the dipolar energy changes little across the band, while the other energies have indeed changed much more substantially. We can therefore conclude that the introduction of the dipolar field causes the energy barrier to change due to its affect on the size of the skyrmion rather than the difference between the dipolar field energies in the initial and final states. This means that a greater number

of spins are forced to rotate together in the no-defect case, increasing the energy barrier, and fewer spins are required to in the large defect case, decreasing the barrier.

6.3.2 Rectangular and Bezier Defects

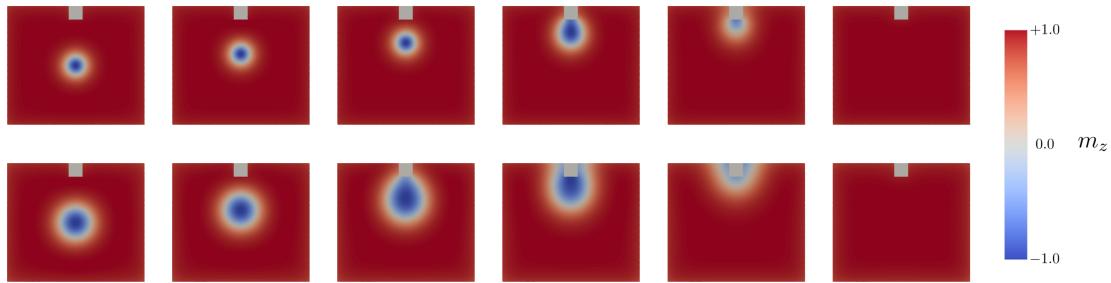


FIGURE 6.8: Path to destruction for a rectangular defect of width $w = 5.0$ nm and height $h = 5.0$ nm, with (above) and without (below) dipolar interaction.

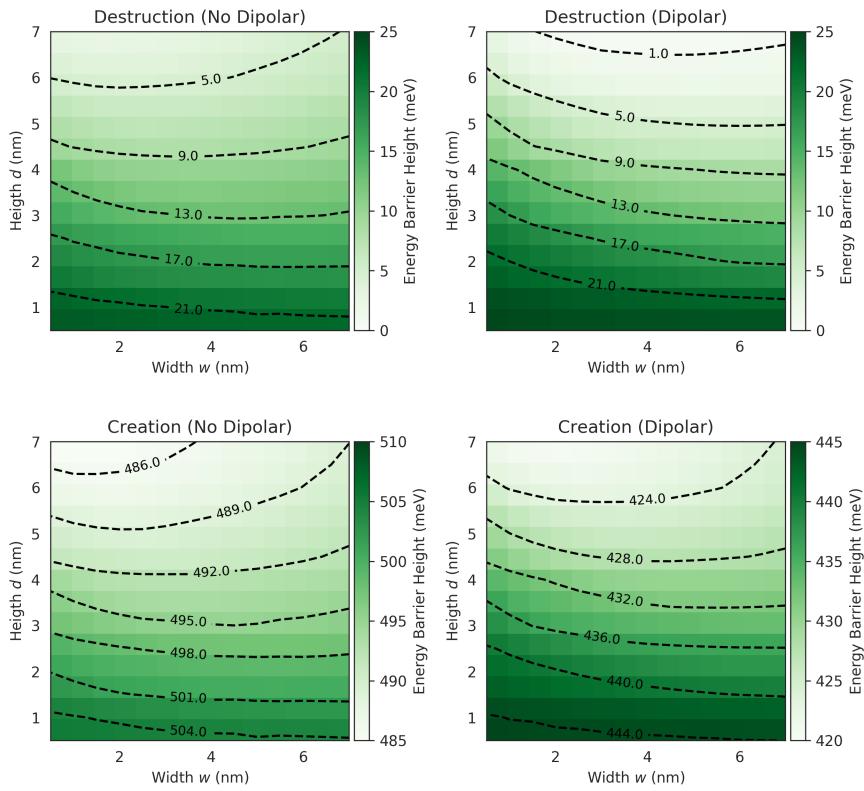


FIGURE 6.9: Rectangular defect barrier heights for creation and destruction of the skyrmion.

We now consider rectangular and Bezier defects, and for these systems both the width and height of the defect are freely varied, in order to study the combination

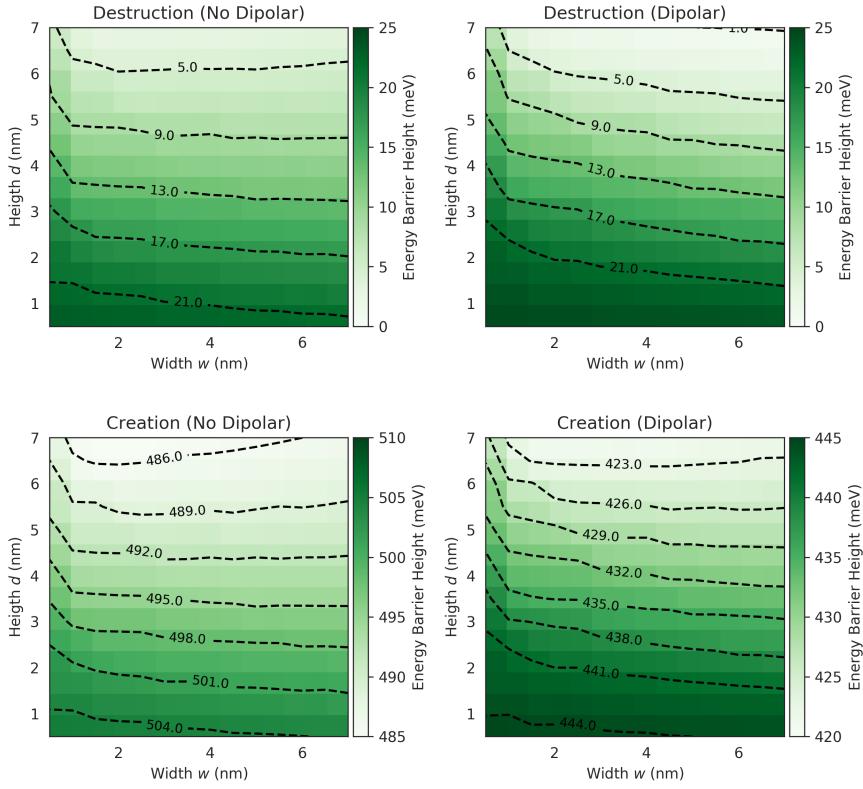


FIGURE 6.10: Bezier defect barrier heights for creation and destruction of the skyrmion.

of the effect of both the defect shape and size on the energy barrier. The energy barrier heights from these rectangular and Bezier defect simulations are shown in Figures 6.9 and 6.10 respectively. These show the activation energy barrier height for both the dipolar and non-dipolar systems and for the destruction and creation mechanisms.

In common with the triangular system, we find that the size of the defect affects the barrier heights, and that for small defects the barriers are not substantially affected. In the range of defect size studied, we find that the depth affects the energy barrier height substantially more than width. For the non-dipolar rectangular case, it is found that the minimum energy barrier found for destruction is nearly an order of magnitude smaller than in the pristine case; for a defect of width $w = 4.0$ nm and height $h = 7.0$ nm, with an energy of 0.280 meV (1.12% of the pristine case). For the dipolar case, the smallest energy barrier is found with a much narrower defect of size $w = 1.5$ nm and $h = 7.0$ nm. In the Bezier systems, we see very similar behaviour to the rectangular case, and so we focus on this.

There are several things of note in the results shown here. In the non-dipolar

rectangular system, for a defect depth of $h > 1.5$ nm, the destruction barrier height is found to start to slowly increase with the width w above intermediate sized defects, the point at which this occurs varying for each depth. We can explain this behaviour by considering what would happen if we continued increasing w beyond what we have shown here - it is clear that from the perspective of a skyrmion, a wide defect appears like a flat surface - i.e. the system is simply equivalent to one with fewer rows of atoms. Thus, given that we know the limiting behaviour, and that a defect can reduce the energy barrier as argued above, it is obvious that the energy barrier height must start to increase beyond a certain width of defect. We note that this behaviour is much more substantial with no dipolar field present; the same behaviour here does occur with a dipolar interaction, but only when the defect is much deeper as $h > 4.0$ nm, and with the inflection point occurring for much wider defects.

In accordance with the results for the skyrmion destruction mechanism, we find that the creation energy barrier also drops with the introduction of a defect. In this case, with increasing depth, we see that the reduction in the barrier is substantial for larger systems, though the activation energy for creation in this case is still substantial. We find here that again, with increasing defect width, the energy barrier to creation inflects, leading to an increase in the barrier height with increasing defect width.

6.4 Conclusion

In conclusion, we have seen through the systematic study of a cobalt nanolayer with the nudged elastic band method that edge defects act to catalyse the destruction of skyrmions. We find that the size and shape of such defects substantially affects the energy barrier, with larger defects affecting the barrier more than small defects. We seek to provide an explanation of the mechanism behind such behaviour as a combination of the reduction in the Dzyaloshinskii-Moriya interaction and reduction in the effective field for atoms on the boundary, considering that a larger defect leads to a greater number of spins on the edge of the system. These results explain further the behaviour of skyrmion formation and destruction at system edges in thin-film systems.

The results also have an important implication for simulation studies of thin films skyrmions when edge defects and notches are present in the system. For the non-dipolar case, it is found that the minimum energy barrier size for square defects

when no dipolar interaction is present is nearly an order of magnitude larger than the dipolar case, and that the barrier heights are substantially affected for anything other than small sizes. This has important implications for the simulations of skyrmion motion around defects. Many authors make the argument that the dipolar field is negligible in thin systems, and therefore neglect it from spin dynamics calculations involving Skyrmions. Our results show that for an accurate picture of skyrmion motion and destruction in spin dynamics simulations, one must consider the dipolar contribution in order to correctly model the energy landscape around defects in such systems.

Though we have not included direct contributions to the magnetic energy by the defects themselves, we can consider how these might affect the results. Generally, uniaxial anisotropies manifest in a shrinking of the skyrmion, because it is less energetically favourable to align against this axis. In magnetostrictive materials, stress can induce additional anisotropy contributions [262]. Around defects, therefore, we may expect additional local contributions to the anisotropy energy, which would have the effect of small shifts in the relative energy contributions, and thus affect the energy barrier.

Chapter 7

Bloch Points in Chiral Systems

7.1 Background

In micromagnetics, a constant saturation magnetisation M_s is assumed everywhere within a ferromagnetic material, with the magnetisation continuous and differentiable. However, in 1965, Feldtkeller [263, 264] posited that for a cylinder, with the magnetisation at each end pointing in the direction of the surface normal, only through the introduction of a point singularity, known as a Bloch point (BP) could the internal magnetisation structure be described. On a sphere surrounding such a singularity, the magnetisation vectors cover the surface once, even as the sphere radius tends to zero. The result of this is that around this point, the magnetisation vector is not defined, and the micromagnetic assumption of constant M_s is broken. In the same work, Feldtkeller described how through the introduction

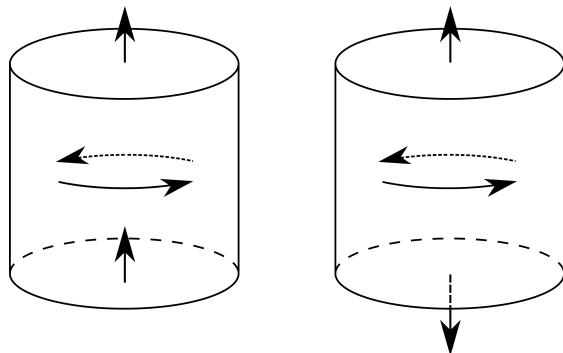


FIGURE 7.1: Feldtkeller's idea [263] leads to the necessity of a singularity in the magnetisation. For the cylinder on the left, the external observed magnetisation can be met with a continuous internal distribution. For the cylinder on the right, it is only through the introduction of a point singularity that surface state can exist.

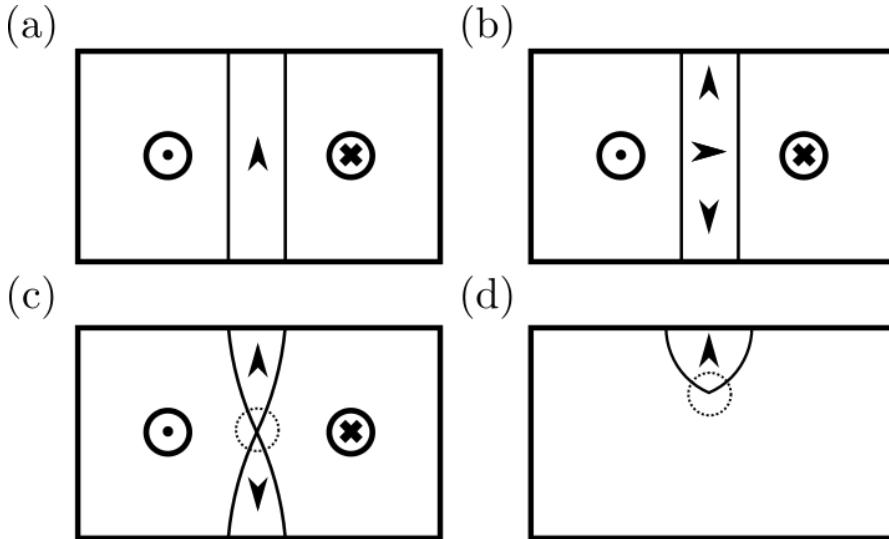


FIGURE 7.2: Here are shown four types of magnetic state which can be observed in ferromagnetic materials. A Bloch Wall is shown in (a); these are formed through competition between the magnetostatic and exchange interactions. In (b), a Bloch Line is shown, and in (c) a Bloch Point within a Bloch Line. These states can form to reduce the energy of a Bloch Wall when it is sufficiently large. In (d) is shown a Bobber, a recent state observed in magnetic materials with a Dzyaloshinskii-Moriya Interaction.

of such point singularities reversal mechanisms which are discontinuous could be admitted. Subsequent analytic work on Bloch points was carried out by Döring [265], who calculated energy bounds on the introduction of a Bloch point, and by Hubert [266], who showed that the introduction of a Bloch point into a Bloch line could reduce the magnetostatic energy in sufficiently thick films.

Slonczewski made major contributions to the study of Bloch Points, and showed that the formation of Bloch point pairs in bubble domains was feasible [267], computing the energy barrier which must be overcome in order to do so and later computing the mobility of a Bloch point along a Bloch line [268]. Subsequently, higher than expected magnetic bubble resonances were explained via contributions from Bloch line and Bloch point contributions in garnet films [269], and today, bubbles containing Bloch Lines are widely known as Type-II bubbles. It was shown by Arrott [270] that by the introduction of two Bloch points at the surface of a ferromagnetic cylinder aligned along the long axis, the magnetisation can undergo reversal, and subsequently this idea has been utilised by other authors to explain reversal mechanisms of other states.

The computational study of Bloch Points started in the late 1980s, with Nakatani and Hayashi calculating the magnetisation distribution around both a Bloch Line

and a Bloch Point (Figure 7.2). Little research followed until Thiaville [50] calculated through micromagnetic simulations the dynamics of two reversal mechanisms for vortices through either Bloch Point pair creation, or via the transition of a Bloch Point from one surface to another. The Bloch Point structure was studied and the energy minima was found to switch between the ‘radial’ hedgehog Bloch Point and the ‘twisted’ Bloch Point state under fixed when the size of a nanosphere was varied [271]. More recent work has shown that Bloch Points can be stabilised in spherical nanoparticles under an applied field gradient [272]. Some studies have made use of novel multiphysics simulations, whereby micromagnetic and atomistic modelling are combined, to explore ultra-fast dynamics of Bloch Points in the reversal mechanisms of nanotubes and nanowires [273–275]. Numerical work on skyrmions states has shown that these, like vortices, can undergo reversal mediated by the introduction of a Bloch point [32, 131]. Other states, christened ‘bobbers’ due to the their superficial resemblance to a fishing float (named a bobber in American English), which are bounded inside a material by a Bloch Point (Figure 7.2 (d)) and which are metastable in chiral systems, were predicted through micromagnetic simulations [276]. Subsequently, experimental work identified such states in FeGe lamellae through differential phase contrast in LTEM imaging [246]. Meanwhile, direct observation of Bloch Points in ordinary ferromagnetic materials has been shown in many studies [277]

In order to explain the vortex reversal process, Miltat and Thiaville [278] utilised an analogy; by taking a copy of a system with a vortex, reversing the core magnetisation and chirality, and stacking the two structures, the magnetisation is continuous at all points except for at the core of the sample where the two vortices meet, at which point there is a Bloch Point. If such a structure were to be produced, the Bloch point would be unstable and exit the sample, leading to a final single vortex state. However, in materials which host a Dzyaloshinskii-Moriya interaction, a fixed sense of rotation is enforced due to the chiral nature of the interaction. Because of this, it is easy to see that if the top and bottom structures have Dzyaloshinskii-Moriya constant D of opposite sign, then vortices with the same chirality but opposite core orientations would be of minimum energy in the top and bottom layers, keeping a Bloch Point fixed in place. This is fundamentally a different mechanism of stabilising a vortex than has been previously considered. In previous work [279], the study of such a system was shown. In a stacked layer disk structure of the bulk chiral material FeGe, a Bloch Point state was found to be realisable at the centre of a nanodisk where two opposite orientation vortex cores meet, and the field driven reversal mechanism from a

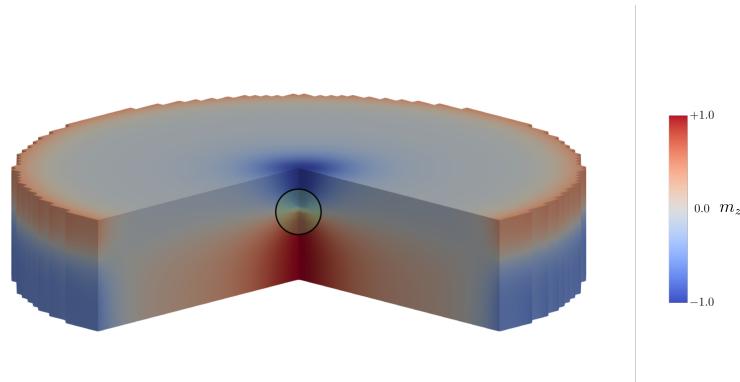


FIGURE 7.3: OOMMF Simulation of a 150 nm nanodisk, with bottom and top layer thicknesses of 12 nm and 20 nm and Dzyaloshinskii-Moriya respectively. The Bloch Point is located at the centre of the sample, with a circle indicating the position.

‘head-to-head’ to ‘tail-to-tail’ Bloch point was shown through a hysteresis loop. In this prior work, the main contribution of the present author was the verification of finite element simulations through finite difference simulations with OOMMF [126]. It is worth considering how such a structure can be created. Recent work by Spencer, et. al. [280] showed that zero crossing of the DMI constant could occur through changing the concentration of Co in $Fe_{1-y}Co_yGe$. During thin-film deposition, by doing this in a controlled manner, it is conceivable to consider building up an Fe-rich layer and then a second Co-rich layer, in which each layer then has a bulk Dzyaloshinskii-Moriya interaction with opposite sign and preferred sense of rotation.

In this chapter, finite element simulations are performed on nanodisks and nanotracks. The equilibrium behaviour of a Bloch Point in a nanodisk under an in-plane applied field is determined, and the effect of thickness studied, showing that the Bloch Point state is robust under weak perturbation, and the strength of applied field needed to expel a Bloch Point is determined. Manipulation of Bloch Points using an applied field in a long nanotrack is then studied, and it is shown that while in motion, the vortices comprising the Bloch Point state are distorted but remains intact under a laterally applied field. These results show that the Bloch Point shows promise as a candidate for future device technologies in a similar manner to other topological magnetic states.

The present author and Dr. Marijan Beg conceived of this study. The present author wrote and performed the simulations and analysis and interpreted the results, and the results were discussed with Prof. Hans Fangohr and Dr. Marijan Beg. The data from this study has been made available [281].

7.2 Field Induced Bloch Point Motion

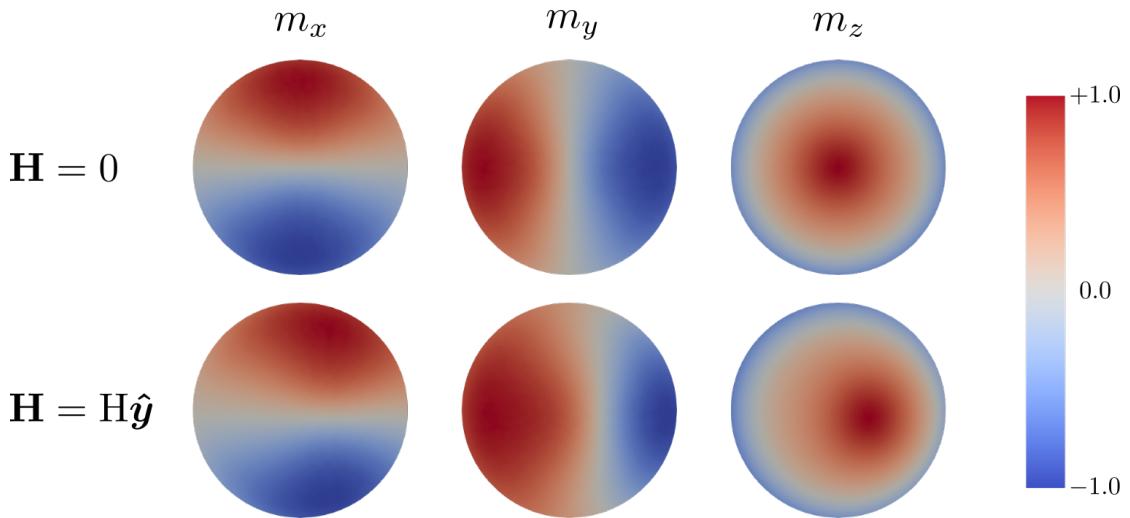


FIGURE 7.4: We show here the components of the magnetisation for a vortex state in a nanodisk of FeGe of diameter 80 nm and with thickness $t = 10$ nm. We see that under the application of an applied magnetic field, along the y direction, the vortex core is displaced, with the region of the magnetisation pointing along the direction of the field growing in size.

7.2.1 Bloch Point in Equilibrium under an Applied Field

It was shown in previous work [279] that in a nanodisk system hosting a bulk Dzyaloshinskii-Moriya interaction with $D > 0$ in a top layer and $D < 0$ in a bottom layer, the magnetisation in the top layer will follow the chirality of the bottom layer if the top layer is sufficiently thin in comparison. This holds even though it is energetically unfavourable to do so with respect to that layer's Dzyaloshinskii-Moriya interaction. To understand how such a layered system will respond to an in-plane applied field, a non-layered nanodisk system is first considered. The nanodisk has a diameter of 80 nm and 10 nm thickness, and the material parameters are set to those of FeGe with a bulk Dzyaloshinskii-Moriya interaction, as in Chapter 5. The magnetisation is initialised with the magnetisation aligned along the z axis. The system is relaxed under highly-damped ($\alpha = 1$) Landau-Lifshitz-Gilbert dynamics, and the system relaxes to an ‘incomplete skyrmion’ or radial vortex. With no applied field, there is a symmetry in the system, and the core of this vortex can point in either direction with degenerate energy. However, the chirality of the nanodisk is fixed with respect to the core orientation because of the Dzyaloshinskii-Moriya interaction and is controlled by the sign of the Dzyaloshinskii-Moriya constant.

Consider now an applied field incident on this on the sample along the $+y$ direction. Figure 7.4 shows that the vortex core moves, and this motion acts to maximise the region of the magnetisation aligned along $+y$, and correspondingly minimises the region of the system which is aligned in the $-y$ direction. Now consider, applying a magnetic field in the same way as in the above vortex system, in a bilayered nanodisk. It is expected that in the same way, the region of the magnetisation aligned with the applied field should grow, while the region aligned anti-parallel should shrink. This is tested by simulating such a system. The diameter of the system is again 80 nm. This time, a top layer of the nanodisk is kept at a fixed thickness, as $h_t = 10$ nm, and the bottom layer thickness h_b is varied from 12 nm to 16 nm. The initial magnetisation is set to point along $+\hat{z}$ direction, and then relax the system under Landau-Lifshitz-Gilbert dynamics with $\alpha = 1$ in order to stabilise a Bloch Point state. A magnetic field is then applied along the $+\hat{y}$ direction, and the system relaxed to equilibrium repeatedly, with the field increased in steps $\mu_0 dH = 0.01$ T until the applied field reaches $\mu_0 H = 0.1$ T.

Before analysing the results, it is important to consider how a Bloch Point can be detected in a simulation. The winding number defined by Feldtkeller [263, 282], determines the number of times the magnetisation winds over a surface:

$$W = \frac{1}{A} \int_S \mathbf{m} \cdot \hat{\mathbf{n}} \, dS \quad (7.1)$$

where A is the surface area of the bounding region. For convenience in the finite-element simulations, we can recast this using Stokes' theorem:

$$W = \frac{1}{A} \int_V \nabla \cdot \mathbf{m} \, dV \quad (7.2)$$

Given that the entire sample geometry encloses a Bloch Point if it is present, it can thus countably be known the number of Bloch Points inside a sample by reference to the magnetisation on the surface of the total geometry. It is shown in Figure 7.6 that the Bloch Point behaves in a similar manner to an ordinary radial vortex. As the field is increased, the equilibrium position of the Bloch Point shifts, and beyond a certain field strength, the Bloch Point is expelled from the sample. In order to calculate programmatically the presence of the Bloch Point and the expulsion field, the winding number is computed for the system at each stage. It is found that due to the numerical discretisation and twisting at the sample edges, this value is not integer quantised when a Bloch Point is present, but is very close to zero when no Bloch Point is present. In Figure 7.6, the winding number is shown for each of the nanodisk thicknesses, and one can see that as the thickness

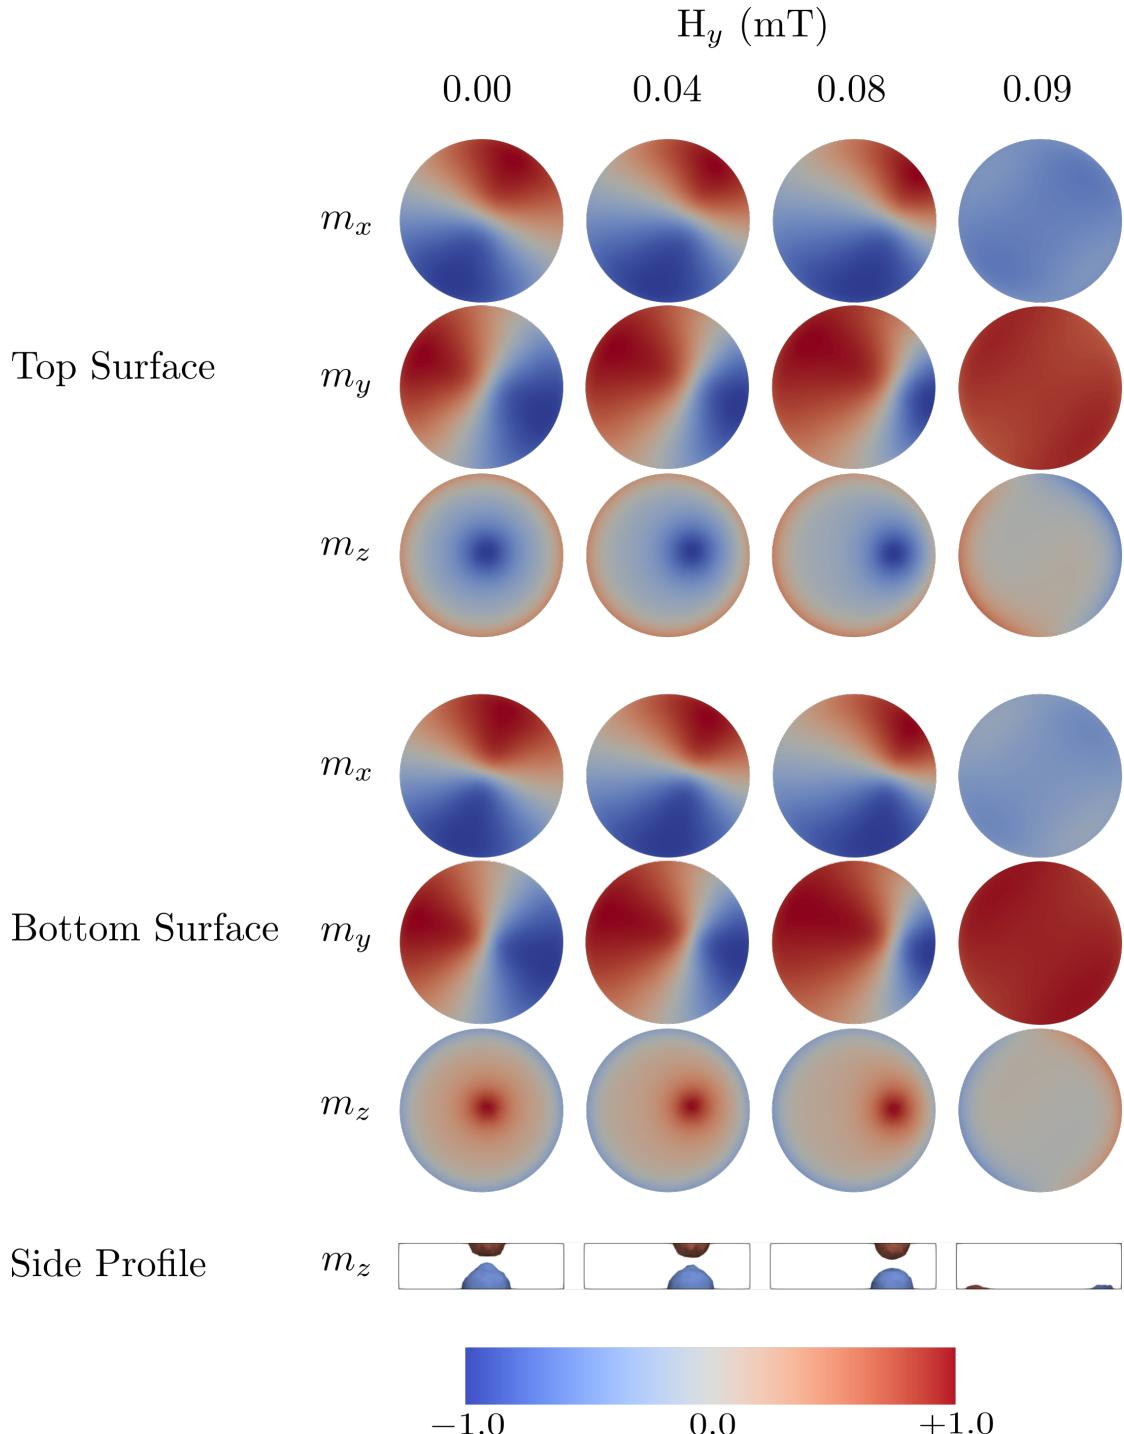


FIGURE 7.5: Here, the surface magnetisation on the top and bottom of an FeGe nanocylinder with two layers is shown for different applied field strengths, along with a side profile showing magnetisation contours at $m_z = \pm 0.6$. When $H_y = 0.9$, the Bloch Point state is found to have been expelled from the system.

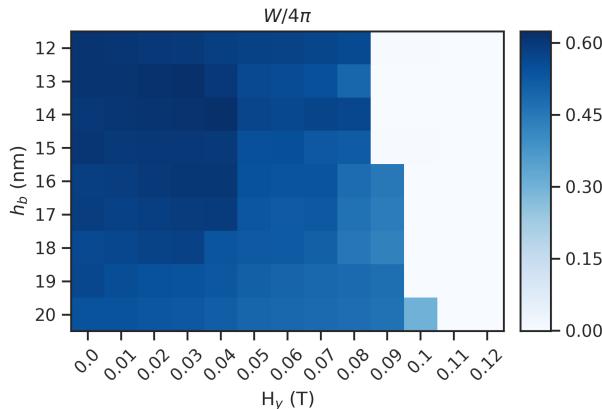


FIGURE 7.6: Winding Number $W/4\pi$ for the disk systems at a given applied field. When the winding number drops to zero, the Bloch Point has been expelled from the sample edge.

of the bottom layer is increased, the field required to expel the Bloch Point from the disk increases.

In order to explain this behaviour, one must consider the relative energies in a given system. When the Bloch Point is present within the system, the magnetisation is varying relatively slowly aside from in the vicinity of the Bloch Point itself, and the magnetostatic surface contribution to the energy is small. When the Bloch Point is moved closer to the edge of the system by an applied field, the magnetisation becomes more aligned in a single direction, decreasing the exchange energy, while increasing the Dzyaloshinskii-Moriya energy. Meanwhile, the magnetostatic energy increases due to surface charges, and the field acts to push the Bloch Point back into the sample centre, thus increasing the magnetostatic contribution to the energy. Kicking the Bloch Point from the sample then, happens when the decrease in the energy from the Zeeman contribution as the magnetisation aligns with the applied field is sufficient to overcome the increase in energy contribution from the other terms. In thicker samples, a higher field is thus required in particular because the magnetostatic field acting to demagnetise the sample is larger.

7.2.2 Bloch Point Velocity

FeGe nanotrack geometries which are sufficiently long as to observe motion without interference from the boundary are now considered, in order to try and measure the velocity of the Bloch Point under an applied field. A system is constructed in the same manner as before, with a bottom layer with Dzyaloshinskii-Moriya constant $-D$, and an upper layer with $+D$. The thicknesses of the bottom and

top layers are now kept fixed as $h_b = 14$ nm and $h_t = 10$ nm respectively. First, in order to investigate whether the size of the system has any bearing on the formation of the Bloch Point state within nanotacks, the width L_x and depth L_y of the nanotacks between 20 nm to 280 nm (note that the helical length of FeGe is 70 nm). Here, only half of the phase diagram is computed, with the other half mirrored. Initially, the magnetisation is set to point along the $+z$ direction. Each system is then relaxed under heavily damped ($\alpha = 1.0$) Landau-Lifshitz-Gilbert dynamics until $|\partial\mathbf{m}/\partial t| < 10^{-6}$ deg / ns, in order that the system finds a metastable state. In Figure 7.7, the winding number is shown along with the average magnetisation in each layer, which proves to be a useful proxy for the presence of the Bloch Point. It is notable samples up to three times the helical length in lateral dimensions still relax into the Bloch Points state, suggesting that the coupling between layers suppresses the formation of helices. For systems in which the lateral dimensions are < 40 nm, we do not always observe a Bloch Point state.

With the knowledge that a Bloch Point can be observed in large nanotrack systems, an intermediate sized sample is chosen such that studying the sustained motion of the Bloch Point is feasible. A nanotrack of 200 nm in length is studied, with the width L_y varied between 60, 80 and 100 nm. When an-plane magnetic field is applied to the system in the $+y$ direction, the Bloch Point in the centre of the sample is induced to move along the $-x$ direction. In order to suppress oscillatory modes along the edges of the sample which substantially increased the calculation time due to the adaptive time step integration method, we again choose a high damping value of $\alpha = 1.0$, though we note that in tests we found that the qualitative behaviour was not altered by this.

In Figure 7.8, we show snapshots of the Bloch Point state at different times for the 60 nm width nanotrack at two different fields. For the low field of $H_y = 0.05$ T, the Bloch Point initially moves before reaching an equilibrium position, consistent with the behaviour found in the nanodisk where the expulsion field was found to be between 0.09 and 0.10 T at the same layer thicknesses. For the larger field of $H_y = 0.20$ T, the Bloch Point moves substantially faster, and is ejected from the side of the sample at $t = 0.26$ ns. In both cases, on the application of an applied field, the vortex cores which make up the magnetisation are distorted, with this distortion strongest in the thicker layer. Notably, when the Bloch Point slows down through the force from the boundary, the distortion is reduced, and the bottom and top vortex cores again become near symmetric.

The equilibrium position of the Bloch Point is calculated for each width of nanotrack. In order to track this, the m_z component of the magnetisation is sampled over a uniform grid on the top and bottom surfaces of the mesh. From this, the maximum and minimum of the m_z component on each are found, which gives a coordinate for the vortex core position on each surface. Figure 7.9 shows this for each of the three nanotacks.

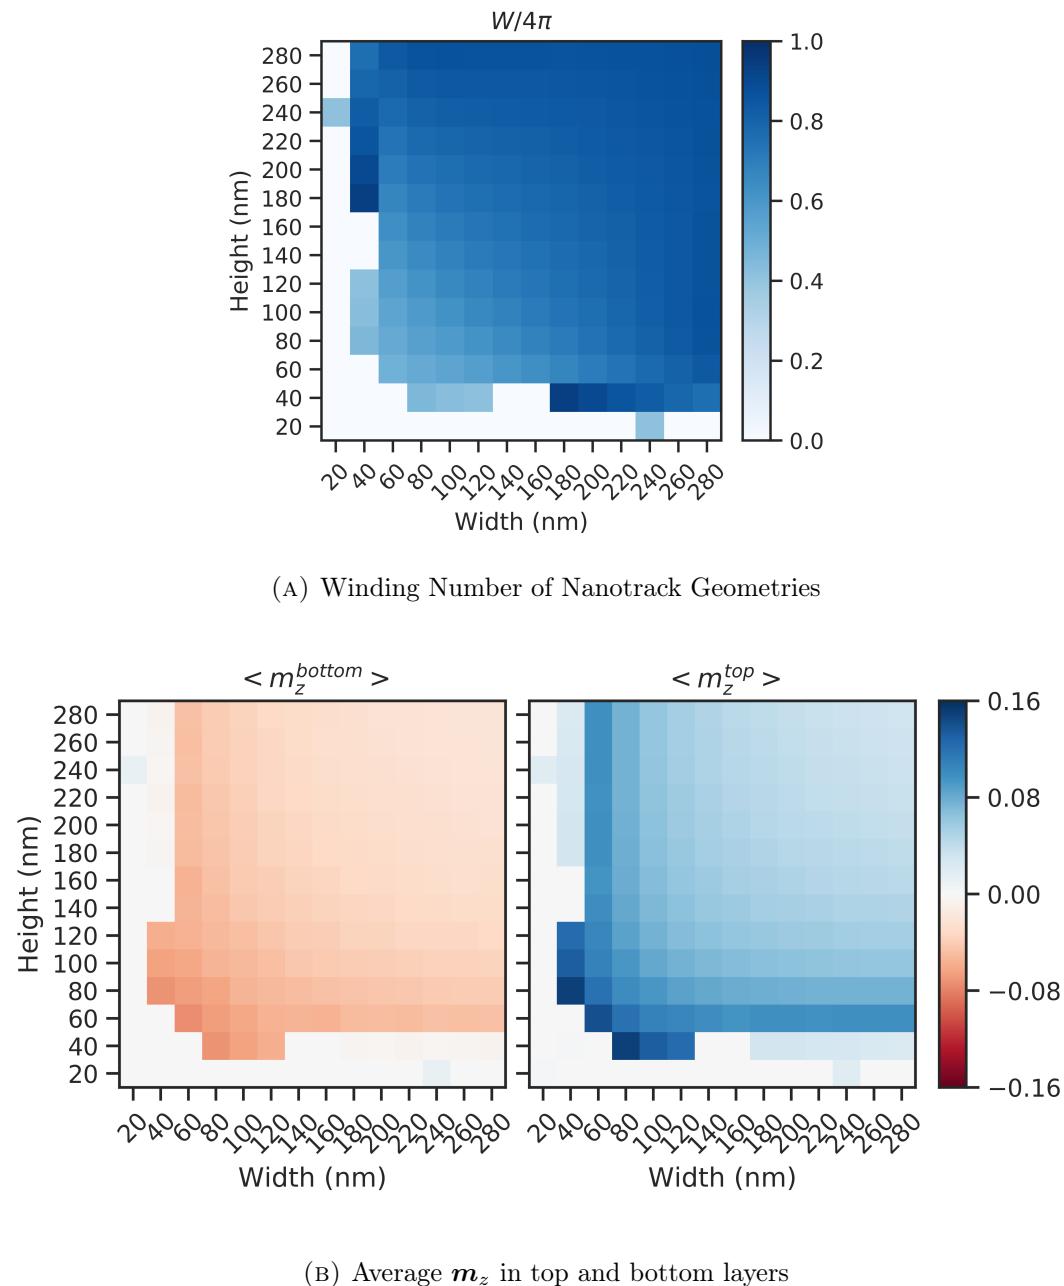


FIGURE 7.7: We see here that in both diagrams a jump in the average magnetisation and skyrmion number in each layer. Where this jump occurs, Bloch Points are found to be the observed state in the larger part of the phase diagram. We note that we see sudden jumps in winding number indicating the presence of a Bloch Point at width/height of 40 nm and height/width of 180 nm. We believe that this may be simply our simulation procedure not finding Bloch Points in the intermediate range, even though these are energetically possible.

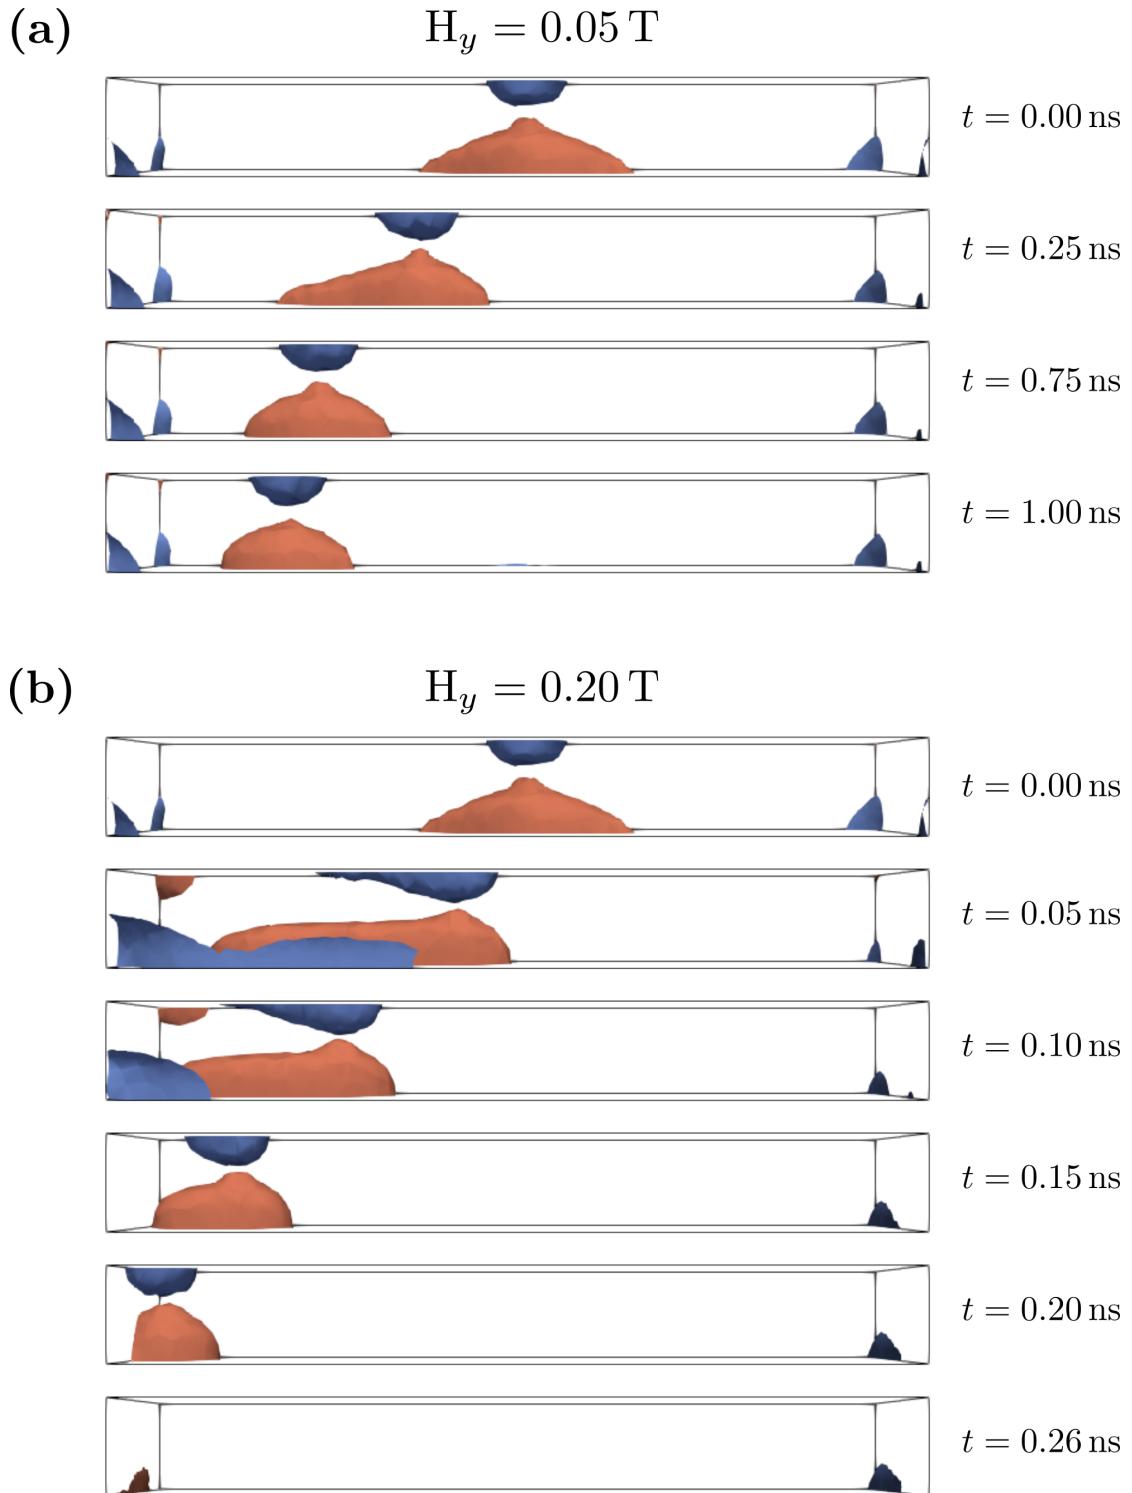


FIGURE 7.8: Snapshots of the Bloch Point Motion in the $L_y = 60$ nm system. Here, contours of the magnetisation are taken at $m_z = -0.6$ (blue) and 0.6 (red).

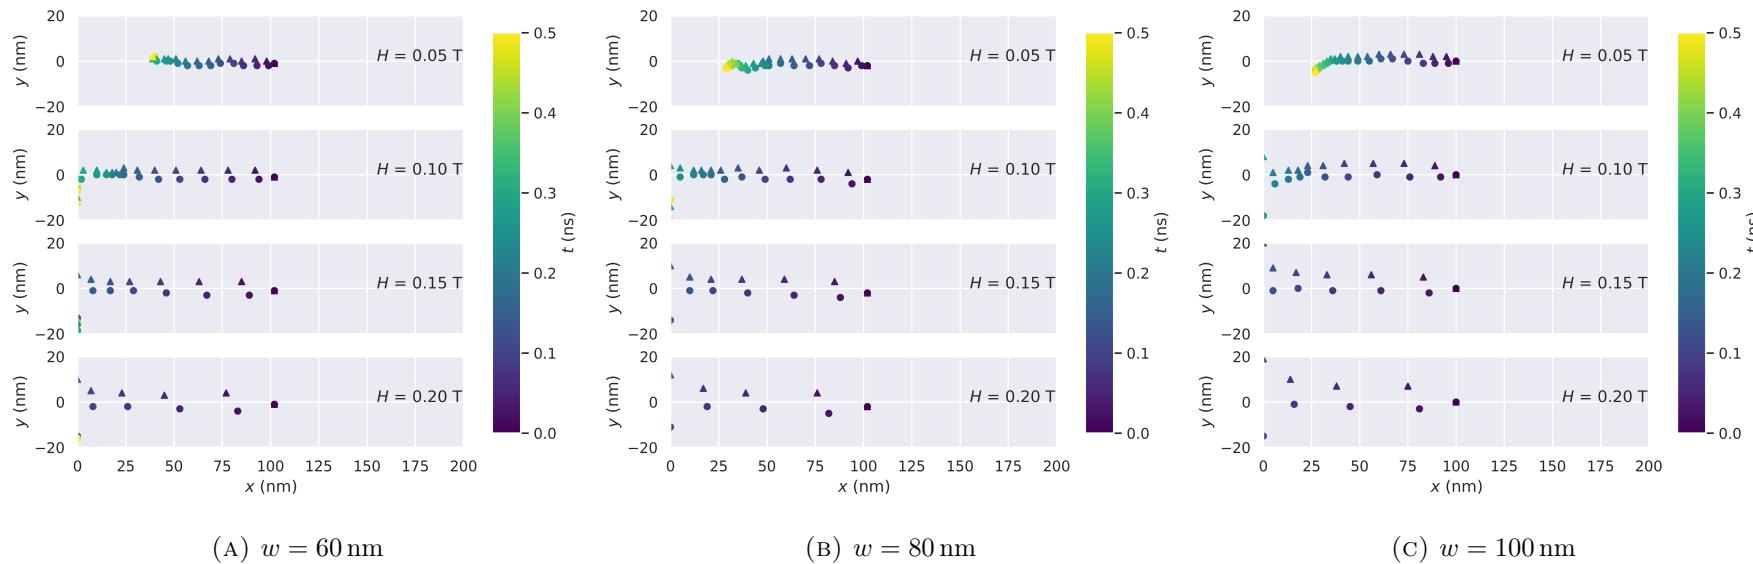


FIGURE 7.9: Positions identified as the core of the vortices in the top (circular point) and bottom surfaces (triangular point) of the system in the $x - y$ plane. The colours of each scatter point indicates the time t at which the Bloch Point is located at that position.

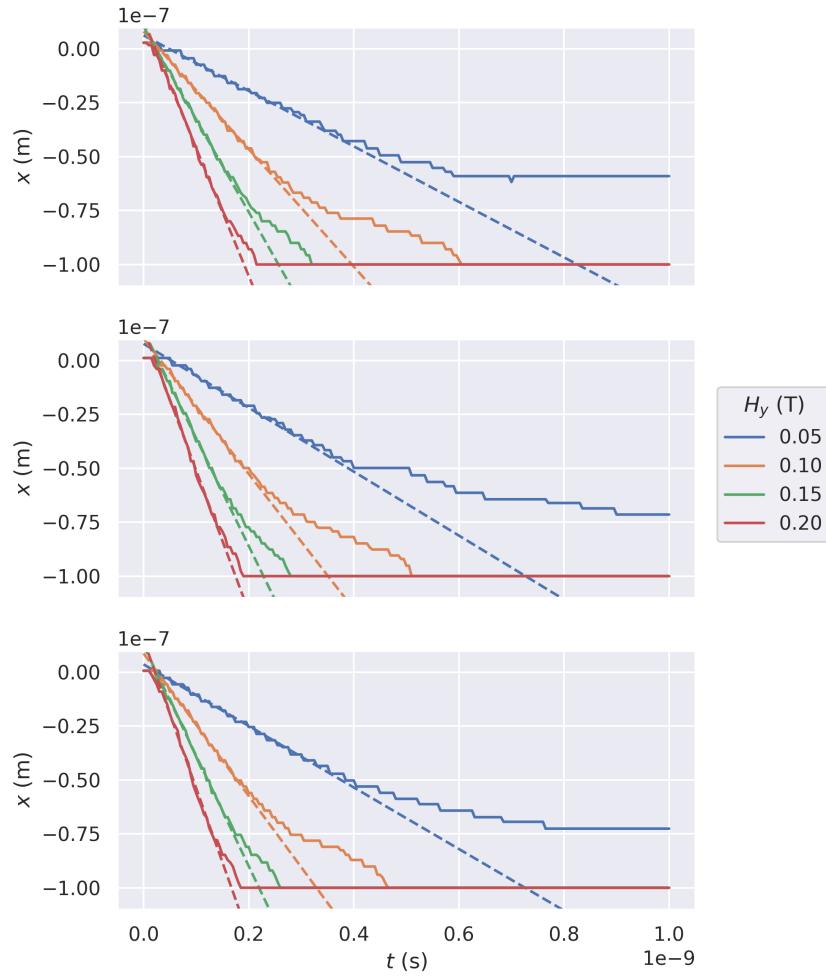
We show the two vortex core positions for each of the systems in Figure 7.9. The behaviour observed whereby the vortex cores are distorted can clearly be seen through the separation of the vortex core maxima and minima, and is stronger with increased applied field. It can be observed that when the Bloch Point is located close to the centre of the sample, the motion of the Bloch Point appears to be linear with regards to the strength of the applied field H , while beyond this the Bloch Point slows down as it approaches the boundary of the sample. In order to calculate the equilibrium velocity, the average position of the vortex cores is taken, and a straight line is fitted with least squares over the linear region of the velocity. This approach is necessary as due to the numerical discretisation of the field, the vortex core position appears to make small jumps over time rather than smoothly varying. We plot the x position with the fitted lines in Figure 7.10a.

We show the magnitudes of the calculated velocities against the applied field in each of the nanotracks in Figure 7.10b. For the 60 nm track, the velocity increases at a lower rate in comparison to the other two tracks, suggesting that for smaller widths, where the vortex cores are close to the boundary, the Bloch Point interacts with the track boundary which acts to resist the motion. This behaviour of the system is consistent with that of vortices; potential curves have been previously calculated for vortices in elliptical nanodisk which show different potentials along the semi-major and semi-minor axes, [283], while the oscillatory motion of the vortex core caused by applied currents was found to be damped in thin nanoelements as the system size decreased due to the edge potential [284].

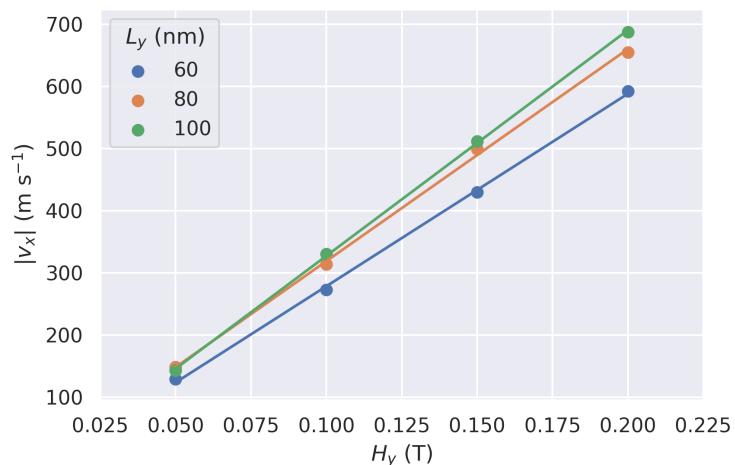
7.3 Conclusion

In this chapter, we have considered the motion of Bloch Points, showing that the Bloch Point motion is in many ways similar to that of vortices. We calculate the displacement of the Bloch Point from the natural position in a nanodisk, and as a result calculate the required field to expel the Bloch Point from the system for varying thicknesses. We show that Bloch Points can be hosted in a wide range of nanotrack samples, and then study the sustained motion of a Bloch Point in such a nanotrack, showing that in narrow nanotracks, the Bloch Point velocity is linear with the strength of the applied field, and that when the width is increased, the motion of the Bloch Point is faster.

This work is important because it shows that the Bloch Point state can be manipulated in a similar manner to a vortex. With that in mind, there are studies of



(A) Bloch Point position taken as an average between the top and bottom cores positions. A linear fit was applied to the data through least-squares regression for each series in the linear region of the motion.



(B) Bloch Point velocity from least squares fits. A linear regression was fitted to each of the data series, and the gradients k found were $k_{60\text{ nm}} = 3096.87 \text{ m s}^{-1}\text{T}^{-1}$, $k_{80\text{ nm}} = 3411.82 \text{ m s}^{-1}\text{T}^{-1}$ and $k_{100\text{ nm}} = 3631.17 \text{ m s}^{-1}\text{T}^{-1}$.

FIGURE 7.10: Position and Velocity determination for the Bloch Point in the three nanotacks

vortices which can inform and help to predict the behaviour of a Bloch Point in a layered system, but there are many differences and hence open questions remain as to how such systems will behave. As an example, a topic of interest in vortex research has been dynamic switching, whereby an applied field pulse induces a reversal of either the core orientation, or the circulation, reducing the field needed to do so over a static applied field in this study by the excitation of modes in the system [285, 286]. In the BP systems discussed in this work, the top and bottom layers are each in a vortex-like state, but because of the fixed chirality, there are only two states which are possible, and it is therefore likely that dynamic switching may occur differently.

Chapter 8

Conclusion

In this thesis, we consider the Barnes-Hut and Fast Multipole Methods as used in long range calculations. A software framework for generating operator functions for both methods, which can be used in implementations was developed and released publicly as ‘fmmgen’, and test implementations of both methods were developed as part of this work. We show that the use of common subexpression elimination, which removes repeated operations from calculations, significant performance increases can be seen, and compare this the algebraically more complicated precomputation of derivatives using the harmonic properties of $1/r$ for the method. We show that the performance here is strongly dependent on the compiler chosen; the Intel C++ compiler outperforms the GNU compiler in all cases studied, but also reduces the impact of the harmonic derivative optimisation. We note that this type of performance study must be done on the target hardware, due to inaccurate measures of operation counts in symbolic algebra libraries which have no bearing on the real computational cost. We implement the dipolar field calculation via the FMM into the simulation package Fidimag. We find that contrary to the behaviour shown in some publications, the FMM does not outperform the common FFT convolution method as implemented in most atomistic simulation software over the numbers of particles as can be currently modelled.

We then consider several systems comprised of the Dzyaloshinskii-Moriya interaction hosting material FeGe through micromagnetic simulations, and investigate these in different conditions. In Chapter 5, we investigated how size affects the formation of skyrmions in thin ($t = 10$ nm) FeGe samples. We demonstrate that skyrmions, incomplete skyrmions (radial vortices) and helices can all form the ground state under different geometric conditions, and that a wide variety of other

states can be observed. We suggest an experimental procedure for observing both skyrmion and target skyrmions by applying and removing an applied field, and show the expected states which would form under such a procedure.

We then study the presence of Bloch Points in FeGe system, taking further the work in a previous study [279]. The field driven behaviour of Bloch points in bilayered FeGe nanostructures in which each layer has an opposite sign Dzyaloshinskii-Moriya are investigated for nanodisks and nanotracks. The mechanism for field driven motion of a Bloch Point in such systems is considered, and shown to be identical to that of a vortex due to each layers vortex-like magnetisation texture. It is found that in nanodisks, the Bloch Point can be expelled through the side of the sample by the application of a sufficiently strong applied field, but that below this, the Bloch point can sit in equilibrium with the field. The relative thickness of the layers in the sample is shown to have an effect on the expulsion field, and this is explained by considering the relative micromagnetic energy contributions.

The formation of Bloch points in nanotracks is considered, and we show that the Bloch point state forms from an initially uniformly magnetised sample for a wide range of side lengths. The motion of the Bloch point in 200 nm nanotrack systems is then considered, and the Bloch Point position is tracked by calculating the vortex core positions for each layer. We show that the Bloch point velocity is linear with the applied field when far enough from the sample edge that no force is felt from the boundary, and that consistent with the results in the nanodisk, for sufficiently strong applied fields, the Bloch point is expelled from the system. These results prove further that Bloch points are an object worthy of study.

Finally, we investigate the energy barriers for Skyrmion destruction and creation via edge defects in Chapter 6 by means of atomistic simulations and the use of the Nudged Elastic Band Method. We show that the defect size and shape has an effect on the energy barrier height, with large, deep, defects generally reducing the skyrmion annihilation activation energy to nearly zero in Cobalt monolayers. We show that contrary to what may be expected, increasing the defect width actually causes the energy barrier height to increase. In these studies, it is shown that including the dipolar field is crucial for determining the barrier height for large defects accurately.

References

- [1] T. Lindh, P. Salminen, J. Pyrhonen, M. Niemela, J. Kinnunen, and J. Haataja. Permanent Magnet Generator Designing Guidelines. In *2007 International Conference on Power Engineering, Energy and Electrical Drives*. IEEE, April 2007. <https://dx.doi.org/10.1109/powereng.2007.4380172>.
- [2] K.T. Chau, C.C. Chan, and Chunhua Liu. Overview of Permanent-Magnet Brushless Drives for Electric and Hybrid Electric Vehicles. *IEEE Transactions on Industrial Electronics*, 55(6):2246–2257, June 2008. <https://dx.doi.org/10.1109/tie.2008.918403>.
- [3] M. J. Perenson. The hard drive turns 50, Sep 2006. <https://www.pcworld.com/article/127104/article.html>.
- [4] M. N. Baibich, J. M. Broto, A. Fert, F. Nguyen Van Dau, F. Petroff, P. Eitenne, G. Creuzet, A. Friederich, and J. Chazelas. Giant magnetoresistance of (001)Fe/(001)Cr magnetic superlattices. *Physical Review Letters*, 61(21):2472–2475, 1988. ISSN 00319007. <https://dx.doi.org/10.1103/PhysRevLett.61.2472>.
- [5] G. Binasch, P. Grünberg, F. Saurenbach, and W. Zinn. Enhanced magnetoresistance in layered magnetic structures with antiferromagnetic interlayer exchange. *Physical Review B*, 39(7):4828–4830, 1989. ISSN 01631829. <https://dx.doi.org/10.1103/PhysRevB.39.4828>.
- [6] H. J. Richter, A. Lyberatos, U. Nowak, R. F. L. Evans, and R. W. Chantrell. The thermodynamic limits of magnetic recording. *Journal of Applied Physics*, 111(3):033909, February 2012. <https://dx.doi.org/10.1063/1.3681297>.

- [7] Y. Tanaka, S. Iwasaki, and H. Muraoka. Perpendicular Magnetic Recording. In *2017 IEEE HISTory of Electrotechnology CONference (HISTELCON)*. IEEE, August 2017. <https://dx.doi.org/10.1109/histelcon.2017.8535815>.
- [8] H. J. Richter. The transition from longitudinal to perpendicular recording. *Journal of Physics D: Applied Physics*, 40(9):R149–R177, 2007. ISSN 0022-3727. <https://dx.doi.org/10.1088/0022-3727/40/9/R01>.
- [9] D. Klein. The history of semiconductor memory: From magnetic tape to NAND flash memory. *IEEE Solid-State Circuits Magazine*, 8(2):16–22, 2016. <https://dx.doi.org/10.1109/mssc.2016.2548422>.
- [10] Y. Shiroishi, K. Fukuda, I. Tagawa, H. Iwasaki, S. Takenoiri, H. Tanaka, H. Mutoh, and N. Yoshikawa. Future options for HDD storage. *IEEE Transactions on Magnetics*, 45(10):3816–3822, October 2009. <https://dx.doi.org/10.1109/tmag.2009.2024879>.
- [11] T. R. Albrecht, H. Arora, V. Ayanoor-Vitikkate, J-M. Beaujour, D. Bedau, D. Berman, A. L. Bogdanov, Y-A. Chapuis, J. Cushen, E. E. Dobisz, G. Doerk, H. Gao, M. Grobis, B. Gurney, W. Hanson, O. Hellwig, T. Hirano, P-O. Jubert, D. Kercher, J. Lille, Z. Liu, C. M. Mate, Y. Obukhov, K. C. Patel, K. Rubin, R. Ruiz, M. Schabes, L. Wan, D. Weller, T-W. Wu, and E. Yang. Bit-Patterned Magnetic Recording: Theory, Media Fabrication, and Recording Performance. *IEEE Transactions on Magnetics*, 51(5):1–42, May 2015. <https://dx.doi.org/10.1109/tmag.2015.2397880>.
- [12] T. Rausch, E. Gage, and J. Dykes. Heat assisted magnetic recording. *Springer Proceedings in Physics*, 159(11):200–202, 2015. ISSN 18674941. https://dx.doi.org/10.1007/978-3-319-07743-7_63.
- [13] M.T. Kief and R.H. Victora. Materials for heat-assisted magnetic recording. *MRS Bulletin*, 43(2):87–92, February 2018. <https://dx.doi.org/10.1557/mrs.2018.2>.
- [14] A. Nordrum. Lasers vs. Microwaves: The Billion-Dollar Bet on the Future of Magnetic Storage. *IEEE Spectrum*, 2019. <https://spectrum.ieee.org/computing/hardware/lasers-vs-microwaves-the-billion-dollar-bet-on-the-future-of-magnetic-storage>.
- [15] X. Yang, Y. Xu, K. Lee, S. Xiao, D. Kuo, and D. Weller. Advanced lithography for bit patterned media. *IEEE Transactions on Magnetics*, 45

(2):833–838, 2009. ISSN 00189464. <https://dx.doi.org/10.1109/TMAG.2008.2010647>.

[16] R. A. Griffiths, A. Williams, C. Oakland, J. Roberts, A. Vijayaraghavan, and T. Thomson. Directed self-assembly of block copolymers for use in bit patterned media fabrication. *Journal of Physics D: Applied Physics*, 46(50), 2013. ISSN 00223727. <https://dx.doi.org/10.1088/0022-3727/46/50/503001>.

[17] C. Vogler, C. Abert, F. Bruckner, D. Suess, and D. Praetorius. Heat-assisted magnetic recording of bit-patterned media beyond 10 tb/in². *Applied Physics Letters*, 108(10):102406, March 2016. <https://dx.doi.org/10.1063/1.4943629>.

[18] K. Hono, Y.K. Takahashi, Ganping Ju, Jan-Ulrich Thiele, Antony Ajan, XiaoMin Yang, Ricardo Ruiz, and Lei Wan. Heat-assisted magnetic recording media materials. *MRS Bulletin*, 43(2):93–99, February 2018. <https://dx.doi.org/10.1557/mrs.2018.5>.

[19] E. A. Torrero. Memories: Bubbles rise from the lab: Developmental prototypes and even limited production of 65-kb and larger bubble memories are the result of advances on a broad front. *IEEE Spectrum*, 13(9):29–32, 1976. <https://dx.doi.org/10.1109/mspec.1976.6367514>.

[20] D. Toombs. Solid state: CCD and bubble memories system implications: Mainframe computers and intelligent terminals are predicted to be primary beneficiaries. *IEEE Spectrum*, 15(5):36–39, May 1978. <https://dx.doi.org/10.1109/mspec.1978.6367696>.

[21] J. Gesner and B. Morrison. Bubble memory applications for aircraft systems. In *Proceedings of the IEEE National Aerospace and Electronics Conference*. IEEE, 1989. <https://dx.doi.org/10.1109/naecon.1989.40192>.

[22] S. S. P. Parkin, M. Hayashi, and L. Thomas. Magnetic Domain-Wall Racetrack Memory. *Science*, 320(5873):190–194, 2008. ISSN 0036-8075. <https://dx.doi.org/10.1126/science.1145799>.

[23] M. Hayashi, L. Thomas, R. Moriya, C. Rettner, and S. S. P. Parkin. Current-Controlled Magnetic Domain-Wall Nanowire Shift Register. *Science*, 320(5873):209–211, April 2008. <https://dx.doi.org/10.1126/science.1154587>.

[24] I. Polenciu, A. J. Vick, D. A. Allwood, T. J. Hayward, G. Vallejo-Fernandez, K. O'Grady, and A. Hirohata. Domain wall pinning for racetrack memory using exchange bias. *Applied Physics Letters*, 105(16):162406, October 2014. <https://dx.doi.org/10.1063/1.4899134>.

[25] R. Tomasello, E. Martinez, R. Zivieri, L. Torres, M. Carpentieri, and G. Finocchio. A strategy for the design of skyrmion racetrack memories. *Scientific Reports*, 4(1), October 2014. <https://dx.doi.org/10.1038/srep06784>.

[26] A. Fert, V. Cros, and J. Sampaio. Skyrmions on the track. *Nature Nanotechnology*, 8(3):152–156, 2013. ISSN 1748-3387. <https://dx.doi.org/10.1038/nnano.2013.29>.

[27] A. N. Bogdanov and D. A. Yablonskii. Thermodynamically stable "vortices" in magnetically ordered crystals. The mixed state of magnets. *Zh. Eksp. Teor. Fiz.*, 95(1):178, 1989. ISSN 0044-4510. [https://dx.doi.org/10.1016/0304-8853\(94\)90046-9](https://dx.doi.org/10.1016/0304-8853(94)90046-9).

[28] A. Bogdanov and A. Hubert. Thermodynamically stable magnetic vortex states in magnetic crystals, 1994. ISSN 03048853. [https://dx.doi.org/10.1016/0304-8853\(94\)90046-9](https://dx.doi.org/10.1016/0304-8853(94)90046-9).

[29] U. K. Rößler, A. N. Bogdanov, and C. Pfleiderer. Spontaneous skyrmion ground states in magnetic metals. *Nature*, 442(7104):797–801, 2006. ISSN 0028-0836. <https://dx.doi.org/10.1038/nature05056>.

[30] S. Rohart and A. Thiaville. Skyrmion confinement in ultrathin film nanostructures in the presence of Dzyaloshinskii-Moriya interaction. *Physical Review B*, 88:184422, nov 2013. ISSN 10980121. <https://dx.doi.org/10.1103/PhysRevB.88.184422>.

[31] J. Sampaio, V. Cros, S. Rohart, A. Thiaville, and A. Fert. Nucleation, stability and current-induced motion of isolated magnetic skyrmions in nanostructures. *Nature Nanotechnology*, 8(11):839–844, 2013. ISSN 1748-3387. <https://dx.doi.org/10.1038/nnano.2013.210>.

[32] M. Beg, R. Carey, W. Wang, D. Cortés-Ortuño, M. Vousden, M-A. Bisotti, M. Albert, D. Chernyshenko, O. Hovorka, R. L. Stamps, and H. Fangohr. Ground state search, hysteretic behaviour, and reversal mechanism of skyrmionic textures in confined helimagnetic nanostructures. *Scientific*

Reports, 5(October):17137, 2015. ISSN 2045-2322. <https://dx.doi.org/10.1038/srep17137>.

[33] N. Romming, C. Hanneken, M. Menzel, J. E. Bickel, B. Wolter, K. von Bergmann, A. Kubetzka, and R. Wiesendanger. Writing and Deleting Single Magnetic Skyrmions. *Science*, 341(6146):636–639, 2013. ISSN 0036-8075. <https://dx.doi.org/10.1126/science.1240573>.

[34] R. Wiesendanger. Nanoscale magnetic skyrmions in metallic films and multilayers: a new twist for spintronics. *Nature Reviews Materials*, 1(7), June 2016. <https://dx.doi.org/10.1038/natrevmats.2016.44>.

[35] A. Soumyanarayanan, M. Raju, A. L. Gonzalez Oyarce, A. K. C. Tan, Mi-Young Im, A. P. Petrović, P. Ho, K. H. Khoo, M. Tran, C. K. Gan, F. Ernult, and C. Panagopoulos. Tunable room-temperature magnetic skyrmions in ir/fe/co/pt multilayers. *Nature Materials*, 16(9):898–904, July 2017. <https://dx.doi.org/10.1038/nmat4934>.

[36] A. Fert, N. Reyren, and V. Cros. Magnetic skyrmions: advances in physics and potential applications. *Nature Reviews Materials*, 2(7), June 2017. <https://dx.doi.org/10.1038/natrevmats.2017.31>.

[37] X. Zhang, G. P. Zhao, H. Fangohr, J. Ping Liu, W. X. Xia, J. Xia, and F. J. Morvan. Skyrmion-skyrmion and skyrmion-edge repulsions in skyrmion-based racetrack memory. *Scientific Reports*, 5(1), January 2015. <https://dx.doi.org/10.1038/srep07643>.

[38] W. Kang, Y. Huang, X. Zhang, Y. Zhou, and W. Zhao. Skyrmion-Electronics: An Overview and Outlook. *Proceedings of the IEEE*, 104(10):2040–2061, October 2016. <https://dx.doi.org/10.1109/jproc.2016.2591578>.

[39] W. Kang, Y. Huang, C. Zheng, W. Lv, Na Lei, Y. Zhang, X. Zhang, Y. Zhou, and W. Zhao. Voltage Controlled Magnetic Skyrmion Motion for Racetrack Memory. *Scientific Reports*, 6(1), March 2016. <https://dx.doi.org/10.1038/srep23164>.

[40] W. Kang, X. Chen, D. Zhu, X. Zhang, Y. Zhou, K. Qiu, Y. Zhang, and W. Zhao. A Comparative Study on Racetrack Memories: Domain Wall vs. Skyrmion. In *2018 IEEE 7th Non-Volatile Memory Systems and Applications Symposium (NVMSA)*. IEEE, August 2018. <https://dx.doi.org/10.1109/nvmsa.2018.00009>.

[41] W. Kang, B. Wu, X. Chen, D. Zhu, Z. Wang, X. Zhang, Y. Zhou, Y. Zhang, and W. Zhao. A Comparative Cross-layer Study on Racetrack Memories. *ACM Journal on Emerging Technologies in Computing Systems*, 16(1):1–17, October 2019. <https://dx.doi.org/10.1145/3333336>.

[42] N. Nagaosa and Y. Tokura. Topological properties and dynamics of magnetic skyrmions. *Nature Nanotechnology*, 8(12):899–911, December 2013. <https://dx.doi.org/10.1038/nnano.2013.243>.

[43] W. Jiang, X. Zhang, G. Yu, W. Zhang, X. Wang, M. Benjamin Jungfleisch, J. E. Pearson, X. Cheng, O. Heinonen, K. L. Wang, Y. Zhou, A. Hoffmann, and S. G. E. te Velthuis. Direct observation of the skyrmion Hall effect. *Nature Physics*, 13(2):162–169, September 2016. <https://dx.doi.org/10.1038/nphys3883>.

[44] D. Suess, C. Vogler, F. Bruckner, P. Heistracher, F. Slanovc, and C. Abert. Spin Torque Efficiency and Analytic Error Rate Estimates of Skyrmion Racetrack Memory. *Scientific Reports*, 9(1), 2019. <https://dx.doi.org/10.1038/s41598-019-41062-y>.

[45] C. Hanneken, F. Otte, A. Kubetzka, B. Dupé, N. Romming, K. von Bergmann, R. Wiesendanger, and S. Heinze. Electrical detection of magnetic skyrmions by tunnelling non-collinear magnetoresistance. *Nature Nanotechnology*, 10(12):1039–1042, October 2015. <https://dx.doi.org/10.1038/nnano.2015.218>.

[46] D. Maccariello, W. Legrand, N. Reyren, K. Garcia, K. Bouzehouane, S. Collin, V. Cros, and A. Fert. Electrical detection of single magnetic skyrmions in metallic multilayers at room temperature. *Nature Nanotechnology*, 13 (3):233–237, January 2018. <https://dx.doi.org/10.1038/s41565-017-0044-4>.

[47] C. Back, V. Cros, H. Ebert, K. Everschor-Sitte, A. Fert, M. Garst, Tianping Ma, S. Mankovsky, T. L. Monchesky, M. Mostovoy, N. Nagaosa, S. S. P. Parkin, C. Pfleiderer, N. Reyren, A. Rosch, Y. Taguchi, Y. Tokura, K. von Bergmann, and Jiadong Zang. The 2020 skyrmionics roadmap, 2019.

[48] D. Klein and J. Engemann. Bloch line memory: Dams for stripe domain confinement. *Journal of Applied Physics*, 57(8):4071–4073, April 1985. <https://dx.doi.org/10.1063/1.334675>.

[49] V.G. Red'ko and V.A. Skidanov. Magnetic bubble and vertical Bloch-line neural associative memory. In *[Proceedings] 1992 RNNS/IEEE Symposium on Neuroinformatics and Neurocomputers*. IEEE, 1992. <https://dx.doi.org/10.1109/rnns.1992.268576>.

[50] A. Thiaville, J. García, R. Dittrich, J. Miltat, and T. Schrefl. Micromagnetic study of Bloch-point-mediated vortex core reversal. *Physical Review B*, 67(9):094410, Mar 2003. <https://dx.doi.org/10.1103/PhysRevB.67.094410>.

[51] S. Rohart, J. Miltat, and A. Thiaville. Path to collapse for an isolated néel skyrmion. *Physical Review B*, 93(21), 2016. <https://dx.doi.org/10.1103/physrevb.93.214412>.

[52] A. Wartelle, B. Trapp, M. Staňo, C. Thirion, S. Bochmann, J. Bachmann, M. Foerster, L. Aballe, T. O. Mentes, A. Locatelli, A. Sala, L. Cagnon, J.-C. Toussaint, and O. Fruchart. Bloch-point-mediated topological transformations of magnetic domain walls in cylindrical nanowires. *Physical Review B*, 99(2), January 2019. <https://dx.doi.org/10.1103/physrevb.99.024433>.

[53] G. P. Müller, F. N. Rybakov, H. Jónsson, S. Blügel, and N. S. Kiselev. Coupled quasimonopoles in chiral magnets, 2019.

[54] G. Bihlmayer. Density Functional Theory for Magnetism and Magnetic Anisotropy. In *Handbook of Materials Modeling*, pages 1–23. Springer International Publishing, 2018. https://dx.doi.org/10.1007/978-3-319-42913-7_73-1.

[55] E. Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik*, 31(1):253–258, February 1925. <https://dx.doi.org/10.1007/bf02980577>.

[56] L. Onsager. Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition. *Physical Review*, 65(3-4):117–149, February 1944. <https://dx.doi.org/10.1103/physrev.65.117>.

[57] U. Nowak. Classical Spin Models. In *Handbook of magnetism and advanced magnetic materials*, chapter Classical, pages 858–876. Wiley, 2007.

[58] R. F. L Evans, W. J Fan, P. Chureemart, T. A. Ostler, M. O. A. Ellis,

and R. W. Chantrell. Atomistic spin model simulations of magnetic nanomaterials. *Journal of Physics: Condensed Matter*, 26(10):103202, 2014. <https://dx.doi.org/10.1088/0953-8984/26/10/103202>.

[59] D. Meilak, S. Jenkins, R. Pond, and R. F. L. Evans. Massively parallel atomistic simulation of ultrafast thermal spin dynamics of a permalloy vortex, 2019.

[60] L. D. Landau and E. Lifshitz. On the theory of the dispersion of magnetic permeability in ferromagnetic bodies. *Phys. Z. Sowjet.*, 8:153, 1935.

[61] W. F. Brown. *Micromagnetics*. John Wiley & Sons, New York, 1st edition, 1963. ISBN 978-0470110379.

[62] T. L. Gilbert. A phenomenological theory of damping in ferromagnetic materials. *IEEE Transactions on Magnetics*, 40(6):3443–3449, nov 2004. <https://dx.doi.org/10.1109/TMAG.2004.836740>.

[63] C. Kittel. Theory of the Structure of Ferromagnetic Domains in Films and Small Particles. *Physical Review*, 70(11-12):965–971, December 1946. <https://dx.doi.org/10.1103/physrev.70.965>.

[64] A. Bernand-Mantel, C. B. Muratov, and T. M. Simon. Unraveling the role of dipolar versus Dzyaloshinskii-Moriya interaction in stabilizing compact magnetic skyrmions, 2019.

[65] S. W. Yuan and H. N. Bertram. Fast adaptive algorithms for micromagnetics. *IEEE Transactions on Magnetics*, 28(5 pt I):2031–2036, 1992. ISSN 00189464. <https://dx.doi.org/10.1109/20.179394>.

[66] W. Zhang and S. Haas. Adaptation and performance of the Cartesian coordinates fast multipole method for nanomagnetic simulations. *Journal of Magnetism and Magnetic Materials*, 321(22):3687–3692, nov 2009. ISSN 03048853. <https://dx.doi.org/10.1016/j.jmmm.2009.07.016>.

[67] P. B. Visscher and D. M. Apalkov. Simple recursive implementation of fast multipole method. *Journal of Magnetism and Magnetic Materials*, 322(2): 275–281, 2010. ISSN 03048853. <https://dx.doi.org/10.1016/j.jmmm.2009.09.033>.

[68] R. A. Pepper and Fangohr H. F. fmmgen, 2020. <https://github.com/rpep/fmmgen>.

[69] J. C. Maxwell. VIII. a dynamical theory of the electromagnetic field. *Philosophical Transactions of the Royal Society of London*, 155:459–512, Jan 1865. <https://dx.doi.org/10.1098/rstl.1865.0008>.

[70] W. Gerlach and O. Stern. Der experimentelle Nachweis der Richtungsquantelung im Magnetfeld. *Zeitschrift für Physik*, 9(1):349–352, December 1922. <https://dx.doi.org/10.1007/bf01326983>.

[71] R. Eisberg and R. Resnick. *Quantum Physics of Atoms, Molecules, Solids, Nuclei and Particles*. Wiley, 2nd edition, 1985. ISBN 978-0-471-873730. <https://store.doverpublications.com/0486435032.html>.

[72] T. E. Phipps and J. B. Taylor. The Magnetic Moment of the Hydrogen Atom. *Physical Review*, 29(2):309–320, February 1927. <https://dx.doi.org/10.1103/physrev.29.309>.

[73] P. A. M. Dirac. The Quantum Theory of the Electron. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 117 (778):610–624, February 1928. <https://dx.doi.org/10.1098/rspa.1928.0023>.

[74] B. Odom, D. Hanneke, B. D’Urso, and G. Gabrielse. New Measurement of the Electron Magnetic Moment Using a One-Electron Quantum Cyclotron. *Physical Review Letters*, 97(3), July 2006. <https://dx.doi.org/10.1103/physrevlett.97.030801>.

[75] L. H. Thomas. The Motion of the Spinning Electron. *Nature*, 117(2945): 514–514, April 1926. <https://dx.doi.org/10.1038/117514a0>.

[76] P.C.W. Davies and D. S. Betts. *Quantum Mechanics*. Chapman and Hall, 2nd edition, 1994. ISBN 0-412-57900-6.

[77] P. Zeeman. XXXII. on the influence of magnetism on the nature of the light emitted by a substance. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 43(262):226–239, March 1897. <https://dx.doi.org/10.1080/14786449708620985>.

[78] P. Echenique and J. L. Alonso. A mathematical and computational review of Hartree–Fock SCF methods in quantum chemistry. *Molecular Physics*, 105(23-24):3057–3098, December 2007. <https://dx.doi.org/10.1080/00268970701757875>.

[79] A. D. Becke. Perspective: Fifty years of density-functional theory in chemical physics. *The Journal of Chemical Physics*, 140(18):18A301, May 2014. <https://dx.doi.org/10.1063/1.4869598>.

[80] R. O. Jones. Density functional theory: Its origins, rise to prominence, and future. *Reviews of Modern Physics*, 87(3):897–923, August 2015. <https://dx.doi.org/10.1103/revmodphys.87.897>.

[81] R. P. Joshi, J. J. Phillips, K. J. Mitchell, G. Christou, K. A. Jackson, and J. E. Peralta. Accuracy of density functional theory methods for the calculation of magnetic exchange couplings in binuclear iron(III) complexes. *Polyhedron*, page 114194, November 2019. <https://dx.doi.org/10.1016/j.poly.2019.114194>.

[82] S. G. Brush. History of the Lenz-Ising Model. *Reviews of Modern Physics*, 39(4):883–893, October 1967. <https://dx.doi.org/10.1103/revmodphys.39.883>.

[83] M. Krawczyk, M. L. Sokolovskyy, J. W. Klos, and S. Mamica. On the Formulation of the Exchange Field in the Landau-Lifshitz Equation for Spin-Wave Calculation in Magnonic Crystals. *Advances in Condensed Matter Physics*, 2012:1–14, 2012. <https://dx.doi.org/10.1155/2012/764783>.

[84] I. E. Dzyaloshinskii. Theory of Helicoidal Structures in Antiferromagnets. I. Nonmetals. *J. Exptl. Theoret. Phys. (U.S.S.R.)*, 19(4):960–971, 1964. http://www.jetp.ac.ru/cgi-bin/dn/e_019_04_0960.pdf.

[85] T. Moriya. Anisotropic Superexchange Interaction and Weak Ferromagnetism. *Physical Review*, 120(1):91–98, 1960. ISSN 0031-899X. <https://dx.doi.org/10.1103/PhysRev.120.91>.

[86] D. Cortés-Ortuño, M. Beg, V. Nehruji, L. Breth, R. Pepper, T. Kluyver, G. Downing, T. Hesjedal, P. Hatton, T. Lancaster, R. Hertel, O. Hovorka, and H. Fangohr. Proposal for a micromagnetic standard problem for materials with dzyaloshinskii–moriya interaction. *New Journal of Physics*, 20(11):113015, November 2018. <https://dx.doi.org/10.1088/1367-2630/aaea1c>.

[87] H. Kronmueller. Mikromagnetische berechnung der magnetisierung in der umgebung unmagnetischer einschlösse in ferromagnetika. *Zeitschrift für Physik*, 168(5):478–494, October 1962. <https://dx.doi.org/10.1007/bf01378144>.

[88] G. S. Abo, Y-K. Hong, J. Park, J. Lee, W. Lee, and B-C. Choi. Definition of Magnetic Exchange Length. *IEEE Transactions on Magnetics*, 49(8):4937–4939, August 2013. <https://dx.doi.org/10.1109/tmag.2013.2258028>.

[89] Dina Niculaes, Aidin Lak, George C. Anyfantis, Sergio Marras, Oliver Laslett, Sahitya K. Avugadda, Marco Cassani, David Serantes, Ondrej Hovorka, Roy Chantrell, and Teresa Pellegrino. Asymmetric Assembling of Iron Oxide Nanocubes for Improving Magnetic Hyperthermia Performance. *ACS Nano*, 11(12):12121–12133, November 2017. <https://dx.doi.org/10.1021/acsnano.7b05182>.

[90] T. L. Gilbert and J. M. Kelly. Anomalous rotational damping in ferromagnetic sheets. Proc. Conference on Magnetism and Magnetic Materials, Washington, DC, Oct 1955.

[91] A. A. Pervishko, M. I. Baglai, O. Eriksson, and D. Yudin. Another view on Gilbert damping in two-dimensional ferromagnets. *Scientific Reports*, 8(1), November 2018. <https://dx.doi.org/10.1038/s41598-018-35517-x>.

[92] H. T. Simensen, A. Kamra, R. E. Troncoso, and A. Brataas. Magnon decay theory of Gilbert damping in metallic antiferromagnets. *Physical Review B*, 101(2), January 2020. <https://dx.doi.org/10.1103/physrevb.101.020403>.

[93] J.-M. Beaujour, D. Ravelosona, I. Tudosa, E. E. Fullerton, and A. D. Kent. Ferromagnetic resonance linewidth in ultrathin films with perpendicular magnetic anisotropy. *Physical Review B*, 80(18), November 2009. <https://dx.doi.org/10.1103/physrevb.80.180415>.

[94] M. Beg, M. Albert, M-A. Bisotti, D. Cortés-Ortuño, W. Wang, R. Carey, M. Vousden, O. Hovorka, C. Ciccarelli, C. S. Spencer, C. H. Marrows, and H. Fangohr. Dynamics of skyrmionic states in confined helimagnetic nanostructures. *Physical Review B*, 95(1), January 2017. <https://dx.doi.org/10.1103/physrevb.95.014433>.

[95] M-A. Bisotti, D. Cortés-Ortuño, Pepper R., W. Wang, M. Beg, T. Kluyver, and H. Fangohr. Fidimag – a finite difference atomistic and micromagnetic simulation package. *Journal of Open Research Software*, 6, 2018. <https://dx.doi.org/10.5334/jors.223>.

[96] Ryan Pepper, Weiwei Wang, Marc-Antonio Bisotti, David Ignacio Cortes, Hans Fangohr, Thomas Kluyver, Mark Vousden, Marijan Beg, Rlc2v07, and

Oliver Laslett. computationalmodelling/fidimag: v3.0, 2020. <https://dx.doi.org/10.5281/ZENODO.3841935>.

[97] M. Mansuripur and R. Giles. Demagnetizing field computation for dynamic simulation of the magnetization reversal process. *IEEE Transactions on Magnetics*, 24(6):2326–2328, 1988. ISSN 19410069. <https://dx.doi.org/10.1109/20.92100>.

[98] U. Nowak. Thermally Activated Reversal In Magnetic Nanostructures. In *Annual Reviews Of Computational Physics IX*, pages 105–151. World Scientific, April 2001. https://dx.doi.org/10.1142/9789812811578_0002.

[99] J.W. Eastwood. Optimal particle-mesh algorithms. *Journal of Computational Physics*, 18(1):1–20, May 1975. [https://dx.doi.org/10.1016/0021-9991\(75\)90099-6](https://dx.doi.org/10.1016/0021-9991(75)90099-6).

[100] J. Tranchida, S.J. Plimpton, P. Thibaudeau, and A.P. Thompson. Massively parallel symplectic algorithm for coupled magnetic spin dynamics and molecular dynamics. *Journal of Computational Physics*, 372:406–425, November 2018. <https://dx.doi.org/10.1016/j.jcp.2018.06.042>.

[101] J. W. Cooley and J. W. Tukey. An Algorithm for the Machine Calculation of Complex Fourier Series. *Mathematics of Computation*, 19(90):297, 1965. ISSN 00255718. <https://dx.doi.org/10.2307/2003354>.

[102] J. E. Miltat and M. J. Donahue. Numerical Micromagnetics: Finite Difference Methods, Dec 2007. <https://dx.doi.org/10.1002/9780470022184.hmm202>.

[103] M.J. Donahue and D.G. Porter. Exchange energy formulations for 3d micromagnetics. *Physica B: Condensed Matter*, 343(1-4):177–183, January 2004. <https://dx.doi.org/10.1016/j.physb.2003.08.090>.

[104] A. Vansteenkiste, J. Leliaert, M. Dvornik, M. Helsen, F. Garcia-Sanchez, and B. Van Waeyenberge. The design and verification of MuMax3. *AIP Advances*, 4(10), 2014. ISSN 21583226. <https://dx.doi.org/10.1063/1.4899186>.

[105] A. J. Newell, W. Williams, and D. J. Dunlop. A generalization of the demagnetizing tensor for nonuniform magnetization. *Journal of Geophysical Research*, 98(B6):9551–9555, 1993. ISSN 0148-0227. <https://dx.doi.org/10.1029/93JB00694>.

[106] M.J. Donahue. Accurate computation of the demagnetization tensor. Proc. 7th Symposia on Hysteresis Modelling and Micromagnetics, Gaithersburg, MD, 2007. "https://math.nist.gov/~MDonahue/talks/hmm2007-MBO-03-accurate_demag.pdf".

[107] D. Chernyshenko and H. Fangohr. Computing the demagnetizing tensor for finite difference micromagnetic simulations via numerical integration. *Journal of Magnetism and Magnetic Materials*, 381:440–445, May 2015. <https://dx.doi.org/10.1016/j.jmmm.2015.01.013>.

[108] H. P. Langtangen and A. Logg. Solving PDEs in Minutes - The FEniCS Tutorial Volume I. *Manual*, I, 2016. <https://dx.doi.org/10.1007/978-3-319-52462-7>.

[109] H.P. Langtangen and K.A. Mardal. *Introduction to Numerical Methods for Variational Problems*. Texts in Computational Science and Engineering. Springer International Publishing, 2019. ISBN 9783030237882. <https://books.google.co.uk/books?id=X0WyDwAAQBAJ>.

[110] D. Fredkin and T. Koehler. Numerical micromagnetics by the finite element method. *IEEE Transactions on Magnetics*, 23(5):3385–3387, September 1987. <https://dx.doi.org/10.1109/tmag.1987.1065578>.

[111] W. Chen, D. R. Fredkin, and T. R. Koehler. A new finite element method in micromagnetics. *IEEE Transactions on Magnetics*, 29(3):2124–2128, 1993. ISSN 00189464. <https://dx.doi.org/10.1109/20.221033>.

[112] W. Scholz, J. Fidler, T. Schrefl, D. Suess, R. Dittrich, H. Forster, and V. Tsiantos. Scalable parallel micromagnetic solvers for magnetic nanostructures. *Computational Materials Science*, 28(2):366–383, October 2003. [https://dx.doi.org/10.1016/s0927-0256\(03\)00119-8](https://dx.doi.org/10.1016/s0927-0256(03)00119-8).

[113] T. Fischbacher and H. Fangohr. Continuum multi-physics modeling with scripting languages: the Nsim simulation compiler prototype for classical field theory, 2009. <https://arxiv.org/abs/0907.1587>.

[114] H. Fangohr, M. Albert, and M. Franchin. Nmag micromagnetic simulation tool. In *Proceedings of the International Workshop on Software Engineering for Science - SE4Science '16*. ACM Press, 2016. <https://dx.doi.org/10.1145/2897676.2897677>.

[115] M-A. Bisotti, M. Beg, W. Wang, M. Albert, Dmitri Chernyshenko, David

Cortés-Ortuño, R. A. Pepper, M. Voussden, R. Carey, H. Fuchs, A. Johansen, G. Balaban, L. Breth, T. Kluyver, and H. Fangohr. FinMag: finite-element micromagnetic simulation tool, 2018. <https://dx.doi.org/10.5281/zenodo.1216011>.

[116] Claas Abert, Lukas Exl, Florian Bruckner, André Drews, and Dieter Suess. magnum.fe: A micromagnetic finite-element simulation code based on FEniCS. *Journal of Magnetism and Magnetic Materials*, 345:29–35, November 2013. <https://dx.doi.org/10.1016/j.jmmm.2013.05.051>.

[117] D.R. Fredkin and T.R. Koehler. Hybrid method for computing demagnetizing fields. *IEEE Transactions on Magnetics*, 26(2):415–417, 1990. ISSN 00189464. <https://dx.doi.org/10.1109/20.106342>.

[118] J. D. Jackson. *Classical Electrodynamics*. John Wiley & Sons, New York, 3rd edition, 1999. ISBN 0-471-30932-X.

[119] C. Abert, L. Exl, G. Selke, A. Drews, and T. Schrefl. Numerical methods for the stray-field calculation: A comparison of recently developed algorithms. *Journal of Magnetism and Magnetic Materials*, 326:176–185, January 2013. <https://dx.doi.org/10.1016/j.jmmm.2012.08.041>.

[120] L. Nagy. *Parallelisation of micromagnetic simulations*. PhD thesis, University of Edinburgh, 2016. <http://hdl.handle.net/1842/20433>.

[121] W. Wang. *Computer simulation studies of complex magnetic materials*. PhD thesis, University of Southampton, 2015. <http://eprints.soton.ac.uk/id/eprint/386147>.

[122] P. Palmesi, L. Exl, F. Bruckner, C. Abert, and D. Suess. Highly parallel demagnetization field calculation using the fast multipole method on tetrahedral meshes with continuous sources. *Journal of Magnetism and Magnetic Materials*, 442:409–416, November 2017. <https://dx.doi.org/10.1016/j.jmmm.2017.06.128>.

[123] J-C. Toussaint. FeeLLGood - An Open-Source Finite-Element Micromagnetic Software, 2018. <https://github.com/feellgood/FeeLLGood>.

[124] V. D. Tsiantos, T. Schrefl, J. Fidler, and A. Bratsos. Cost-effective way to speed up micromagnetic simulations in granular media. *Applied Numerical Mathematics*, 39(2):191–204, November 2001. [https://dx.doi.org/10.1016/s0168-9274\(01\)00097-6](https://dx.doi.org/10.1016/s0168-9274(01)00097-6).

[125] D. Shepherd, J. Miles, M. Heil, and M. Mihajlovic. Discretization-Induced Stiffness in Micromagnetic Simulations. *IEEE Transactions on Magnetics*, 50(11):1–4, November 2014. <https://dx.doi.org/10.1109/tmag.2014.2325494>.

[126] M.J. Donahue and D.G. Porter. OOMMF User’s Guide, Version 1.0. Technical report, Interagency Report NISTIR 6376, National Institute of Standards and Technology, Gaithersburg, MD, 1999. <http://math.nist.gov/oommf>.

[127] 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*, 209(2):730–753, November 2005. <https://dx.doi.org/10.1016/j.jcp.2005.04.001>.

[128] L. Exl, N. J. Mauser, T. Schrefl, and D. Suess. The extrapolated explicit midpoint scheme for variable order and step size controlled integration of the landau–lifschitz–gilbert equation. *Journal of Computational Physics*, 346:14–24, October 2017. <https://dx.doi.org/10.1016/j.jcp.2017.06.005>.

[129] D. Praetorius, M. Ruggeri, and B. Stiftner. Convergence of an implicit–explicit midpoint scheme for computational micromagnetics. *Computers & Mathematics with Applications*, 75(5):1719–1738, March 2018. <https://dx.doi.org/10.1016/j.camwa.2017.11.028>.

[130] L. Exl, S. Bance, F. Reichel, T. Schrefl, H. P. Stimming, and N. J. Mauser. LaBonte’s method revisited: An effective steepest descent method for micromagnetic energy minimization. *Journal of Applied Physics*, 115(17):17D118, May 2014. <https://dx.doi.org/10.1063/1.4862839>.

[131] P. F. Bessarab, V. M. Uzdin, and H. Jónsson. Method for finding mechanism and activation energy of magnetic transitions, applied to skyrmion and antivortex annihilation. *Computer Physics Communications*, 196:335–347, 2015. ISSN 0010-4655. <https://dx.doi.org/http://dx.doi.org/10.1016/j.cpc.2015.07.001>.

[132] D. Cortés-Ortuño, W. Wang, M. Beg, R. A. Pepper, M-A. Bisotti, R. Carey, M. Vousden, T. Kluyver, O. Hovorka, and H. Fangohr. Thermal stability and topological protection of skyrmions in nanotacks. *Scientific Reports*, 7(1):4060, 2017. ISSN 2045-2322. <https://dx.doi.org/10.1038/s41598-017-03391-8>.

[133] R. Dittrich, T. Schrefl, D. Suess, W. Scholz, H. Forster, and J. Fidler. A path method for finding energy barriers and minimum energy paths in complex micromagnetic systems. *Journal of Magnetism and Magnetic Materials*, 250: 12–19, September 2002. [https://dx.doi.org/10.1016/s0304-8853\(02\)00388-8](https://dx.doi.org/10.1016/s0304-8853(02)00388-8).

[134] P. F. Bessarab. Comment on "path to collapse for an isolated Néel skyrmion. *Physical Review B - Condensed Matter and Materials Physics*, 95(13):2–3, 2017. ISSN 1550235X. <https://dx.doi.org/10.1103/PhysRevB.95.136401>.

[135] D. Cortés-Ortuño, N. Romming, M. Beg, K. von Bergmann, A. Kubetzka, O. Hovorka, H. Fangohr, and R. Wiesendanger. Nanoscale magnetic skyrmions and target states in confined geometries. *Physical Review B*, 99(21), June 2019. <https://dx.doi.org/10.1103/physrevb.99.214408>.

[136] M. Menarini, R. Medapalli, E. E. Fullerton, and V. Lomakin. Micromagnetic simulation of THz signals in antiferromagnetic FeRh by sub-picosecond thermal pulses. *AIP Advances*, 9(3):035040, March 2019. <https://dx.doi.org/10.1063/1.5080170>.

[137] A. Witt, K. Fabian, and U. Bleil. Three-dimensional micromagnetic calculations for naturally shaped magnetite: Octahedra and magnetosomes. *Earth and Planetary Science Letters*, 233(3-4):311–324, May 2005. <https://dx.doi.org/10.1016/j.epsl.2005.01.043>.

[138] R. Chang, S. Li, M. V. Lubarda, B. Livshitz, and V. Lomakin. FastMag: Fast micromagnetic simulator for complex magnetic structures (invited). *Journal of Applied Physics*, 109(7):07D358, April 2011. <https://dx.doi.org/10.1063/1.3563081>.

[139] P. Ó Conbhúí, W. Williams, K. Fabian, P. Ridley, L. Nagy, and A. R. Muxwworthy. MERRILL: Micromagnetic earth related robust interpreted language laboratory. *Geochemistry, Geophysics, Geosystems*, 19(4):1080–1106, April 2018. <https://dx.doi.org/10.1002/2017gc007279>.

[140] L. Nagy, W. Williams, and M. Valdez-Grijalva. Micromag (A micromagnetics code for high performance computing), 2019. <https://dx.doi.org/10.5281/ZENODO.2584378>.

[141] L. Nagy, W. Williams, L. Tauxe, and A. R. Muxwworthy. From Nano to Micro: Evolution of Magnetic Domain Structures in Multidomain Magnetite.

Geochemistry, Geophysics, Geosystems, 20(6):2907–2918, June 2019. <https://dx.doi.org/10.1029/2019gc008319>.

[142] Martin S. Alnæs, Jan Blechta, Johan Hake, August Johansson, Benjamin Kehlet, Anders Logg, Chris Richardson, Johannes Ring, Marie E. Rognes, and Garth N. Wells. The fenics project version 1.5. *Archive of Numerical Software*, 3(100), 2015. <https://dx.doi.org/10.11588/ans.2015.100.20553>.

[143] Anders Logg, Kent-Andre Mardal, Garth N. Wells, et al. *Automated Solution of Differential Equations by the Finite Element Method*. Springer, 2012. ISBN 978-3-642-23098-1. <https://dx.doi.org/10.1007/978-3-642-23099-8>.

[144] Anders Logg and Garth N. Wells. DOLFIN: Automated Finite Element Computing. *ACM Transactions on Mathematical Software*, 37(2), 2010. <https://dx.doi.org/10.1145/1731022.1731030>.

[145] Anders Logg, Garth N. Wells, and Johan Hake. *DOLFIN: a C++/Python Finite Element Library*, chapter 10. Springer, 2012.

[146] Robert C. Kirby and Anders Logg. A Compiler for Variational Forms. *ACM Transactions on Mathematical Software*, 32(3), 2006. <https://dx.doi.org/10.1145/1163641.1163644>.

[147] Martin S. Alnæs, Anders Logg, and Kent-Andre Mardal. *UFC: a Finite Element Code Generation Interface*, chapter 16. Springer, 2012.

[148] Robert C. Kirby. Algorithm 839: FIAT, a New Paradigm for Computing Finite Element Basis Functions. *ACM Transactions on Mathematical Software*, 30(4):502–516, 2004. <https://dx.doi.org/10.1145/1039813.1039820>.

[149] Robert C. Kirby. FIAT: numerical construction of finite element basis functions. In *Automated Solution of Differential Equations by the Finite Element Method*, pages 247–255. Springer Berlin Heidelberg, 2012. https://dx.doi.org/10.1007/978-3-642-23099-8_13.

[150] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.13, Argonne National Laboratory, 2020. <https://www.mcs.anl.gov/petsc>.

[151] Martin S. Alnæs, Anders Logg, Kent-Andre Mardal, Ola Skavhaug, and Hans Petter Langtangen. Unified Framework for Finite Element Assembly. *International Journal of Computational Science and Engineering*, 4(4):231–244, 2009. <https://dx.doi.org/10.1504/IJCSE.2009.029160>.

[152] J. Leliaert and J. Mulders. Tomorrow’s micromagnetic simulations. *Journal of Applied Physics*, 125(18):180901, May 2019. <https://dx.doi.org/10.1063/1.5093730>.

[153] Wojciech Śmigaj, Timo Betcke, Simon Arridge, Joel Phillips, and Martin Schweiger. Solving boundary integral problems with BEM++. *ACM Transactions on Mathematical Software*, 41(2):1–40, February 2015. <https://dx.doi.org/10.1145/2590830>.

[154] Josh Barnes and Piet Hut. A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature*, 324(6096):446–449, December 1986. <https://dx.doi.org/10.1038/324446a0>.

[155] P. Fortin, E. Athanassoula, and J.-C. Lambert. Comparisons of different codes for galactic N-body simulations. *Astronomy & Astrophysics*, 531:A120, June 2011. <https://dx.doi.org/10.1051/0004-6361/201015933>.

[156] Tingyu W., S. K. Layton, and L. A. Barba. Inexact Krylov iterations and relaxation strategies with fast-multipole boundary element method, 2015. <https://arxiv.org/abs/1506.05957>.

[157] W. Dehnen. A fast multipole method for stellar dynamics. *Computational Astrophysics and Cosmology*, 1(1), 2014. <https://dx.doi.org/10.1186/s40668-014-0001-7>.

[158] Z. V. Gareeva and K. Y. Guslienko. Collective magnetic skyrmion gyrotrropic modes in a dot chain. *Journal of Physics Communications*, 2(3):035009, 2018. <https://dx.doi.org/10.1088/2399-6528/aab169>.

[159] M. S. Alnæs, A. Logg, K-A. Mardal, O. Skavhaug, and H. P. Langtangen. Unified Framework for Finite Element Assembly. *International Journal of Computational Science and Engineering*, 4(4):231–244, 2009. <https://dx.doi.org/10.1504/IJCSE.2009.029160>.

[160] A. Logg and G. N. Wells. DOLFIN: Automated Finite Element Computing. *ACM Transactions on Mathematical Software (TOMS)*, 37(2):2, 2010. ISSN 00983500. <https://dx.doi.org/10.1145/1731022.1731030>.

[161] F. Rathgeber, D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. T. Mcrae, G-T. Bercea, G. R. Markall, and P. H. J. Kelly. Firedrake. *ACM Transactions on Mathematical Software*, 43(3):1–27, 2016. <https://dx.doi.org/10.1145/2998441>.

[162] D. J. Lusher, S. P. Jammy, and N. D. Sandham. Shock-wave/boundary-layer interactions in the automatic source-code generation framework OpenSBLI. *Computers & Fluids*, 173:17–21, 2018. <https://dx.doi.org/10.1016/j.compfluid.2018.03.081>.

[163] G. R. Mudalige, I. Z. Reguly, S. P. Jammy, C. T. Jacobs, M. B. Giles, and N. D. Sandham. Large-scale performance of a DSL-based multi-block structured-mesh application for direct numerical simulation. *Journal of Parallel and Distributed Computing*, 131:130–146, 2019. <https://dx.doi.org/10.1016/j.jpdc.2019.04.019>.

[164] C. T. Jacobs, S. P. Jammy, and N. D. Sandham. OpenSBLI: A framework for the automated derivation and parallel execution of finite difference solvers on a range of computer architectures. *Journal of Computational Science*, 18: 12–23, 2017. <https://dx.doi.org/10.1016/j.jocs.2016.11.001>.

[165] I. Z. Reguly, G. R. Mudalige, and M. B. Giles. Loop tiling in large-scale stencil codes at run-time with OPS. *IEEE Transactions on Parallel and Distributed Systems*, 29(4):873–886, 2018. <https://dx.doi.org/10.1109/tpds.2017.2778161>.

[166] F. Luporini, M. Lange, M. Louboutin, N. Kukreja, J. Hückelheim, C. Yount, P. Witte, P. H. J. Kelly, F. J. Herrmann, and G. J. Gorman. Architecture and performance of Devito, a system for automated stencil computation . *CoRR*, abs/1807.03032, 2018. <https://dx.doi.org/10.1145/3374916>.

[167] J. P. Coles and M. Masella. The fast multipole method and point dipole moment polarizable force fields. *The Journal of Chemical Physics*, 142(2): 024109, 2015. ISSN 0021-9606. <https://dx.doi.org/10.1063/1.4904922>.

[168] Jonathan P. Coles and Rebekka Bieri. An optimizing symbolic algebra approach for generating fast multipole method operators. *Computer Physics Communications*, 251:107081, June 2020. <https://dx.doi.org/10.1016/j.cpc.2019.107081>.

[169] J. Applequist. Fundamental relationships in the theory of electric multipole moments and multipole polarizabilities in static fields. *Chemical Physics*, 85

(2):279–290, 1984. [https://dx.doi.org/10.1016/0301-0104\(84\)85039-9](https://dx.doi.org/10.1016/0301-0104(84)85039-9).

[170] J. Applequist. Traceless cartesian tensor forms for spherical harmonic functions: new theorems and applications to electrostatics of dielectric media. *Journal of Physics A: Mathematical and General*, 22(20):4303–4330, 1989. ISSN 0305-4470. <https://dx.doi.org/10.1088/0305-4470/22/20/011>.

[171] Ryan Alexander Pepper and Hans Fangohr. fmmgen: Automatic Code Generation of Operators for Cartesian Fast Multipole and Barnes-Hut Methods, 2020. <https://arxiv.org/abs/2005.12351>.

[172] Ryan Alexander Pepper and Hans Fangohr. Dataset for "Fmmgen: Automatic Code Generation of Operators for Cartesian Fast Multipole and Barnes-Hut Methods", 2020. <https://dx.doi.org/10.5281/ZENODO.3842584>.

[173] A. Meurer, C. P. Smith, M. Paprocki, O. Čertík, S. B. Kirpichev, M. Rocklin, A. Kumar, S. Ivanov, J. K. Moore, S. Singh, T. Rathnayake, S. Vig, B. E. Granger, R. P. Muller, F. Bonazzi, H. Gupta, S. Vats, F. Johansson, F. Pedregosa, M. J. Curry, A. R. Terrel, Š. Roučka, A. Saboo, I. Fernando, S. Kulal, R. Cimrman, and A. Scopatz. SymPy: symbolic computing in Python. *PeerJ Computer Science*, 3:e103, 2017. ISSN 2376-5992. <https://dx.doi.org/10.7717/peerj-cs.103>.

[174] J. Cipriani and B. Silvi. Cartesian expressions for electric multipole moment operators. *Molecular Physics*, 45(2):259–272, 1982. <https://dx.doi.org/10.1080/00268978200100211>.

[175] M. Challacombe, E. Schwegler, and J. Almlöf. Recurrence relations for calculation of the Cartesian multipole tensor. *Chemical Physics Letters*, 241 (1-2):67–72, 1995. [https://dx.doi.org/10.1016/0009-2614\(95\)00597-w](https://dx.doi.org/10.1016/0009-2614(95)00597-w).

[176] B. Wicht. C11 performance tip: When to use std::pow?, 2017. <https://baptiste-wicht.com/posts/2017/09/cpp11-performance-tip-when-to-use-std-pow.html>.

[177] U. Lauther. Multipole-based force approximation revisited – a simple but fast implementation using a dynamized enclosing-circle-enhanced k-d-tree. In *Graph Drawing*, pages 20–29. Springer Berlin Heidelberg, 2006. https://dx.doi.org/10.1007/978-3-540-70904-6_4.

[178] R. H. C. Lopes, I. D. Reid, and P. R. Hobson. A well-separated pairs decomposition algorithm for k-d trees implemented on multi-core architectures. *Journal of Physics: Conference Series*, 513(5):052011, June 2014. <https://dx.doi.org/10.1088/1742-6596/513/5/052011>.

[179] Jinshi Zhu, Yongmei Lei, and Jianchen Shan. Parallel FMM algorithm based on space decomposition. In *2010 Ninth International Conference on Grid and Cloud Computing*. IEEE, November 2010. <https://dx.doi.org/10.1109/gcc.2010.43>.

[180] Rio Yokota and Lorena A Barba. A tuned and scalable fast multipole method as a preeminent algorithm for exascale systems. *The International Journal of High Performance Computing Applications*, 26(4):337–346, January 2012. <https://dx.doi.org/10.1177/1094342011429952>.

[181] R. Yokota. An FMM based on dual tree traversal for many-core architectures. *Journal of Algorithms & Computational Technology*, 7(3):301–324, 2013. <https://dx.doi.org/10.1260/1748-3018.7.3.301>.

[182] N. A. Simakov, M. D. Innus, M. D. Jones, J. P. White, S. M. Gallo, R. L. DeLeon, and T. R. Furlani. Effect of Meltdown and Spectre Patches on the Performance of HPC Applications, 2018. <https://arxiv.org/abs/arXiv:1801.04329>.

[183] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C (2Nd Ed.): The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 1992. ISBN 0-521-43108-5.

[184] O. Hovorka, J. Barker, G. Friedman, and R. W. Chantrell. Role of geometrical symmetry in thermally activated processes in clusters of interacting dipolar moments. *Physical Review B*, 89(10), March 2014. <https://dx.doi.org/10.1103/physrevb.89.104410>.

[185] J. Leliaert, A. Vansteenkiste, A. Coene, L. Dupré, and B. Van Waeyenberge. Vinamax: a macrospin simulation tool for magnetic nanoparticles. *Medical & Biological Engineering & Computing*, 53(4):309–317, January 2015. <https://dx.doi.org/10.1007/s11517-014-1239-6>.

[186] Intel. Intel Intrinsics Guide. <https://software.intel.com/sites/landingpage/IntrinsicsGuide/>, 2019.

[187] A. Fogg. vectorclass. <https://github.com/vectorclass>, 2019.

- [188] QuantStack. xsimd, 2019. <https://github.com/QuantStack/xsimd>.
- [189] W. Dehnen. A Hierarchical (N) Force Calculation Algorithm. *Journal of Computational Physics*, 179(1):27–42, 2002. <https://dx.doi.org/10.1006/jcph.2002.7026>.
- [190] J. P. Coles and R. Bieri. A Fully Traceless Cartesian Multipole Formulation for the Distributed Fast Multipole Method, 2018. <https://arxiv.org/abs/1811.06332>.
- [191] J. P. Coles. Mosaic, 2019. <https://github.com/jpcoles/mosaic>.
- [192] T. H. R. Skyrme. A unified field theory of mesons and baryons. *Nuclear Physics*, 31:556–569, 1962. ISSN 00295582. [https://dx.doi.org/10.1016/0029-5582\(62\)90775-7](https://dx.doi.org/10.1016/0029-5582(62)90775-7).
- [193] H-B. Braun. Topological effects in nanomagnetism: from superparamagnetism to chiral quantum solitons. *Advances in Physics*, 61(1):1–116, 2012. ISSN 0001-8732. <https://dx.doi.org/10.1080/00018732.2012.663070>.
- [194] A. N. Bogdanov and D. A. Yablonskii. Thermodynamically stable "vortices" in magnetically ordered crystals. The mixed state of magnets. *Zh. Eksp. Teor. Fiz.*, 95(1):178, 1989. ISSN 0044-4510. [https://dx.doi.org/10.1016/0304-8853\(94\)90046-9](https://dx.doi.org/10.1016/0304-8853(94)90046-9).
- [195] M. Tinkham. *Introduction to Superconductivity*. McGraw-Hill Book Co., 2nd edition, 1996. ISBN 978-0-486-43503-9. <https://store.doverpublications.com/0486435032.html>.
- [196] A. Crépieux and C. Lacroix. Dzyaloshinsky–Moriya interactions induced by symmetry breaking at a surface. *Journal of Magnetism and Magnetic Materials*, 182(3):341–349, 1998. ISSN 03048853. [https://dx.doi.org/10.1016/S0304-8853\(97\)01044-5](https://dx.doi.org/10.1016/S0304-8853(97)01044-5).
- [197] K. Xia, W. Zhang, M. Lu, and H. Zhai. Noncollinear interlayer exchange coupling caused by interface spin-orbit interaction. *Physical Review B*, 55(18):12561–12565, 1997. ISSN 0163-1829. <https://dx.doi.org/10.1103/PhysRevB.55.12561>.
- [198] A. N. Bogdanov and U. K. Rößler. Chiral Symmetry Breaking in Magnetic Thin Films and Multilayers. *Physical Review Letters*, 87(3):037203, 2001. ISSN 0031-9007. <https://dx.doi.org/10.1103/PhysRevLett.87.037203>.

[199] S. Muhlbauer, B. Binz, F. Jonietz, C. Pfleiderer, A. Rosch, A. Neubauer, R. Georgii, and P. Boni. Skyrmion Lattice in a Chiral Magnet. *Science*, 323(5916):915–919, 2009. ISSN 0036-8075. <https://dx.doi.org/10.1126/science.1166767>.

[200] A. Neubauer, C. Pfleiderer, B. Binz, A. Rosch, R. Ritz, P. G. Niklowitz, and P. Boni. Topological hall effect in the a phase of MnSi. *Physical Review Letters*, 102(18):1–4, 2009. ISSN 00319007. <https://dx.doi.org/10.1103/PhysRevLett.102.186602>.

[201] X. Z. Yu, Y. Onose, N. Kanazawa, J. H. Park, J. H. Han, Y. Matsui, N. Nagaosa, and Y. Tokura. Real-space observation of a two-dimensional skyrmion crystal. *Nature*, 465(7300):901–904, June 2010. <https://dx.doi.org/10.1038/nature09124>.

[202] X. Z. Yu, N. Kanazawa, Y. Onose, K. Kimoto, W. Z. Zhang, S. Ishiwata, Y. Matsui, and Y. Tokura. Near room-temperature formation of a skyrmion crystal in thin-films of the helimagnet FeGe. *Nature Materials*, 10(2):106–109, 2011. ISSN 1476-1122. <https://dx.doi.org/10.1038/nmat2916>.

[203] F. N. Rybakov, A. B. Borisov, and A. N. Bogdanov. Three-dimensional skyrmion states in thin films of cubic helimagnets. *Physical Review B - Condensed Matter and Materials Physics*, 87(9):1–5, 2013. ISSN 10980121. <https://dx.doi.org/10.1103/PhysRevB.87.094424>.

[204] A. Tonomura, X. Yu, K. Yanagisawa, T. Matsuda, Y. Onose, N. Kanazawa, H. S. Park, and Y. Tokura. Real-space observation of skyrmion lattice in helimagnet MnSi thin samples. *Nano Letters*, 12(3):1673–1677, 2012. ISSN 15306984. <https://dx.doi.org/10.1021/nl300073m>.

[205] S. Seki, X. Z. Yu, S. Ishiwata, and Y. Tokura. Observation of Skyrmions in a Multiferroic Material. *Science*, 336(6078):198–201, April 2012. <https://dx.doi.org/10.1126/science.1214143>.

[206] J. C. Loudon, A. O. Leonov, A. N. Bogdanov, M. Ciomaga Hatnean, and G. Balakrishnan. Direct observation of attractive skyrmions and skyrmion clusters in the cubic helimagnet cu2oseo3. *Physical Review B*, 97(13), April 2018. <https://dx.doi.org/10.1103/physrevb.97.134403>.

[207] G. Finocchio, F. Büttner, R. Tomasello, M. Carpentieri, and M. Kläui. Magnetic skyrmions: from fundamental to applications. *Journal of Physics D: Applied Physics*, 51(38), December 2018. <https://dx.doi.org/10.1088/1361-6463/aae001>.

Applied Physics, 49(42):423001, September 2016. <https://dx.doi.org/10.1088/0022-3727/49/42/423001>.

[208] W. Jiang, G. Chen, K. Liu, J. Zang, S. G.E. te Velthuis, and A. Hoffmann. Skyrmions in magnetic multilayers. *Physics Reports*, 704:1–49, August 2017. <https://dx.doi.org/10.1016/j.physrep.2017.08.001>.

[209] K. Everschor-Sitte, J. Masell, R. M. Reeve, and M. Kläui. Perspective: Magnetic skyrmions—overview of recent progress in an active research field. *Journal of Applied Physics*, 124(24):240901, December 2018. <https://dx.doi.org/10.1063/1.5048972>.

[210] S. Heinze, K. von Bergmann, M. Menzel, J. Brede, A. Kubetzka, R. Wiesendanger, G. Bihlmayer, and S. Blügel. Spontaneous atomic-scale magnetic skyrmion lattice in two dimensions. *Nature Physics*, 7(9):713–718, 2011. ISSN 1745-2473. <https://dx.doi.org/10.1038/nphys2045>.

[211] C. Moreau-Luchaire, C. Moutafis, N. Reyren, J. Sampaio, C. A. F. Vaz, N. Van Horne, K. Bouzehouane, K. Garcia, C. Deranlot, P. Wernicke, P. Wohlhüter, J.-M. George, M. Weigand, J. Raabe, V. Cros, and A. Fert. Additive interfacial chiral interaction in multilayers for stabilization of small individual skyrmions at room temperature. *Nature Nanotechnology*, 11(5):444–448, January 2016. <https://dx.doi.org/10.1038/nnano.2015.313>.

[212] S. Woo, K. Litzius, B. Krüger, M-y. Im, L. Caretta, K. Richter, M. Mann, A. Krone, R. M. Reeve, M. Weigand, P. Agrawal, I. Lemesh, M-A. Mawass, P. Fischer, M. Kläui, and G. S. D. Beach. Observation of room-temperature magnetic skyrmions and their current-driven dynamics in ultrathin metallic ferromagnets. *Nature Materials*, 15(5):501–506, 2016. ISSN 1476-1122. <https://dx.doi.org/10.1038/nmat4593>.

[213] L. Caretta, M. Mann, F. Büttner, K. Ueda, B. Pfau, C. M. Günther, P. Hessing, A. Churikova, C. Klose, M. Schneider, D. Engel, C. Marcus, D. Bono, K. Bagschik, S. Eisebitt, and G. S. D. Beach. Fast current-driven domain walls and small skyrmions in a compensated ferrimagnet. *Nature Nanotechnology*, 13(12):1154–1160, September 2018. <https://dx.doi.org/10.1038/s41565-018-0255-3>.

[214] F. Jonietz, S. Muhlbauer, C. Pfleiderer, A. Neubauer, W. Munzer, A. Bauer, T. Adams, R. Georgii, P. Boni, R. A. Duine, K. Everschor, M. Garst, and A. Rosch. Spin Transfer Torques in MnSi at Ultralow Current Densities.

Science, 330(6011):1648–1651, 2010. ISSN 0036-8075. <https://dx.doi.org/10.1126/science.1195709>.

[215] X. Z. Yu, N. Kanazawa, W.Z. Zhang, T. Nagai, T. Hara, K. Kimoto, Y. Matsui, Y. Onose, and Y. Tokura. Skyrmion flow near room temperature in an ultralow current density. *Nature Communications*, 3:988, 2012. ISSN 2041-1723. <https://dx.doi.org/10.1038/ncomms1990>.

[216] P. Heistracher, C. Abert, F. Bruckner, C. Vogler, and D. Suess. GPU-accelerated atomistic energy barrier calculations of skyrmion annihilations. *IEEE Transactions on Magnetics*, 54(11):1–5, November 2018. <https://dx.doi.org/10.1109/tmag.2018.2847446>.

[217] Durgesh Kumar, Tianli Jin, S. Al Risi, Rachid Sbiaa, W. S. Lew, and S. N. Piramanayagam. Domain Wall Motion Control for Racetrack Memory Applications. *IEEE Transactions on Magnetics*, 55(3):1–8, March 2019. <https://dx.doi.org/10.1109/tmag.2018.2876622>.

[218] Jan Müller. Magnetic skyrmions on a two-lane racetrack. *New Journal of Physics*, 19(2):025002, February 2017. <https://dx.doi.org/10.1088/1367-2630/aa5b55>.

[219] Xichao Zhang, Motohiko Ezawa, and Yan Zhou. Thermally stable magnetic skyrmions in multilayer synthetic antiferromagnetic racetracks. *Physical Review B*, 94(6), August 2016. <https://dx.doi.org/10.1103/physrevb.94.064406>.

[220] R Tomasello, V Puliafito, E Martinez, A Manchon, M Ricci, M Carpentieri, and G Finocchio. Performance of synthetic antiferromagnetic racetrack memory: domain wall versus skyrmion. *Journal of Physics D: Applied Physics*, 50(32):325302, July 2017. <https://dx.doi.org/10.1088/1361-6463/aa7a98>.

[221] P. F. Bessarab, G. P. Müller, I. S. Lobanov, F. N. Rybakov, N. S. Kiselev, H. Jónsson, V. M. Uzdin, S. Blügel, L. Bergqvist, and A. Delin. Lifetime of racetrack skyrmions. *Scientific Reports*, 8(1), February 2018. <https://dx.doi.org/10.1038/s41598-018-21623-3>.

[222] Karin Everschor-Sitte and Matthias Sitte. Real-space Berry phases: Skyrmion soccer (invited). *Journal of Applied Physics*, 115(17):172602, May 2014. <https://dx.doi.org/10.1063/1.4870695>.

[223] X. Zhang, W. Cai, X. Zhang, Z. Wang, Z. Li, Y. Zhang, K. Cao, N. Lei, W. Kang, Y. Zhang, H. Yu, Y. Zhou, and W. Zhao. Skyrmions in Magnetic Tunnel Junctions. *ACS Applied Materials & Interfaces*, 10(19):16887–16892, April 2018. <https://dx.doi.org/10.1021/acsami.8b03812>.

[224] S. Wang, J. Tang, W. Wang, L. Kong, M. Tian, and H. Du. Electrical Detection of Magnetic Skyrmions. *Journal of Low Temperature Physics*, 197(3-4):321–336, June 2019. <https://dx.doi.org/10.1007/s10909-019-02202-w>.

[225] M. Leroux, M. J. Stolt, S. Jin, D. V. Pete, C. Reichhardt, and B. Maiorov. Skyrmion Lattice Topological Hall Effect near Room Temperature. *Scientific Reports*, 8(1), October 2018. <https://dx.doi.org/10.1038/s41598-018-33560-2>.

[226] K. Everschor, M. Garst, R. A. Duine, and A. Rosch. Current-induced rotational torques in the skyrmion lattice phase of chiral magnets. *Physical Review B*, 84(6), August 2011. <https://dx.doi.org/10.1103/physrevb.84.064401>.

[227] K. Everschor, M. Garst, B. Binz, F. Jonietz, S. Mühlbauer, C. Pfleiderer, and A. Rosch. Rotating skyrmion lattices by spin torques and field or temperature gradients. *Physical Review B*, 86(5), August 2012. <https://dx.doi.org/10.1103/physrevb.86.054432>.

[228] M. Beg, M. Albert, M-A. Bisotti, D. Cortés-Ortuño, W. Wang, R. Carey, M. Vousden, O. Hovorka, C. Ciccarelli, C. S. Spencer, C. H. Marrows, and H. Fangohr. Dynamics of skyrmionic states in confined helimagnetic nanostructures. *Physical Review B*, 95(1):014433, 2017. ISSN 2469-9950. <https://dx.doi.org/10.1103/PhysRevB.95.014433>.

[229] A. O. Leonov and M. Mostovoy. Edge states and skyrmion dynamics in nanostripes of frustrated magnets. *Nature Communications*, 8(1), February 2017. <https://dx.doi.org/10.1038/ncomms14394>.

[230] J. Müller, A. Rosch, and M. Garst. Edge instabilities and skyrmion creation in magnetic layers. *New Journal of Physics*, 18(6):065006, June 2016. <https://dx.doi.org/10.1088/1367-2630/18/6/065006>.

[231] A. O. Leonov, T. L. Monchesky, N. Romming, A. Kubetzka, A. N. Bogdanov, and R. Wiesendanger. The properties of isolated chiral skyrmions in thin

magnetic films. *New Journal of Physics*, 18(6):065003, May 2016. <https://dx.doi.org/10.1088/1367-2630/18/6/065003>.

[232] A. Aharoni. *Introduction to the Theory of Ferromagnetism*. Oxford University Press Inc., New York, 2nd edition, 2000. ISBN ISBN 0 19 850809.

[233] H. Du, X. Zhao, F. N. Rybakov, A. B. Borisov, S. Wang, J. Tang, C. Jin, C. Wang, W. Wei, N. S. Kiselev, Y. Zhang, R. Che, S. Blügel, and M. Tian. Interaction of Individual Skyrmions in a Nanostructured Cubic Chiral Magnet. *Physical Review Letters*, 120(19), May 2018. <https://dx.doi.org/10.1103/physrevlett.120.197203>.

[234] R. A. Pepper, M. Beg, D. Cortés-Ortuño, T. Kluyver, M-A. Bisotti, R. Carey, M. Vousden, M. Albert, W. Wang, O. Hovorka, and H. Fangohr. Skyrmion states in thin confined polygonal nanostructures. *Journal of Applied Physics*, 123(9):093903, 2018. ISSN 0021-8979. <https://dx.doi.org/10.1063/1.5022567>.

[235] R. A. Pepper, M. Beg, D. Cortes-Ortuno, T. Kluyver, M-A. Bisotti, R. Carey, M. Vousden, M. Albert, W. Wang, O. Hovorka, and H. Fangohr. Dataset for "skyrmion states in thin confined polygonal nanostructures", 2017. <https://dx.doi.org/10.5281/zenodo.1066792>.

[236] F. N. Rybakov, A. B. Borisov, S. Blügel, and N. S. Kiselev. New spiral state and skyrmion lattice in 3D model of chiral magnets. *New Journal of Physics*, 18(4):045002, 2016. ISSN 1367-2630. <https://dx.doi.org/10.1088/1367-2630/18/4/045002>.

[237] M. Vousden, M. Albert, M. Beg, M-A. Bisotti, R. Carey, D. Chernyshenko, D. Cortés-Ortuño, W. Wang, O. Hovorka, C. H. Marrows, and H. Fangohr. Skyrmions in thin films with easy-plane magnetocrystalline anisotropy. *Applied Physics Letters*, 108(13):132406, 2016. ISSN 0003-6951. <https://dx.doi.org/10.1063/1.4945262>.

[238] A. O. Leonov, Y. Togawa, T. L. Monchesky, A. N. Bogdanov, J. Kishine, Y. Kousaka, M. Miyagawa, T. Koyama, J. Akimitsu, Ts Koyama, K. Harada, S. Mori, D. McGrouther, R. Lamb, M. Krajnak, S. McVitie, R. L. Stamps, and K. Inoue. Chiral Surface Twists and Skyrmion Stability in Nanolayers of Cubic Helimagnets. *Physical Review Letters*, 117(8):1–5, 2016. ISSN 10797114. <https://dx.doi.org/10.1103/PhysRevLett.117.087202>.

[239] S. Schneider, D. Wolf, M. J. Stolt, S. Jin, D. Pohl, B. Rellinghaus, M. Schmidt, B. Büchner, S. T. B. Goennenwein, K. Nielsch, and A. Lubk. Induction mapping of the 3d-modulated spin texture of skyrmions in thin helimagnets. *Physical Review Letters*, 120(21), 2018. <https://dx.doi.org/10.1103/physrevlett.120.217201>.

[240] B. Lebech, J. Bernhard, and T. Freltoft. Magnetic structures of cubic FeGe studied by small- angle neutron scattering. *Journal of Physics: Condensed Matter*, 1:6105–6122, 1989. <https://dx.doi.org/10.1088/0953-8984/1/35/010>.

[241] F. Zheng, H. Li, S. Wang, D. Song, C. Jin, W. Wei, A. Kovács, J. Zang, M. Tian, Y. Zhang, H. Du, and R. E. Dunin-Borkowski. Direct Imaging of a Zero-Field Target Skyrmion and Its Polarity Switch in a Chiral Magnetic Nanodisk. *Physical Review Letters*, 119(19):197205, 2017. ISSN 0031-9007. <https://dx.doi.org/10.1103/PhysRevLett.119.197205>.

[242] 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*, 31(3):363–396, sep 2005. <https://dx.doi.org/10.1145/1089014.1089020>.

[243] A.R. Aranda, A. Hierro-Rodriguez, G.N. Kakazei, O. Chubykalo-Fesenko, and K.Y. Guslienko. Magnetic skyrmion size and stability in ultrathin nanodots accounting Dzyaloshinskii-Moriya exchange interaction. *Journal of Magnetism and Magnetic Materials*, 465:471–479, November 2018. <https://dx.doi.org/10.1016/j.jmmm.2018.05.074>.

[244] H. Du, W. Ning, M. Tian, and Y. Zhang. Magnetic vortex with skyrmionic core in a thin nanodisk of chiral magnets. *EPL (Europhysics Letters)*, 101(3):37001, February 2013. <https://dx.doi.org/10.1209/0295-5075/101/37001>.

[245] F. N. Rybakov, A. B. Borisov, S. Blügel, and N. S. Kiselev. New Type of Stable Particlelike States in Chiral Magnets. *Physical Review Letters*, 115(11), September 2015. <https://dx.doi.org/10.1103/physrevlett.115.117201>.

[246] F. Zheng, F. N. Rybakov, A. B. Borisov, D. Song, S. Wang, Z-A. Li, H. Du, N. S. Kiselev, J. Caron, A. Kovács, M. Tian, Y. Zhang, S. Blügel, and R. E.

Dunin-Borkowski. Experimental observation of chiral magnetic bobbers in b20-type FeGe. *Nature Nanotechnology*, 13(6):451–455, April 2018. <https://dx.doi.org/10.1038/s41565-018-0093-3>.

[247] R.J Enkin and D. J. Dunlop. A micromagnetic study of pseudo single-domain remanence in magnetite. *Journal of Geophysical Research*, 92(B12):12726, 1987. <https://dx.doi.org/10.1029/jb092ib12p12726>.

[248] R. J. Enkin and W. Williams. Three-dimensional micromagnetic analysis of stability in fine magnetic grains. *Journal of Geophysical Research: Solid Earth*, 99(B1):611–618, January 1994. <https://dx.doi.org/10.1029/93jb02637>.

[249] Y. Zhang and H. N. Bertram. Thermal decay in high density disk media. *IEEE Transactions on Magnetics*, 34(5):3786–3793, 1998. <https://dx.doi.org/10.1109/20.718543>.

[250] D. Suess, T. Schrefl, and J. Fidler. Reversal modes, thermal stability and exchange length in perpendicular recording media. *IEEE Transactions on Magnetics*, 37(4):1664–1666, July 2001. <https://dx.doi.org/10.1109/20.950931>.

[251] R. Dittrich, T. Schrefl, A. Thiaville, J. Miltat, V. Tsiantos, and J. Fidler. Comparison of Langevin dynamics and direct energy barrier computation. *Journal of Magnetism and Magnetic Materials*, 272-276:747–749, May 2004. <https://dx.doi.org/10.1016/j.jmmm.2003.11.274>.

[252] D. Suess, T. Schrefl, S. Fähler, M. Kirschner, G. Hrkac, F. Dorfbauer, and J. Fidler. Exchange spring media for perpendicular recording. *Applied Physics Letters*, 87(1):012504, July 2005. <https://dx.doi.org/10.1063/1.1951053>.

[253] E. Paz, F. Garcia-Sanchez, and O. Chubykalo-Fesenko. Numerical evaluation of energy barriers in nano-sized magnetic elements with Lagrange multiplier technique. *Physica B: Condensed Matter*, 403(2-3):330–333, February 2008. <https://dx.doi.org/10.1016/j.physb.2007.08.042>.

[254] S. Rohart, J. Miltat, and A. Thiaville. Reply to “comment on ‘path to collapse for an isolated néel skyrmion’ ”. *Physical Review B*, 95(13), April 2017. <https://dx.doi.org/10.1103/physrevb.95.136402>.

[255] P. F. Bessarab, D. Yudin, D. R. Gulevich, P. Wadley, M. Titov, and O. A.

Tretiakov. Stability and lifetime of antiferromagnetic skyrmions. *Physical Review B*, 99(14), April 2019. <https://dx.doi.org/10.1103/physrevb.99.140411>.

[256] L. Desplat, C. Vogler, J.-V. Kim, R. L. Stamps, and D. Suess. Path sampling for lifetimes of metastable magnetic skyrmions and direct comparison with kramers' method. *Physical Review B*, 101(6), February 2020. <https://dx.doi.org/10.1103/physrevb.101.060403>.

[257] C. Hanneken, A. Kubetzka, K. von Bergmann, and R. Wiesendanger. Pinning and movement of individual nanoscale magnetic skyrmions via defects. *New Journal of Physics*, 18(5):055009, May 2016. <https://dx.doi.org/10.1088/1367-2630/18/5/055009>.

[258] V. M. Uzdin, M. N. Potkina, I. S. Lobanov, P. F. Bessarab, and H. Jónsson. The effect of confinement and defects on the thermal stability of skyrmions. *Physica B: Condensed Matter*, 549:6–9, November 2018. <https://dx.doi.org/10.1016/j.physb.2017.09.040>.

[259] D. Stosic, T. B. Ludermir, and M. V. Milošević. Pinning of magnetic skyrmions in a monolayer co film on pt(111): Theoretical characterization and exemplified utilization. *Physical Review B*, 96(21), December 2017. <https://dx.doi.org/10.1103/physrevb.96.214403>.

[260] Ryan Alexander Pepper, David Cortés-Ortuño, and Hans Fangohr. Data-set for "Edge Defects and Energy Barriers in Skyrmion Monolayers", 2020. <https://dx.doi.org/10.5281/ZENODO.4058558>.

[261] Ryan Pepper. Destruction of a skyrmion via a triangular edge defect, 2020. https://youtu.be/z-NHjRK3x_0.

[262] Robert C. Wetherhold and Victor H. Guerrero. Effect of substrate anisotropy on magnetic state and stress state in magnetostrictive multilayers. *Journal of Magnetism and Magnetic Materials*, 262(2):218–229, June 2003. [https://dx.doi.org/10.1016/s0304-8853\(02\)01412-9](https://dx.doi.org/10.1016/s0304-8853(02)01412-9).

[263] E. Feldtkeller. Mikromagnetisch stetige und unstetige Magnetisierungskonfigurationen. *Zeitschrift für Angewandte Physik*, 19(6):530–536, 1965.

[264] E. Feldtkeller. Continuous and Singular Micromagnetic Configurations. *IEEE Transactions on Magnetics*, 53(10):1–8, Oct 2017. <https://dx.doi.org/10.1109/tmag.2017.2737403>.

- [265] W. Döring. Point Singularities in Micromagnetism. *Journal of Applied Physics*, 39(2):1006–1007, February 1968. <https://dx.doi.org/10.1063/1.1656144>.
- [266] A Hubert. Mikromagnetisch singuläre punkte in bubbles. *Journal of Magnetism and Magnetic Materials*, 2(1-3):25–31, December 1975. [https://dx.doi.org/10.1016/0304-8853\(75\)90100-6](https://dx.doi.org/10.1016/0304-8853(75)90100-6).
- [267] J. C. Slonczewski. Properties of Bloch points in bubble domains. *AIP Conference Proceedings*, 24:613–614, 1975. <https://dx.doi.org/10.1063/1.30197>.
- [268] M. Margulies and J. C. Slonczewski. Abstract: Energy and mobility of Bloch point singularities. *Journal of Applied Physics*, 49(3):1912–1912, March 1978. <https://dx.doi.org/10.1063/1.324802>.
- [269] W. Jantz, J.C. Slonczewski, and B.E. Argyle. Effects of Bloch lines and Bloch points on resonances of magnetic bubbles. *Journal of Magnetism and Magnetic Materials*, 23(1):8–14, March 1981. [https://dx.doi.org/10.1016/0304-8853\(81\)90061-5](https://dx.doi.org/10.1016/0304-8853(81)90061-5).
- [270] A. Arrott, B. Heinrich, and A. Aharoni. Point singularities and magnetization reversal in ideally soft ferromagnetic cylinders. *IEEE Transactions on Magnetics*, 15(5):1228–1235, September 1979. <https://dx.doi.org/10.1109/tmag.1979.1060342>.
- [271] R. G. Elías and A. Verga. Magnetization structure of a Bloch point singularity. *The European Physical Journal B*, 82(2):159–166, July 2011. <https://dx.doi.org/10.1140/epjb/e2011-20146-6>.
- [272] O. V. Pylypovskiy, D. D. Sheka, and Y. Gaididei. Bloch point structure in a magnetic nanosphere. *Physical Review B*, 85(22), June 2012. <https://dx.doi.org/10.1103/physrevb.85.224401>.
- [273] C. Andreas, A. Kákay, and R. Hertel. Multiscale and multimodel simulation of Bloch-point dynamics. *Physical Review B*, 89(13), April 2014. <https://dx.doi.org/10.1103/physrevb.89.134403>.
- [274] R. Hertel and C. Andreas. Multiscale simulation of Bloch point dynamics in thick nanowires. In *Magnetic Nano- and Microwires*, pages 653–677. Elsevier, 2015. <https://dx.doi.org/10.1016/b978-0-08-100164-6.00022-9>.

[275] A. Fernández-Pacheco, R. Streubel, O. Fruchart, R. Hertel, P. Fischer, and R. P. Cowburn. Three-dimensional nanomagnetism. *Nature Communications*, 8(1), June 2017. <https://dx.doi.org/10.1038/ncomms15756>.

[276] Filipp N Rybakov, Aleksandr B Borisov, Stefan Blügel, and Nikolai S Kiselev. New spiral state and skyrmion lattice in 3d model of chiral magnets. *New Journal of Physics*, 18(4):045002, April 2016. <https://dx.doi.org/10.1088/1367-2630/18/4/045002>.

[277] M. Im, H. Han, M-S. Jung, Y-S Yu, S. Lee, S. Yoon, W. Chao, P. Fischer, J-I. Hong, and K-S. Lee. Dynamics of the Bloch point in an asymmetric permalloy disk. *Nature Communications*, 10(1), February 2019. <https://dx.doi.org/10.1038/s41467-019-08327-6>.

[278] J. Miltat. FERROMAGNETISM: Vortex cores-smaller than small. *Science*, 298(5593):555–555, October 2002. <https://dx.doi.org/10.1126/science.1077704>.

[279] M. Beg, R. A. Pepper, D. Cortés-Ortuño, B. Atie, M-A. Bisotti, G. Downing, T. Kluyver, O. Hovorka, and H. Fangohr. Stable and manipulable Bloch point. *Scientific reports*, 9(1):7959, 2019.

[280] C. S. Spencer, J. Gayles, N. A. Porter, S. Sugimoto, Z. Aslam, C. J. Kinane, T. R. Charlton, F. Freimuth, S. Chadov, S. Langridge, J. Sinova, C. Felser, S. Blügel, Y. Mokrousov, and C. H. Marrows. Helical magnetic structure and the anomalous and topological hall effects in epitaxial b20 $Fe_{1-y}Co_yGe$ films. *Physical Review B*, 97(21), June 2018. <https://dx.doi.org/10.1103/physrevb.97.214406>.

[281] Ryan Alexander Pepper, Marijan Beg, and Hans Fangohr. Dataset for "Field-driven Bloch Point Motion in Chiral Systems", 2020. <https://dx.doi.org/10.5281/ZENODO.4016634>.

[282] P.R. Kotiuga. The algebraic topology of Bloch points. *IEEE Transactions on Magnetics*, 25(5):3476–3478, 1989. <https://dx.doi.org/10.1109/20.42340>.

[283] G. M. Wysin and. Magnetic vortex dynamics in the non-circular potential of a thin elliptic ferromagnetic nanodisk with applied fields. *AIMS Materials Science*, 4(2):421–438, 2017. <https://dx.doi.org/10.3934/matersci.2017.2.421>.

[284] B. Krüger, A. Drews, M. Bolte, U. Merkt, D. Pfannkuche, and G. Meier. Harmonic oscillator model for current- and field-driven magnetic vortices. *Physical Review B*, 76(22), December 2007. <https://dx.doi.org/10.1103/physrevb.76.224426>.

[285] B. Van Waeyenberge, A. Puzic, H. Stoll, K. W. Chou, T. Tyliszczak, R. Hertel, M. Fähnle, H. Brückl, K. Rott, G. Reiss, I. Neudecker, D. Weiss, C. H. Back, and G. Schütz. Magnetic vortex core reversal by excitation with short bursts of an alternating field. *Nature*, 444(7118):461–464, November 2006. <https://dx.doi.org/10.1038/nature05240>.

[286] M. Urbánek, V. Uhlíř, C-H. Lambert, J. J. Kan, N. Eibagi, M. Vaňatka, L. Flajšman, R. Kalousek, M-Y. Im, P. Fischer, T. Šikola, and E. E. Fullerton. Dynamics and efficiency of magnetic vortex circulation reversal. *Physical Review B*, 91(9), March 2015. <https://dx.doi.org/10.1103/physrevb.91.094415>.