

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Author (Year of Submission) "Full thesis title", University of Southampton, name of the University Faculty or School or Department, PhD Thesis, pagination.

Data: Author (Year) Title. URI [dataset]

University of Southampton

Faculty of Engineering and Physical Sciences

Computational Engineering and Design

**Simulation of Thermally Fluctuating
Magnetic Nanoparticles for
Hyperthermia**

by

Oliver William Laslett

ORCID ID 0000-0001-9661-4637

Doctor of Philosophy

30th June 2019

University of Southampton

Abstract

Faculty of Engineering and Physical Sciences
Computational Engineering and Design

Doctor of Philosophy

Simulation of Thermally Fluctuating Magnetic Nanoparticles for Hyperthermia

by Oliver William Laslett

This thesis details the design, implementation, and application of numerical methods for simulating the stochastic dynamics of magnetic nanoparticles. In particular, the experiments in this thesis simulate the heat dissipated by magnetic nanoparticles when subjected to alternating magnetic fields. The nonlinear dynamics of these systems are relevant to hyperthermia, a clinical technique that uses localised particles to increase the temperature of tumour tissue and trigger cell death.

The simulations are implemented in an open-source software package, which focuses on reproducibility of results and rigorous testing of the underlying algorithms. In particular, the numerical errors of popular numerical schemes for simulating the stochastic Landau-Lifshitz-Gilbert equation are evaluated. These results suggest a trade-off between accuracy and computational complexity in systems with strong thermal fluctuations. The simulations are used to investigate the effects of novel applied waveform shapes and arbitrary shaped clusters of interacting particles.

The simulation results show that careful consideration of the applied field characteristics and clustering of particles can both enhance and diminish heat dissipation in magnetic nanoparticles. Large chains of particles aligned with the applied field are capable of very efficient heat dissipation but their strong dependence on the orientation of the applied field could lead to spatial variations in heating. Small dense clusters are shown to consistently perform worse than a single, isolated nanoparticle. Additionally, the use of square-wave applied magnetic fields are shown to improve the robustness of hyperthermia procedures and potentially increase heat dissipation when compared with traditional sinusoidal fields.

Research Thesis: Declaration of Authorship

Print name:	Oliver William Laslett
-------------	------------------------

Title of thesis:	Simulation of Thermally Fluctuating Magnetic Nanoparticles for Hyperthermia
------------------	---

I declare that this thesis and the work presented in it is 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. Either none of this work has been published before submission, or parts of this work have been published as: [please list references below]:

Chapter 5 "Design of Magpy" was submitted to the Journal of Open Research Software.

Preprint is available on arXiv <https://arxiv.org/abs/1801.06073>

Signature:		Date:	04/09/19
------------	--	-------	----------

Contents

1	Introduction	1
2	Background	5
2.1	Classical magnetic moment	5
2.1.1	Energy contributions	6
2.1.2	Effective magnetic field	7
2.1.3	Landau-Lifshitz-Gilbert dynamics	7
2.2	Magnetic domains	8
2.2.1	Hysteresis in bulk materials	8
2.2.2	Single-domain nanoparticles	9
2.3	Stochastic processes	12
2.3.1	Stochastic differential equations	12
2.3.2	Fokker-Planck equation	13
2.3.3	Jump processes	14
2.3.4	Master equation	15
2.3.5	Canonical ensemble	16
2.3.6	Simulation of stochastic processes	16
2.4	Finite-temperature magnetism	20
2.4.1	Stochastic Landau-Lifshitz-Gilbert equation	21
2.4.2	Thermal equilibrium	21
2.4.3	Thermal activation	22
2.4.4	Linear response theory	25
2.4.5	Heat and work	25
3	Literature Review: Magnetic particle hyperthermia	27
3.1	Therapeutic applications	27
3.2	Current challenges	28
3.3	Heat dissipation in magnetic nanoparticles	29
3.3.1	Optimising particle and field properties	29
3.3.2	Effect of interactions	33
3.4	Scope of present study	35
4	Simulating the Landau-Lifshitz-Gilbert equation	37
4.1	Equation properties	38
4.1.1	Reduced equations	39
4.2	Review of numerical methods	40
4.3	Technical definition and implementation	43

4.3.1	Euler-Maruyama scheme	43
4.3.2	Stochastic Heun Scheme	44
4.3.3	Fully implicit midpoint	44
4.4	Results and validation	48
4.4.1	Linear test equation	48
4.4.2	Deterministic Landau-Lifshitz-Gilbert equation	49
4.4.3	Stochastic Landau-Lifshitz-Gilbert equation	53
4.4.4	Performance	57
4.4.5	Equilibrium	58
4.5	Conclusions	61
5	Design of Magpy	63
5.1	Introduction	63
5.2	A model for magnetic nanoparticle dynamics	63
5.3	Thermal activation for single particles	65
5.4	Alternative software	68
5.5	Implementation and architecture	69
5.6	Quality control	70
5.7	Example: the first passage time distribution	73
5.7.1	Results	74
6	Nanoparticle cluster geometries and heating effects	79
6.1	System definition	79
6.1.1	Cluster geometries	80
6.1.2	Random anisotropy axes	81
6.1.3	Random cluster orientation	81
6.2	Heat dissipation in the major-loop regime	81
6.2.1	Randomly oriented clusters	82
6.2.2	Anisotropy-aligned clusters	85
6.2.3	Discussion	86
6.3	Heat dissipation in the minor-loop regime	87
6.3.1	Effect of increasing field amplitude	87
6.3.2	Effect of increasing cluster size	87
6.3.3	Discussion	89
6.4	Comparison with experimental data	91
6.4.1	Tuning the damping ratio	92
6.4.2	Heat dissipation in dense clusters	92
6.4.3	Heat dissipation in mixed solutions	94
6.4.4	Anisotropy-aligned clusters	96
6.5	Local heating distribution	97
6.5.1	Variation of SPL between clusters	98
6.5.2	SPL variation with shape orientation	98
6.6	Conclusions	101
7	Non-sinusoidal alternating magnetic fields in hyperthermia	103
7.1	Introduction	103
7.2	Heat dissipation models for arbitrary fields	104

7.2.1	Thermal activation model	105
7.2.2	Linear response theory	105
7.3	Comparison of low-amplitude square and sinusoidal AMFs	107
7.4	Comparison of high-amplitude square and sinusoidal AMFs	109
7.4.1	Analytic master-equation dynamics	110
7.4.2	Numerical simulations with high-amplitude fields	113
7.5	Polydispersion and filtered AMFs	115
7.6	Conclusions	118
8	Conclusions and outlook	119
A	Stochastic integrals	123
B	Normalisation	125
C	Arkus cluster definitions	129
D	Low temperature simulations from Chapter 6	137
	References	141

List of Figures

2.1	Hysteresis loops of magnetic materials depict the variation of their equilibrium magnetisation with the amplitude of an applied magnetic field.	9
2.2	The time-evolution of the Wiener process. The black paths show possible realisations (solutions) to the stochastic differential equation describing the Wiener process. The Fokker-Planck equation describes the probability distribution over the possible states at a time t and is shown as a heat map.	14
3.1	The validity of the linear response theory for calculating the SPL of an ensemble of single-domain non-interacting magnetic nanoparticles. The LRT area represents the validity of the linear model for magnetisation response. The NB area represents the validity of the Néel-Brown model for computing the relaxation time of the particle.	31
4.1	Structure of the diffusion Jacobian for a system of 3 magnetic moments. $0_{3,3,3}$ represents the 3-dimensional cube of zeroes of size $3 \times 3 \times 3$. $B^n(\hat{\mathbf{m}}^n, t)$ denotes the diffusion of the n^{th} magnetic moment in the system.	46
4.2	Comparison of Heun and fully implicit methods for approximating the linear test equation with different coefficients μ, σ . Setting σ higher introduces stiffness in the diffusion term.	49
4.3	Example solution of the deterministic Landau Lifshitz Gilbert equation for a single spin.	50
4.4	Path-wise error of numerical approximations to the deterministic LLG equation using the Heun and fully implicit schemes. The error is shown for three different time steps $\Delta t = 1, 10, 100$	50
4.5	Empirical strong convergence of the Euler, Heun, and implicit ($\epsilon = 10^{-9}$) methods on the deterministic LLG.	51
4.6	Empirical strong convergence of the implicit scheme on the deterministic LLG with varying QNS tolerance ϵ	52
4.7	Global error in the magnetic moment magnitude as a function of the time step for Euler, Heun, and implicit (with varying tolerance ϵ) methods. The implicit scheme shows an unexpected results at $\Delta t = 1$	52
4.8	A single continuous Wiener process sampled at three different intervals. The Wiener processes used to compute the convergence rates must be identical when using different time steps.	54
4.9	Comparison of the Heun and implicit schemes for the simplified stochastic LLG equation. The two schemes were used to estimate the trajectory of the equation with $H = 0.01, \alpha = 0.1, \sigma = 0.05$ (true solution shown in left panel). The empirical strong convergence of the schemes was estimated (middle panel). The global magnetic moment error was computed for each time step Δt (right panel).	55

4.10	Comparison of the Heun and implicit schemes for the simplified stochastic LLG equation. The experiment from Figure 4.9 was repeated with the parameters $H = 0.01, \alpha = 0.1, \sigma = 0.25$	55
4.11	Comparison of the implicit scheme and the Heun scheme with an added renormalisation step using the same parameters from Figure 4.9. The missing data points (right panel) correspond to $1 - \ \mathbf{m}\ _2 = 1$ to machine precision.	56
4.12	Empirical strong convergence of the Heun and implicit methods of the full sLLG equation, estimated using the convergence of the relative global error between subsequent time steps (equation (4.32)).	57
4.13	Equilibrium distribution of 10,000 non-interacting magnetic moments simulated with the fully implicit method and $\sigma \approx 35$. The normalised histogram resulting from the simulation closely matches the empirical distribution.	59
4.14	Equilibrium distribution for a system of two interacting magnetic moments computed using stochastic simulation of an ensemble of 50,000 systems and the Metropolis MCMC method.	60
4.15	The marginal equilibrium distribution for a single spin in a system of two magnetic moments. The results compare the results of stochastic simulation using an ensemble of 50,000 systems with and without the interaction term, against the results of MCMC sampling.	60
5.1	A two-dimensional sketch of a nanoparticle. The atoms of a magnetic material are packed into a regular crystal lattice and each is modelled by a three-dimensional magnetic moment that varies with time. Due to strong interactions between these magnetic moments in small particles, Magpy assumes they rotate coherently and are represented by a single macrospin vector.	64
5.2	The three effective field contributions acting upon macrospin i . (1) The macrospin experiences a force towards alignment with the particle anisotropy axis (dashed line) \mathbf{k}_i in either direction. (2) The macrospin is also forced towards aligning with the externally applied field direction (solid arrows) \mathbf{h}_{app} . (3) Finally, each macrospin is repelled and attracted by nearby macrospins. The force of the dipolar-dipolar interaction diminishes with distance (dash-dotted line) between two particles \mathbf{R}_{ij}	65
5.3	Simulating a chain of three particles (created in Magpy using Listing 1). (left) A chain structure of three identical particles are initialised with an external field applied along their anisotropy axis and their magnetisation initially against the applied field. (right) The coloured lines show the total magnetisation in the direction of the applied field for 5 simulations from the same initial condition (see left). The black line is the result of averaging 500 simulations from the same initial condition (i.e. the expected or mean trajectory).	66

5.4	In the thermal activation model, a single nanoparticle may occupy one of only two states: <i>up</i> or <i>down</i> , which each correspond to local energy minima around the anisotropy axis (left). The dynamics of the system are described by a master equation with transition rates $w_{u \rightarrow d/d \rightarrow u}$ between the two states. The solution of the master equation depends on the initial condition, particle properties and applied field. In this example (right) the particle is initialised up with probability 1 and allowed to relax into equilibrium (see Listing 2 for Magpy script). The equilibrium magnetisation depends upon the strength of the constant applied magnetic field.	66
5.5	The flow of data through an ensemble simulation in Magpy. (1) The user instantiates an <code>EnsembleModel</code> object and calls the <code>EnsembleModel.simulate</code> function, specifying the number of CPU cores to utilise. (2) The main Python process spawns a new individual process for each model in the ensemble. (3) The individual processes each call the Magpy C++ library using their respective model parameters. (4) The results from the C++ simulation are returned to the individual python process. (5) When all the processes have finished, the results are collected on the main process to be analysed and plotted.	71
5.6	Mean first passage time for a Stoner-Wohlfarth particle. Comparison of Langevin stochastic simulation and the analytic Néel-Brown formula shows good agreement in the intermediate regime. $\sigma = 8.8$, $\vec{h} = 0$, $\hat{u} = [0 \ 0 \ 1]$, $\hat{m}_0 = [0 \ 0 \ 1]$	75
5.7	The empirical distribution of first passage times from Langevin simulations with decreasing stability ratio σ . Parameters as in Figure 5.6 with $\alpha = 0.1$. The empirical histograms are fitted with parameterised distribution functions by maximising the likelihood.	76
5.8	The KS distance measures the difference between the first passage time distribution of the stochastic simulations and the analytic Néel-Brown result. A lower KS distance is desirable. Parameters as in Figure 5.6.	77
6.1	Dynamic major hysteresis loops for chain and Arkus cluster geometries with 1-8 particles. Results were obtained by stochastic simulation of coupled sLLG equations using Magpy. The vertical and horizontal axes are shared between all figures. Results are shown for randomly oriented clusters with random anisotropy axes and clusters aligned with the applied field with aligned anisotropy axes.	82
6.2	Dynamic hysteresis loop quantities and SPL of randomly-oriented ensembles of chain of increasing length N . The SPL is computed by numerical integration of the hysteresis loops in figure 6.1(a).	83
6.3	Dynamic hysteresis loop quantities and SPL of randomly-oriented ensembles of Arkus clusters of increasing size N . The SPL is computed by numerical integration of the hysteresis loops in Figure 6.1(b).	83
6.4	Arkus clusters of anisotropy-aligned particles. All anisotropy directions (red arrows) are aligned in the same direction as the alternating magnetic field. The figure shows rendered clusters of size 1-6. In the aligned simulation, the relative positions of the particles are also fixed in the arrangements shown. The longitudinal axis of cluster (2) and the plane spanned by the particles in (3) are parallel with the applied field.	85

6.5	SPL obtained from numerical integration of hysteresis loops. The hysteresis loops were simulated for Chain and Arkus geometries of various sizes both randomly-oriented and aligned with the externally applied field. The results are shown as a function of the alternating magnetic field reduced amplitude (normalised by the anisotropy field H_k for a single particle).	88
6.6	Simulated dynamic hysteresis loops for an ensemble of aligned non-interacting nanoparticles. The amplitude of the alternating magnetic field was increased from $0.1H_k$ to $3.0H_k$	88
6.7	Heat dissipation (SPL) resulting from subjecting Arkus clusters and chains of increasing size to an alternating magnetic field of low amplitude $H = 0.5H_k$ where H_k is the anisotropy field of the constituent particles. In all figures, the particles and clusters are randomly oriented.	90
6.8	SPL of a single magnetic nanoparticle simulated with Magpy using varying values of the damping ratio α . High values of α lead to high longitudinal damping of the magnetic moment. The various values are compared to the SPL simulated using kinetic Monte-Carlo simulations (kMC) in [152]	92
6.9	SPL values obtained from kinetic Monte Carlo and Magpy simulations of dense clusters of particles. The anisotropy value of the particles was changed in each simulation, which is represented by the dimensionless stability ratio $\sigma = KV/(k_B T)$	93
6.10	Dynamic hysteresis loops simulated using magpy for $\alpha = 0.1$ and $\sigma = 5.1$ ($K = 50\text{kJ/m}$) compared with simulated loops obtained by kinetic Monte Carlo simulation.	94
6.11	SPL prefactors for each cluster size for three different solutions synthesised in [152] using different amounts of polymer in order to control the degree of clustering in each solution.	95
6.12	Cumulative SPL dissipated by solutions of mixed clusters as described in Figure 6.11. The clusters were simulated with $\alpha = 0.1$ and $\sigma = 5.1$. The simulated SPL from the kMC method were a close match to experimental values but the magpy simulations show substantially lower SPL.	96
6.13	Cumulative SPL dissipated by solutions of mixed clusters as described in Figure 6.11. The kMC simulations used randomly oriented clusters and anisotropy directions. The magpy simulations held the anisotropy axes aligned and randomly oriented the entire clusters. $\alpha = 0.1$ and $\sigma = 5.1$	97
6.14	Total SPL and coefficient of variation for mixed solutions of Arkus and chain clusters of different sizes as a function of the applied field reduced amplitude H/H_k	99
6.15	Distribution of SPL as a function of the cluster orientation in three-dimensional space. Both plots convey the same distribution of SPL in two and three dimensions.	100

- 7.1 Normalised energy dissipated in a single magnetic nanoparticle subjected to a 0.05mT square and sinusoidal alternating magnetic field. The energy has been computed using both the linear response theory (LRT) and numerical simulations (DOM) with $M_s = 10^6 \text{Am}^{-1}$, $K = 10^4 \text{Jm}^{-3}$, $\alpha = 0.1$, $T = 300\text{K}$, $\mu_0 H = 0.05\text{mT}$, $f = 100\text{kHz}$, and r in the range of 8-12nm such that $2\sigma h \ll 1$ for all r . The curves for the two methods match exactly, validating the simulation approach. The results for the sine wave are multiplied by two to be of equivalent power to the square wave. The insets show the two hysteresis loops from the sine (blue) and square (orange) applied magnetic fields resulting from numerical simulations at the indicated value of $\omega\tau$ 108
- 7.2 The magnetisation response of a magnetic nanoparticle to a square alternating magnetic field of period T (left). In dynamic equilibrium the magnetisation is assumed to oscillate between $\pm M_p$. The corresponding hysteresis loop in the M - H plane (right) is always an oblong and is defined by the amplitude of the applied field and the peak value of magnetisation M_p 110
- 7.3 The equilibrium magnetisation as predicted by the linear (dashed) and nonlinear (solid) models, which is proportional to the normalised energy dissipated $\Delta U / (2\sigma h \tanh(\pi / (2\omega\tau)))$. The linear response model overpredicts ΔU for large field amplitudes $2\sigma h > 1$. Markers (circles) mark $2\sigma h = \frac{1}{3}$ 112
- 7.4 Energy dissipated per cycle in a single magnetic nanoparticle subjected to a 5mT alternating magnetic field of sinusoidal and square shape as a function of the particle radius. The results were obtained using numerical simulations with the same material properties as in Figure 7.1. The insets show the dynamic hysteresis loop obtained from the simulations at the indicated radius. 113
- 7.5 The gain in energy dissipated resulting from applying a square alternating magnetic field. Gain is defined as $\log_{10}(\Delta_{\text{square}}) - \log_{10}(\Delta_{\text{sine}})$. A particle of magnetite of varying radius was subjected to each of the alternating magnetic field shapes with increasing field amplitude. Material properties as in figure 7.1. The contour line (0.0) represents the h, r combinations for which the energy dissipated is equal under both field shapes. The sinusoidal optimum radius (dash-dotted line) represents $r_{\text{opt}} = \text{argmax}_r \Delta U_{\text{sine}}(r)$ for each value of the applied field h 114
- 7.6 Energy dissipated in polydisperse ensembles of magnetic nanoparticles subjected to different applied magnetic field waveforms. The distribution of radii in each ensemble is lognormal with a mean parameter \bar{R} and scale parameter S . Simulations used the following parameters: $M_s = 10^6 \text{Am}^{-1}$, $K = 2 \times 10^4 \text{Jm}^{-3}$, $\alpha = 0.1$, $T = 300\text{K}$, $\mu_0 H = 5\text{mT}$, $f = 100\text{kHz}$. The figures (a-d) correspond to simulations using a sinusoidal, square (10 harmonics), square (100 harmonics), and square waveform respectively. The inset plots show a sketch of one period of each applied waveform shape. 116
- 7.7 Increase in maximum energy dissipation for square waves compared to sine waves at various field values for polydisperse ensembles. The LRT region marks the values of h for which the linear response theory is applicable. 117

- D.1 Dynamic major hysteresis loops for chain and Arkus cluster geometries with 1-8 particles at temperature $T = 30\text{K}$. Results were obtained by stochastic simulation of coupled sLLG equations using Magpy. The vertical and horizontal axes are shared between all figures. Results are shown for randomly oriented clusters with random anisotropy axes and clusters aligned with the applied field with aligned anisotropy axes. 138
- D.2 SPL obtained from numerical integration of hysteresis loops at temperature $T = 30\text{K}$. The hysteresis loops were simulated for Chain and Arkus geometries of various sizes both randomly-oriented and aligned with the externally applied field. The results are shown as a function of the alternating magnetic field reduced amplitude (normalised by the anisotropy field H_k for a single particle). 139

List of Tables

4.1	Global mean squared error convergence for popular stochastic numerical schemes.	43
4.2	Performance results for the fully implicit and Heun scheme solving the sLLG equation for a single magnetic moment. The results are in units of CPU time per unit of simulation time (larger is worse). The speed-up is the relative increase in performance from using the Heun scheme.	57
4.3	Speed-up of the Heun scheme relative to the IMP method for interacting particle chains of increasing length.	58

“Those diseases which medicines do not cure, the knife cures; those which the knife cannot cure, fire cures; and those which fire cannot cure are reckoned wholly incurable.”

- Hippocrates of Kos

Chapter 1

Introduction

The first recorded procedure using heat to treat disease is found in the Edwin Smith Papyrus, an ancient document dating back 5000 years. The papyrus documents Egyptian physicians using cauterisation as a means to treat breast tumours. Thermal based therapies, in a variety of forms, are still present in clinical oncology today and have been transformed by the emergence of nanomedicine at the beginning of the millennium. Nanomedicine, and in particular development of nanoparticles, has led to breakthrough technologies in medical imaging, drug delivery, and novel therapeutic techniques. Moreover, recent research shows that nanoparticles have multi-functional properties and may be able to do all this at once.

First demonstrated in 1957, magnetic hyperthermia is a method to induce localised heating in tissue by the use of magnetic nanoparticles subjected to an alternating field. Magnetic hyperthermia has experienced rapid development in the last 20 years, enabled by nanotechnology, particle synthesis methods, and computational simulations. Today, magnetic hyperthermia has clinical approval in Europe and researchers are finding new avenues for improving existing therapies such as chemotherapy and radiotherapy. However, magnetic hyperthermia faces many medical, biological, and engineering challenges. A fundamental design requirement for hyperthermia is to maximise the heat dissipated by the nanoparticles while minimising patient exposure to high amplitude magnetic fields. The optimisation of the material, environmental, and applied field properties to maximise treatment efficacy have benefited from the theory and simulation of magnetic dynamics.

The theory of fluctuating magnetic nanoparticles was pioneered by the work of Brown, who proposed a stochastic differential equation to describe the magnetic dynamics of nanoparticle at finite temperature. Additionally, the stochastic dynamics were shown to reduce to a more simple stochastic jump process in certain cases. Since their initial publication, the stochastic Landau-Lifshitz-Gilbert equation and Néel-Brown models have

contributed substantially to the optimisation of magnetic hyperthermia and other technologies such as magnetic disk storage. Computational simulations are a desirable alternative to clinical experiments because they allow fine grained control and measurement at lower cost and without risk. However, increasingly complex computational experiments must also be held to the same standards of reproducibility as traditional experiments in order to maximise their scientific impact. Chapter 2 collects the detailed mathematical and scientific pre-requisites for the stochastic simulation of magnetic nanoparticles. The nomenclature and equations introduced in this chapter are used throughout the thesis.

A large number of factors influence the heat dissipation of magnetic nanoparticles and the optimum combination of parameters remains unsolved. Chapter 3, reviews the state of magnetic hyperthermia and the search for the optimum parameters. A subset of parameters are well understood; for example the anisotropy strength and size of isolated, uniaxial nanoparticles. However, the effects of interactions in dense clusters of particles are not well established and have only been studied relatively recently. Moreover, non-sinusoidal applied magnetic fields have shown the potential to improve heat dissipation but have received very little attention.

The focus of this thesis is to implement reliable simulations of magnetic nanoparticle dynamics. The simulations are used to experiment with the effects of interactions between particles formed into dense clusters as well as the effects of non-sinusoidal applied fields. Chapter 4, details the selection and implementation of a numerical method to solve the Landau-Lifshitz-Gilbert equation, which is currently an open research question. The mathematical properties of the equation complicate the implementation of reliable numerical methods. The explicit Heun scheme and a fully implicit method are compared using their path-wise error. The implicit method yields substantially lower error than the Heun scheme but is computationally expensive. The numerical tests also show that conservation of the magnetic moment does not necessarily reduce the total error of the simulations. The numerical methods are implemented in an open-source software package, detailed in Chapter 5. The software provides an implementation of complex numerical procedures not currently available under a open-source license. The accompanying documentation and testing suite ensures future researchers can reuse and contribute improvements to the simulation code.

Chapter 6 investigates the effects of the geometry of dense clusters of particles subjected to inter-cluster interactions. Long chains of particles oriented with the applied magnetic field are shown to dissipate heat very efficiently. However, their alignment must be carefully controlled and the applied field must be strong in order to avoid potentially dangerous spatial variations in temperature. Isotropic, sphere-like, clusters of particles are shown to dissipate very little heat due to their paramagnetic-like response. The simulations are compared to an existing experimental study and provides an alternative explanation for the experimentally observed behaviour of clusters of nanocubes.

Chapter 7, uses analytic and simple models of heat dissipation to investigate the effect of non-sinusoidal fields. The linear response theory is revisited for arbitrary fields, which predicts that square-wave fields lead to substantial improvements in heat dissipation. Nonlinear models of thermal activation in single domain, isolated particles show that at larger fields, the benefit of using a square wave diminishes. However, the results also suggest that non-sinusoidal waveforms increase total heat dissipation in size-dispersed ensembles of particles, which could increase the robustness of current magnetic hyperthermia procedures.

Chapter 2

Background

2.1 Classical magnetic moment

Magnetism is a fundamentally quantum phenomenon that arises from an intrinsic property of electrons known as *spin* and the orbital motion of electrons around atoms within a material. In a semi-classical description, the motion of a negatively charged particle forms a current loop, which possesses a magnetic dipole moment due to orbital momentum. However, the total magnetic moment includes an additional contribution due to electron spin, which may only be described with a quantum mechanical description and obtained via experiments or detailed calculations. The resulting total magnetic moment is a vector quantity $\boldsymbol{\mu} \in \mathbb{R}^3$ with fixed magnitude $|\boldsymbol{\mu}| = \mu_s$ and a direction that points towards the positive (north) pole of the dipolar field. In addition to *producing* a magnetic field, the magnetic moment may also be *influenced* by internal or external magnetic fields, just as observed with bar magnets in the classroom. When a magnetic dipole is subjected to an external magnetic field $\mathbf{B} \in \mathbb{R}^3$, a torque is induced in the magnetic moment resulting in a precession around \mathbf{B} :

$$\frac{d\boldsymbol{\mu}}{dt} = -\gamma(\boldsymbol{\mu} \times \mathbf{B}) \quad (2.1)$$

where $\omega = -\gamma|\mathbf{B}|$ is the angular frequency and γ is termed the magnetogyric ratio [186]. Electrons, and consequently their associated magnetic moments, move freely through the crystal lattice of magnetic materials. However, in this description, the electrons are assumed to be localised to the atomic sites in the lattice and therefore each atom has an associated magnetic moment (i.e. the atomic magnetic moment) [56].

Ferromagnetic materials (such as iron, cobalt, and nickel) experience strong coupling between magnetic moments, which leads to entire domains that behave as single large magnetic moments. These magnetic domains are described as *uniformly magnetised* and their total magnetic moment a *macrospin*. The magnetisation $\mathbf{M} = \boldsymbol{\mu}/V$ defines the

magnetic moment per unit volume of an electron/atom/domain and has a saturation magnetisation $M_s = |\mathbf{M}|$. The unit vector of the magnetisation is denoted $\mathbf{m} = \mathbf{M}/M_s$.

2.1.1 Energy contributions

In this thesis, we will adopt a semi-classical Heisenberg-like description of a system of N localised magnetic moments. The total free energy of the system is the sum of magnetic moment energy terms: $E = \sum_{i=0}^N E(\mathbf{m}_i)$. The respective energy of each magnetic moment depends on both its state \mathbf{m}_i and the state of particles with which it interacts $\mathbf{m}_j, j \neq i$. We will consider four different energy contributions to the system.

2.1.1.1 Zeeman energy

The Zeeman energy describes the interaction between a spin $\boldsymbol{\mu}_i$ and a (locally) spatially uniform external magnetic field:

$$E_z(\mathbf{m}_i) = -\mu_s \mathbf{m}_i \cdot \mathbf{B} \quad (2.2)$$

The Zeeman energy is minimised when the magnetic moment \mathbf{m}_i aligns with the external field direction \mathbf{B} and therefore the magnetic moment tends to align in the direction \mathbf{B} .

2.1.1.2 Exchange energy

The Heisenberg exchange energy describes a quantum mechanical interaction between magnetic moments, which favours parallel alignment. Heisenberg's classical approximation describes the exchange energy between neighbouring magnetic moments as:

$$E_{\text{ex}}(\mathbf{m}_i) = - \sum_{j \in \mathcal{S}_i} J_{ij} \mathbf{m}_i \cdot \mathbf{m}_j \quad (2.3)$$

where J_{ij} is the exchange constant and \mathcal{S}_i is the set of neighbouring magnetic moments to the moment \mathbf{m}_i . The range (number of nearest neighbours) and strength of the exchange term varies for each material. The definition of \mathcal{S}_i also depends on the dimensionality of the system. For $J > 0$ neighbouring atomistic magnetic moments favour parallel alignment and for $J < 0$ anti-parallel alignment.

2.1.1.3 Dipole-dipole interaction energy

Each magnetic moment produces a dipolar magnetic field, which influences other magnetic moments. The dipole-dipole (or demagnetisation) interaction describes the effect

that a magnetic moment \mathbf{m}_i experiences from all other magnetic moments in the system:

$$E_{\text{dd}}(\mathbf{m}_i) = - \sum_{j \neq i} \mu_0 \frac{\mu_0 \mu_s^2}{4\pi |R_{ij}|^3} (3(\mathbf{m}_i \cdot \mathbf{r}_{ij})(\mathbf{m}_j \cdot \mathbf{r}_{ij}) - \mathbf{m}_i \cdot \mathbf{m}_j) \quad (2.4)$$

$\mathbf{R}_{ij} = |\mathbf{R}_{ij}| \mathbf{r}_{ij}$ is the distance between magnetic moments i and j . The strength of the dipole-dipole interaction reduces with the cube of the distance between the two particles. The result of this long-range interaction term is that magnetic moments tend to align end-to-end in the same direction.

2.1.1.4 Anisotropy energy

The anisotropy energy term describes the preferential direction of the magnetic moment due to the magneto-crystalline orientation. The crystalline anisotropy term depends upon the magnetic material. Commonly found in ferromagnetic materials is the uniaxial anisotropy that acts in the unit direction \mathbf{e} with an anisotropy constant k_i :

$$E_a(\mathbf{m}_i) = -k_i(\mathbf{m}_i \cdot \mathbf{e}_i)^2 \quad (2.5)$$

There are many additional possible anisotropy contributions such as shape, magneto-elastic, and cubic magneto-crystalline, which will not be considered.

2.1.2 Effective magnetic field

The effective magnetic field experienced by a magnetic moment includes contributions from externally applied magnetic fields, long and short range interactions between magnetic moments, and interactions between the magnetic moment and the atomic lattice of the material. These individual contributions together form an effective field, which is obtained from the gradient of the free energy potential of the magnetic moment:

$$\mathbf{B}_{\text{eff}}(\mathbf{m}_i) = -\frac{1}{\mu_s} \frac{\partial E(\mathbf{m}_i)}{\partial \mathbf{m}_i} \quad (2.6)$$

The effective field varies for each magnetic moment in the system and although continuously varying throughout the material will be assumed constant for each local atomic site.

2.1.3 Landau-Lifshitz-Gilbert dynamics

The dynamics of a magnetic moment subjected to a constant magnetic field were described previously by a precession about \mathbf{B}_{eff} (equation (2.6)). However, the Zeeman energy term (equation (2.2)) shows that throughout the precessional motion the energy

is constant and not minimum. This is counter-intuitive to the general principle of that systems should eventually reduce their total energy. Indeed, it is an experimental fact that the magnetic moment does eventually align with the applied field due to dissipative processes. Accounting for this evidence requires the addition of an additional term to equation (2.1) with a torque towards the precessional axis.

A phenomenological damping term, describing an unspecified dissipative process, was first proposed by Landau and Lifshitz [121] and later revised by Gilbert [73]. The resulting dynamical equation for the damped magnetic moment is the Landau-Lifshitz-Gilbert equation (LLG), which is written in terms of the unit vector $\mathbf{m}_i = \boldsymbol{\mu}_i/\mu_s$:

$$\frac{d\mathbf{m}_i}{dt} = \frac{-\gamma}{1+\alpha^2} (\mathbf{m}_i \times \mathbf{B}_{\text{eff},i}) - \frac{\alpha\gamma}{1+\alpha^2} \mathbf{m}_i \times (\mathbf{m}_i \times \mathbf{B}_{\text{eff},i}) \quad (2.7)$$

where α is the dimensionless Gilbert damping constant and $\mathbf{B}_{\text{eff},i}$ is the effective field experienced by particle i . The first term in equation 2.7 describes the precessional motion and the second term describes the damping. Due to the conservation of momentum, the damping process transfers energy to the spin-lattice. Note also that the magnitude of the unit magnetic moment $|\mathbf{m}_i|$ is still preserved.

2.2 Magnetic domains

2.2.1 Hysteresis in bulk materials

The magnetic moments in large samples of material interact with each other and can form local regions of aligned moments termed *magnetic domains*. The emergence of these uniformly magnetised domains is energetically favourable because of the strong exchange interactions between neighbouring magnetic moments (equation (2.3)). Domains of locally aligned moments are separated by *domain walls*. The number and orientation of domains within a sample vary with the material, temperature, and externally applied field. Under the application of a strong enough external field, all magnetic moments will be aligned such that the magnetisation (total moment per unit volume $\mathbf{M} = (1/V) \sum_i \boldsymbol{\mu}_i$) is at a maximum with $|\mathbf{M}| = M_s$, termed the saturation magnetisation. When saturated, the material consists of a single domain. At the other extreme, a material with many randomly oriented magnetic domains will have a magnetisation close to zero.

Magnetic materials are categorised by their equilibrium magnetisation as a function of an externally applied field $M(H)$, which depends on the chemical compounds and crystalline structure of the material as well as external effects such as temperature. The magnetisation is described as a memory-dependent process that shows hysteresis. Hysteresis loops are obtained experimentally by incrementing the field amplitude in

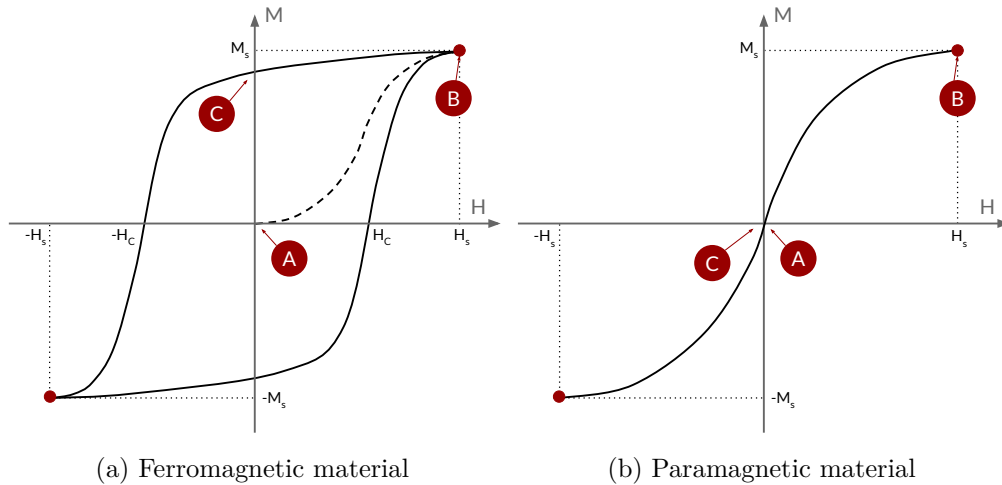


Figure 2.1: Hysteresis loops of magnetic materials depict the variation of their equilibrium magnetisation with the amplitude of an applied magnetic field.

small intervals, allowing the system sufficient time to relax (such that transient effects are diminished) and the magnetisation along the direction \mathbf{H} is recorded.

In the absence of an external field a *paramagnetic* material has randomly oriented domains and thus exhibits a zero magnetisation (point A in Figure 2.1b). The magnetic moments in the material will align if subjected to an external field. As the field strength increases, domains form and merge to create larger domains until eventually the moments are aligned at the saturation field H_s (point B in Figure 2.1b). When the field is removed from a paramagnetic material, the magnetic moments relax into random orientations and the magnetisation returns to zero (point C in Figure 2.1b).

In *ferromagnetic* materials, an externally applied field will also eventually saturate the magnetisation of the material at H_s (point B in Figure 2.1a). However, on removal of the field, the magnetic moments within a ferromagnetic material retain some ordering and hence the material exhibits a remanent magnetisation (point C in Figure 2.1a). The applied field must increase in opposition to the magnetisation until the critical field H_c is reached in order to return the magnetisation to zero. The remanent magnetisation is not permanent but a (potentially very long-lived) meta-stable state due to the effect of thermal fluctuations, described in Section 2.4.

2.2.2 Single-domain nanoparticles

Sufficiently small, nano-sized particles of magnetic material consist of a single domain [32]. Unlike bulk materials in which domains are created and destroyed by magnetic fields, nanoparticles consist of a single uniformly magnetised domain throughout the hysteresis cycle. In other words, the non-equilibrium behaviour of the individual magnetic moments is a *coherent rotation*. The maximum dimensions of a single-domain particle

depend on the geometry and material properties of the nanoparticle [97]. In addition to size limits, the single domain may also be broken by the application of a very strong external field that may induce nucleation, fanning, or curling [14].

A particle that is assumed to consist of perfectly coupled magnetic moments may be represented by a single macrospin (Section 2.1). The theory of coherent rotation in magnetic nanoparticles with uniaxial anisotropy was developed extensively by Stoner and Wohlfarth [187]. Hence the model is often termed the Stoner-Wohlfarth model. The free energy of a non-interacting Stoner-Wohlfarth particle consists of the total Zeeman and anisotropy energy terms of the individual magnetic moments. The state of the magnetic moment is described by two angles θ, ϕ representing the angle between \mathbf{m} and \mathbf{B} and the angle between \mathbf{e} and \mathbf{B} respectively. The free energy of the particle $E(\theta, \phi)$ in state (θ, ϕ) is then given as:

$$\frac{E(\theta, \phi)}{k_B T} = \sigma \sin^2(\theta - \phi) - 2\sigma h \cos \theta \quad (2.8)$$

where T is the temperature (K), k_B is the Boltzmann constant, $\sigma = KV/(k_B T)$ is the stability ratio, and $h = |\mathbf{B}|/B_k$ is the reduced applied magnetic field, which is normalised by the anisotropy field $B_k = 2K/(M_s)$. The first term in equation (2.8) describes the preferential orientation of the magnetic particle moment towards the anisotropy axis and the second term describes preferential alignment with the field direction.

In the special case that the externally applied field \mathbf{B} is applied along \mathbf{e} the equation reduces to a function of θ only:

$$\frac{E(\theta)}{k_B T} = \sigma \sin^2 \theta - 2\sigma h \cos \theta \quad (2.9)$$

When $h(t) < 1$ the particle exhibits two energy minima $\theta_1 = 0, \theta_2 = \pi$ separated by a maximum at $\theta_0 = \cos^{-1}(-h)$. The normalised energy barrier (divided by $k_B T$) separating the two minima is $\sigma(1 \pm h^2)$.

2.2.2.1 Interparticle interactions

The magnetic behaviour of an ensemble of single domain magnetic nanoparticles is influenced by inter-particle interactions. Interactions between individual magnetic nanoparticles lead to collective behaviour, formation into aggregate clusters, and a range of other complex phenomena [14, 104, 135, 145, 162]. The interactions between nanoparticles can be approximated using a number of methods [49], which describe different interactions [14]. However, all these interactions result from the individual magnetic moments within the particles interacting.

If the exchange energy between the intra-particle magnetic moments dominates the dipole-dipole interaction energy between the inter-particle magnetic moments then we

may still assume coherent rotation of the particle macrospin. Consequently, the interaction between two particles is well described by a dipolar interaction between the respective macrospins at the centre of the particles. For perfectly spherical particles, this mean-field assumption is equivalent to the sum of intra-particle magnetic moment dipolar interactions. For other shapes it is only an approximation [48]. This assumption is very common in modelling magnetic nanoparticle dynamics and has successfully described experimental results [19, 98, 161]. However, it is also possible to mathematically derive the particle interactions from the individual magnetic moments, such as for 2D arrays of magnetic dots [163].

Modelling the dipolar interaction between macrospins significantly reduces the number of interactions in the system by neglecting the internal degrees of freedom. However, if the surface-to-surface distance between particles becomes small enough, the magnetic moments in the shell of each particle may interact through the exchange interaction term violating the homogeneity of the macrospin assumption [14, 145].

2.2.2.2 Macrospin dynamics

In a macrospin model in which the magnetic moments are perfectly coupled, the dynamic behaviour of the macrospin is generally assumed to follow the phenomenological Landau-Lifshitz-Gilbert equation. Recall that a fundamental assumption of the LLG equation is that the magnetisation of the macrospin is unchanged. However, if the particle's constituent magnetic moments deviate marginally from the perfectly aligned case, then the magnetisation will be reduced. However, this assumption has proved useful since the earliest work of Brown [22]. The LLG equation and effective fields terms are written in terms of magnetisation of the macrospin \mathbf{M}_i by substituting μ_s for $M_s V_i$. The total free energy of particle i in a system of N single-domain particles (ignoring exchange energy between surfaces, see for example [84]):

$$E(\mathbf{m}_i) = -V_i M_s \mathbf{m}_i \cdot \mathbf{B} - K_i V (\mathbf{m}_i \cdot \mathbf{e}_i)^2 - \sum_{j \neq i} \frac{\mu_0 M_s^2 V_i V_j}{4\pi |R_{ij}|^3} (3(\mathbf{m}_i \cdot \mathbf{r}_{ij})(\mathbf{m}_j \cdot \mathbf{r}_{ij}) - \mathbf{m}_i \cdot \mathbf{m}_j) \quad (2.10)$$

where $K_i = k_i/V_i$ is the anisotropy constant in Jm^{-3} . Note that the model assumes that the applied magnetic field \mathbf{B} is uniform for all particles in the system. The effective field experienced by particle i is obtained from equation (2.6):

$$\mathbf{B}_{\text{eff},i} = -\frac{1}{M_s V} \frac{\partial E(\mathbf{m}_i)}{\partial \mathbf{m}_i} = \mathbf{B} + \frac{2K_i}{M_s} (\mathbf{m}_i \cdot \mathbf{e}_i) \mathbf{e}_i + \sum_{j \neq i} \frac{\mu_0 M_s V_j}{4\pi |R_{ij}|^3} (3(\mathbf{m}_j \cdot \mathbf{r}_{ij}) \mathbf{r}_{ij} - \mathbf{m}_j) \quad (2.11)$$

2.3 Stochastic processes

An overview of the theory of stochastic processes is necessary before discussing the effect of temperature on the dynamics of magnetic moments. Stochastic processes are a mathematical description of randomly fluctuating phenomena, which must be specified using probability theory. More formally, a stochastic process is a set variables $\{X(t)\}_{t \in T}$ where $X(t)$ is a collection of random variables and t usually denotes time where $T = [0, \infty)$. Since $X(t)$ is a collection of random variables the stochastic process has many possible realisations. There are many examples of stochastic processes in physics such as decay in radioactive materials, noise in electronic circuits and the thermal fluctuations of magnetic moments. The prototypical stochastic process is a random walk called *Brownian motion* after Robert Brown observed and described the random paths of pollen in a fluid. These random fluctuations were later described by Einstein as being the result of “exceedingly frequent forces” impacting on the system, which lead to such complicated behaviour that their effect may only be described probabilistically [68].

2.3.1 Stochastic differential equations

Dynamical systems with deterministic dynamics and random fluctuations are described by stochastic differential equations (SDEs). The Langevin equation is an SDE that describes a diffusion process and will be used in Section 2.4.1 to describe the dynamics of a magnetic moment subjected to thermal fluctuations:

$$\frac{d\mathbf{x}}{dt} = \mathbf{a}(\mathbf{x}, t) + B(\mathbf{x}, t)\boldsymbol{\xi}, \quad \boldsymbol{\xi} \sim \mathcal{N}\{\mathbf{0}, 2DI_m\} \quad (2.12)$$

$\mathbf{x} \in \mathbb{R}^n$ is the system state, $\boldsymbol{\xi} \in \mathbb{R}^m$ is a vector of independent normally distributed random variables, I_m is the identity matrix, and $\mathbf{a}: \mathbb{R}^n \times [0, T] \rightarrow \mathbb{R}^n$ and $B: \mathbb{R}^n \times [0, T] \rightarrow \mathbb{R}^{n \times m}$ are called the drift and diffusion respectively. The drift and diffusion are nonlinear functions that describe the deterministic and stochastic dynamics respectively. Equation (2.12) may be found in its integral form:

$$d\mathbf{x} = \mathbf{a}(\mathbf{x}, t)dt + B(\mathbf{x}, t)d\mathbf{W}(t) \quad (2.13)$$

where $\mathbf{W} \in \mathbb{R}^m$ is the m -dimensional Wiener process. The Wiener process can be described intuitively as the continuous limit of a series of random steps. The Wiener process begins with $\mathbf{W}(t_0) = \mathbf{0}$ at $t_0 = 0$ and over a small time δt takes a random step $\mathbf{z} = \mathbf{W}(t + \delta t) - \mathbf{W}(t)$, which is a vector of independent normally distributed variables with zero mean and variance δt (i.e. $\mathbf{z} \sim \mathcal{N}\{\mathbf{0}, \delta t I_m\}$). The Wiener process is obtained by taking the limit $\delta t \rightarrow 0$, which describes a series of small and very frequent random steps. Although the Wiener process is continuous, it is nowhere differentiable and thus equation (2.12) is phenomenological in nature. The solution of the Langevin equation

is:

$$\mathbf{x}(T) = \mathbf{x}(t_0) + \int_{t_0}^T \mathbf{a}(\mathbf{x}, t) dt + \int_{t_0}^T B(\mathbf{x}, t) d\mathbf{W}(t) \quad (2.14)$$

The integral of the Wiener function, the third term of equation (2.14), is a stochastic integral and requires the use of stochastic calculus.

The stochastic integral is the stochastic analogue of the deterministic Riemann integral. However, unlike the Riemann integral, the stochastic integral may be defined in one of two ways, leading to the Itô-Stratonovich dilemma. The mathematical definitions and fundamental differences between the two interpretations is described in Appendix A. One must be careful to choose the correct approach when modelling any diffusion process. The accepted notation to signal the Stratonovich interpretation over the Itô interpretation is to augment the stochastic integral operator in equation (2.13) with ‘ \circ ’ and alter the drift term:

$$d\mathbf{x} = \underline{\mathbf{a}}(\mathbf{x}, t) dt + B(\mathbf{x}, t) \circ d\mathbf{W}(t) \quad (2.15)$$

An Itô stochastic differential equation may be converted into an equivalent Stratonovich stochastic differential equation (and *vice-versa*) using the *Itô correction*:

$$\underline{a}_i(\mathbf{x}, t) = a_i(\mathbf{x}, t) - \frac{1}{2} \sum_{k=1}^m \sum_{j=1}^n \frac{\partial b_{ik}(\mathbf{x}, t)}{\partial x_j} b_{jk}(\mathbf{x}, t) \quad (2.16)$$

An important property of stochastic differential equations is *commutative* noise. The condition for commutative noise is [115]:

$$\sum_{i=1}^d b_{i,j_1} \frac{\partial b_{k,j_2}}{\partial x_i} = \sum_{i=1}^d b_{i,j_2} \frac{\partial b_{k,j_1}}{\partial x_i} \quad (2.17)$$

for all i, j where $b_{i,j}$ and x_i are elements of the diffusion matrix and state vector respectively. SDEs with a commutative noise are often easier to solve analytically and simulate numerically.

2.3.2 Fokker-Planck equation

The state of a system \mathbf{x} with dynamics described by a stochastic differential equation is a random variable. Starting from the same initial condition, a stochastic system can take many possible trajectories. It is rarely useful to obtain just one possible trajectory for a system. In order to answer questions such as: “which trajectory is most likely?”, “which trajectory should we expect?”, and “how much do the possible trajectories vary?” the *probability distribution* of the state trajectories is required $p(\mathbf{x}, t)$.

Each Langevin equation corresponds to a Fokker-Planck equation, which describes the time-evolution of the probability density function of the system state. Recalling the




Figure 2.2: The time-evolution of the Wiener process. The black paths show possible realisations (solutions) to the stochastic differential equation describing the Wiener process. The Fokker-Planck equation describes the probability distribution over the possible states at a time t and is shown as a heat map.

vector-valued drift function and matrix-valued diffusion function from equation (2.12), the corresponding Itô Fokker-Planck equation is:

$$\frac{\partial p(\mathbf{x}, t)}{\partial t} = - \sum_{i=1}^N a_i(\mathbf{x}, t) \frac{\partial p(\mathbf{x}, t)}{\partial x_i} + \sum_{i=1}^N \sum_{j=1}^N \frac{1}{2} b_{ij}(\mathbf{x}, t) \frac{\partial^2 p(\mathbf{x}, t)}{\partial x_i \partial x_j} \quad (2.18)$$

where subscripts denote the vector/matrix indices. Figure 2.2 shows solutions to the SDE and corresponding Fokker-Planck equation for a standard Wiener process. The initial condition of the process is $W(t_0) = 0$ with probability $p(0) = 1$ at time $t_0 = 0$. As time progresses from the initial condition, the probability distribution diffuses - i.e. becomes more flat. Near $t = 0$ the probability that the Wiener process takes a value near zero is quite high. This can be seen from the six different realisations of the Wiener process SDE solution (black lines). The initial state of the system is $W(0) = 0$ and the solutions remain close initially but diverge from one another as they evolve. Note that the Fokker-Planck equation is a *deterministic* partial differential equation, although it represents a *stochastic* process.

2.3.3 Jump processes

The Fokker-Planck equation represents a stochastic process that diffuses with time but a process may also exhibit “randomness” in the form of discrete jumps. Jump processes

are applicable to both systems with a continuous state $\mathbf{x} \in \mathbb{R}^n$ and systems with a set of N discrete states $\mathbf{x} \in \mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2 \dots \mathbf{x}_N\}$. The trajectory of a jump processes is continuous in time but its value can take discrete jumps. The process is defined by the transition rate from each possible state to all other states. Informally, generating a trajectory from a jump process can be described as follows:

1. Begin with an initial state \mathbf{x}_{now} at time $t = t_0$
2. Draw a random variable Δt (the waiting time) from a continuous probability distribution over time, parameterised by the transition rates.
3. Draw a second random variable \mathbf{x}_{next} from a continuous probability distribution (or discrete probability mass function) over the states, parameterised by the transition rates.
4. At time $t = t_0 + \Delta t$ the system jumps from \mathbf{x}_{now} to \mathbf{x}_{next}

An exact algorithm to generate trajectories from a jump process over a discrete state space was given by Gillespie [76] and such methods are more generally known as kinetic Monte-Carlo methods [197]. Given an initial state and a set of transition rates between states, kinetic Monte-Carlo methods simulate realisations from the jump process. Simulation of stochastic processes are discussed more generally in Section 2.3.6.

2.3.4 Master equation

Just as the Fokker-Planck equation describes the probability distribution of the trajectories for a diffusion process, the master equation describes the probability distribution (or probability mass function for a discrete state-space) of the trajectories for a jump process. The remainder of this thesis will only consider jump processes that take values on a set of discrete possible states. Each possible state is labelled x_i where $i = 1, 2, \dots, N$. In this case, the master equation reduces to a set of N first order linear ordinary differential equations:

$$\frac{dp(x_i, t)}{dt} = \sum_j [w_{j \rightarrow i} p(x_j, t) - w_{i \rightarrow j} p(x_i, t)] \quad (2.19)$$

A jump process with N possible states will have 2^N possible transition paths and therefore 2^N transition rates. Since, the probability may be interpreted as the proportion of members of a statistical ensemble occupying the state x_i , the product $w_{i \rightarrow j} p(x_i, t)$ represents the number of members that transition from state x_i to state x_j at time t . Therefore the probability flow into state x_i in equation (2.19) is the total number of members that transition into state x_i minus the total number of members that transition out of state x_i to another state x_j .

2.3.5 Canonical ensemble

An ensemble of identical isolated systems that are in thermal equilibrium with a heat bath of temperature T is termed the canonical ensemble. In other words, the system has reached a statistical equilibrium such that the probability distribution over the states is time-invariant. The probability that the system will be found in state \mathbf{x} with free energy $E(\mathbf{x})$ is in fact given by the Boltzmann distribution [72]:

$$p(\mathbf{x}) = \frac{e^{-\beta E(\mathbf{x})}}{Z} \quad (2.20)$$

where $\beta = k_B T$, k_B is the Boltzmann constant and Z is the scalar-valued partition function:

$$Z = \int_{x_1} \int_{x_2} \cdots \int_{x_n} p(x_1, x_2, \dots, x_n, t) dx_n \cdots dx_2 dx_1 \quad (2.21)$$

For many systems, the Boltzmann distribution represents the stationary distribution $\frac{\partial p(\mathbf{x}, t)}{\partial t} = 0$ of the Fokker-Planck or master equation. The canonical ensemble is a result of equilibrium thermodynamics, which is not applicable to dynamical systems.

2.3.6 Simulation of stochastic processes

2.3.6.1 Evolution of probability distributions

Many stochastic processes do not have known analytic solutions and require numerical methods to simulate approximate solutions. The most descriptive numerical simulation is to approximate the probability distribution of the system states at discrete points in time. Numerical solutions to the Fokker-Planck may be obtained by the use of finite-element [130] and finite-differences methods [128], similarly to any parabolic partial differential equation. For a discrete-state system, the master equation may be solved using techniques for an n -dimensional system of linear ordinary differential equations, such as Euler and Runge-Kutta methods [164].

Obtaining approximate solutions to the Fokker-Planck equation for systems with high dimensionality may be prohibitively expensive. This is a consequence of the increased computational difficulty of solving partial differential equations. However, the Feynmann-Kac theory says that a deterministic partial differential equation problem can be transformed into a stochastic differential equation problem [82]. For Fokker-Planck equations, this amounts to simulating many possible trajectories of the corresponding Langevin equation to obtain an approximate distribution of trajectories. In the limit of simulating infinite different possible trajectories of the stochastic process, the two results are equivalent [82]. This approach to simulating stochastic processes is termed *stochastic*

simulation and is a type of Monte-Carlo method. Similarly, high dimensionality master equations may be solved more efficiently using the kinetic Monte-Carlo or Gillespie methods [75–78], which generate trajectories for jump processes.

2.3.6.2 Trajectories of stochastic differential equation

Numerical methods for simulating stochastic differential equations are analogue to those for deterministic differential equations. Given an initial state, the objective is to obtain an estimate of the state of the system at discrete time intervals. The fundamental difference is that solutions to SDEs are dependent on a specific realisation of a Wiener process. Multiple solutions to the SDE are simulated by using different Wiener processes in each case. In practice, this usually amounts to using a different seed for the random number generator used by the computational method to generate the Wiener process. An advantage of stochastic simulation is that it is inherently *embarrassingly parallel* - i.e. each realisation is independent of the others and therefore they can be computed simultaneously on multiple processors. Specific implementations of numerical methods for obtaining solutions to SDEs are given in Chapter 4.

The aim of any numerical SDE solver is to approximate the true solution $\mathbf{x}(t)$ of the Langevin equation (2.13) at discrete time intervals $t_0 < t_1 < t_2 < \dots < t_N$ given an initial condition $\mathbf{x}_0 = \mathbf{x}(t_0)$. The resulting discrete time approximation is a set of values $\{\mathbf{y}_0, \mathbf{y}_1, \dots, \mathbf{y}_N\}$, which is conditional on a Wiener process \mathbf{W}_t on the period $t_0 \leq t \leq t_N$. The fundamental requirement of the solver is to minimise the error between the approximation \mathbf{y}_i and the true solution $\mathbf{x}(t_i)$ while also minimising computational resources.

The convergence and stability properties of a numerical SDE solver are crucial to ensure robust performance in practice. Numerical methods that have proven convergence and stability may be applied with confidence to problems that have no analytic solution.

2.3.6.3 Convergence

A natural requirement of any numerical method solving an initial value problem is that the approximate solution should approach the true solution with a decreasing time step $\Delta t \rightarrow 0$. A numerical approximation $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n$ to an SDE may converge to the true solution in a *strong* or a *weak* sense. A numerical scheme converges strongly with order γ if the following criterion is satisfied:

$$E \langle \|\mathbf{x}(t_n) - \mathbf{y}_n\| \rangle \leq C \Delta t^\gamma \quad (2.22)$$

for some fixed constant C and γ . In other words, the numerical approximation for the global solution \mathbf{y}_n at time t_n approaches the true value $\mathbf{x}(t_n)$ as the time step

Δt reduces. Strong convergence is analogous to the deterministic case: as the time step decreases, the true and approximate state trajectories converge. Strong solutions, which rely on accurate simulation of the Wiener process, are necessary when path-wise information (such as for hitting times, first-passage times, etc.) is needed. In order to understand the full dynamics of magnetic nanoparticles, path-wise accuracy is an important requirement.

A numerical scheme converges weakly with order β if:

$$\|E \langle \mathbf{x}(t_n) \rangle - E \langle \mathbf{y}_n \rangle\| \leq C \Delta t^\beta \quad (2.23)$$

Equation (2.23) requires that the expected value of the approximation converges to the expected value of the true solution. Note that this is a much weaker requirement than for strong convergence. In the context of weak stochastic simulation, the average value of a large ensemble of trajectories will approach the true expected solution but the individual solutions will not be an accurate representation of the true solutions. Weak numerical solvers, designed only to achieve weak convergence, exploit this weaker constraint by relying on approximate Brownian paths (e.g. a binomial branching processes). These weak solutions are computationally efficient to generate and often converge faster than strong solutions.

Mathematical proofs of the convergence of numerical schemes are challenging to derive. A common approach is to determine the *consistency* of an approximation, which is sufficient to determine convergence - see [115] for more details. A thorough test of a numerical method's implementation is to compute the empirical convergence in a computational experiment and compare the results to the expected theoretical convergence.

2.3.6.4 Stability

The stability of a numerical scheme describes whether the errors of the scheme are bounded or diverge. There are a number of definitions of stability (A-stability, M-stability, T-stability, etc.) and the stability of a solution depends on the numerical scheme, the time step, and the underlying SDE that is being simulated. The basic stability properties of deterministic and stochastic numerical schemes are understood by studying simple test functions.

In the following, we will use the MS-stability (mean-square stability). MS-stability says informally that approximate solutions that begin close to the true solution remain within some bounded error. The test equation is the following scalar, linear, time-invariant, Itô SDE:

$$dy(t) = \mu y(t)dt + \sigma y(t)dW_t, \quad \mu, \sigma \in \mathbb{R} \quad (2.24)$$

The system in equation 2.24 is MS-stable if and only if:

$$\lim_{t \rightarrow \infty} E \langle y^2(t) \rangle = 0, \quad y(0) = \delta_0 \quad (2.25)$$

for any arbitrarily small δ_0 . The test equation can be shown to be MS-stable with initial condition $y(0) = 1$ with probability 1 if $\mu - \sigma^2/2 < 1$.

A stable numerical scheme results in an approximate trajectory of the test equation $\{y_1, y_2, \dots, y_n\}$ that is also MS-stable such that $\lim_{n \rightarrow \infty} y_n = 0$ [23]. A numerical scheme is said to be MS-stable for the test equation (2.24) if and only if:

$$R \triangleq \frac{E \langle y_{n+1}^2 \rangle}{E \langle y_n^2 \rangle} \leq 1 \quad (2.26)$$

where R is defined and called the *stability function*. The region of stability is the three dimensional space $\{(\mu, \sigma, \Delta t) : R < 1\}$ that is the set of parameters for which the condition $R < 1$ holds true.

Region of stability: Euler The simplest Itô numerical scheme is the explicit (or forward) Euler method: $y_{n+1} = y_n + a(y_n, t_n)\Delta t + b(y_n, t_n)\Delta W_n$ where $\Delta W_n = W_{t_{n+1}} - W_{t_n}$ and $a(\cdot), b(\cdot)$ are the drift and diffusion respectively. Substituting the test equation (2.24) into the Euler method obtains the one-step predictor for the method:

$$y_{n+1} = y_n + \mu y_n \Delta t + \sigma y_n \Delta W_n \quad (2.27)$$

The one-step predictor corresponds to the following stability function:

$$R = (1 + p)^2 + q^2 < 1 \quad (2.28)$$

where $p = \mu\Delta t$ and $q = \sigma\sqrt{\Delta t}$. Therefore, it is necessary that $-2 < p < 0$ for the explicit Euler method to be MS-stable; the time step must be well matched to the time scale of the deterministic dynamics such that $\Delta t \approx -1/\mu$. However this condition is not sufficient, the stochastic component must also be low relative to the deterministic dynamics such that $q^2 < -2p - p^2$.

Region of stability: semi-implicit Euler The semi-implicit Euler method introduces implicitness into the deterministic component only: $y_{n+1} = y_n + a(y_{n+1}, t_{n+1})\Delta t + b(y_n, t_n)\Delta W_n$ giving the stability function:

$$R = \frac{1 + q^2}{(1 - p)^2} < 1 \quad (2.29)$$

In contrast to the explicit Euler method, the semi-implicit Euler method requires $p < 0$ such that the time step may be significantly smaller or larger than the deterministic

dynamics. The condition $p < 0$ is true as long as $\mu < 0$ in the test equation. Systems with relatively high noise will still cause instability since $q^2 < p^2 - 2p$. Smaller time steps allow larger noise intensity σ .

Region of stability: fully implicit A naive generalisation of the implicit (backward) Euler scheme for ODEs to SDEs is implicit in the drift and diffusion terms such that $y_{n+1} = y_n + a(y_{n+1}, t_{n+1})\Delta t + b(y_{n+1}, t_{n+1})\Delta W_n$. The one step predictor now gives:

$$y_{n+1} = \frac{1}{1 - \mu\Delta t - \sigma\Delta W_n} y_n \quad (2.30)$$

This step could be undefined or very large for a small ΔW_n . In fact, it can be shown from equation 2.30 that $E\langle y_{n+1} \rangle \rightarrow \infty$ [142]. This presents a fundamental challenge for designing fully implicit methods for SDEs. However, a number of solutions have been proposed to handle this problem [24]. In Section 4.3.3 we describe a fully-implicit method that uses a truncated Wiener process.

2.3.6.5 Systems in equilibrium state

When the dynamical approach to equilibrium is not relevant, Markov-chain Monte-Carlo (MCMC) methods may be used to approximate the equilibrium distribution of the stochastic process. MCMC methods approximate the distribution of the canonical ensemble without requiring the partition function Z be computed (which becomes prohibitively expensive for large systems). MCMC is a sampling method, which draws samples from the equilibrium distribution, which are guaranteed to converge to the true distribution. MCMC may be thought of as a more efficient integration method. Commonly used algorithms for MCMC are the Metropolis method [168], Hamiltonian Monte-Carlo [150], and the No U-Turn sampler [94].

2.4 Finite-temperature magnetism

The non-equilibrium behaviour of a magnetic moment were previously described as a damped precession around an externally applied field, eventually aligning with the field. In reality, a magnetic moment will fluctuate randomly due to thermal effects, which are due to degrees of freedom that are not explicitly present in the physical model. The origins of the thermal fluctuations are the same origin as the damping effect (which is heuristically accounted for in the Landau-Lifshitz-Gilbert equation) and are formally related through the fluctuation-dissipation theorem. The exact origin of these fluctuations are beyond the scope of this thesis. Consequently, the dynamics of a magnetic moment (or single domain magnetic nanoparticle) are described by a stochastic process.

2.4.1 Stochastic Landau-Lifshitz-Gilbert equation

The Landau-Lifshitz-Gilbert equation is augmented with a randomly fluctuating thermal field $\mathbf{H}_{\text{th}} \in \mathbb{R}^3$ to account for thermal fluctuations. The thermal field is a vector of independent and normally distributed random variables such that:

$$\langle \mathbf{H}_{\text{th}}^i(t), \mathbf{H}_{\text{th}}^j(t') \rangle = 2D\delta_{ij}\delta(t-t') \quad (2.31)$$

where the strength of the fluctuations, a scalar D , has yet to be determined. Moreover, the thermal field vector at each atomic site is independent. The stochastic Landau-Lifshitz-Gilbert (sLLG) equation is obtained by adding the thermal field contribution to the effective field (where we have used the H-field $\mathbf{H}_{\text{eff}} = \mu_0\mathbf{B}_{\text{eff}}$):

$$\frac{d\mathbf{m}}{dt} = -\frac{\gamma}{1+\alpha^2} (\mathbf{m} \times (\mathbf{H}_{\text{eff}} + \mathbf{H}_{\text{th}})) - \frac{\alpha\gamma}{1+\alpha^2} \mathbf{m} \times (\mathbf{m} \times (\mathbf{H}_{\text{eff}} + \mathbf{H}_{\text{th}})) \quad (2.32)$$

The stochastic Landau-Lifshitz-Gilbert equation is taken under the Stratonovich interpretation and rearranging into integral form gives:

$$\mathbf{m}(t) = \int \underline{\mathbf{a}}(\mathbf{m}, t) dt + \int B(\mathbf{m}, t) \circ d\mathbf{W}(t) \quad (2.33)$$

where the drift and diffusion coefficients are:

$$\underline{\mathbf{a}}(\mathbf{m}, t) = \frac{-\gamma}{1+\alpha^2} (\mathbf{m} \times \mathbf{H}_{\text{eff}}) - \frac{\alpha\gamma}{1+\alpha^2} \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}}) \quad (2.34)$$

$$B(\mathbf{m}, t) = \sqrt{2D} \frac{-\gamma}{1+\alpha^2} (\mathbf{m} \times) - \sqrt{2D} \frac{\alpha\gamma}{1+\alpha^2} \mathbf{m} \times (\mathbf{m} \times) \quad (2.35)$$

$$(2.36)$$

The stochastic Landau-Lifshitz-Gilbert equation belongs to a class of complex SDEs, which are particularly difficult to simulate and are discussed in more detail in Chapter 4. The corresponding Fokker-Planck equation for the stochastic Landau-Lifshitz-Gilbert equation is [22]:

$$\frac{\partial}{\partial t} p(\mathbf{m}; t) = \frac{\partial}{\partial \mathbf{m}} \left\{ \left[-\frac{-\gamma}{1+\alpha^2} \mathbf{m} \times \mathbf{H}_{\text{eff}} - \frac{\alpha\gamma}{M_s(1+\alpha^2)} \right. \right. \\ \left. \left. + 2D\gamma k_B \mathbf{m} \times \left(\mathbf{m} \times \frac{\partial}{\partial \mathbf{m}} \right) \right] p(\mathbf{m}; t) \right\} \quad (2.37)$$

2.4.2 Thermal equilibrium

At zero-temperature, in the absence of any externally applied forces, a magnetic system will reach a locally stable equilibrium due to the dissipative term in the LLG equation. However, when temperature is introduced, the isolated system will randomly sample states around the equilibrium due to coupling with the heat bath. There is a finite

probability that the system will be found in any one of the possible states. By considering an ensemble of identical systems, each randomly fluctuating, then at any point in time we have a distribution of states. In the absence of any external forces, this distribution becomes stationary: the states of each system are still fluctuating but the distribution over the statistical ensemble is time-invariant.

In thermal equilibrium, a system of interacting magnetic moments satisfy the conditions for a canonical ensemble. Therefore, it is known that the stationary distribution of the system should follow the Boltzmann distribution:

$$\frac{\partial}{\partial \mathbf{m}} \left\{ \left[-\frac{-\gamma}{1+\alpha^2} \mathbf{m} \times \mathbf{H}_{\text{eff}} - \frac{\alpha\gamma}{1+\alpha^2} + 2D\gamma K \mathbf{m} \times \left(\mathbf{m} \times \frac{\partial}{\partial \mathbf{m}} \right) \right] Z^{-1} e^{-E(\mathbf{m},t)/(k_B T)} \right\} = 0 \quad (2.38)$$

Rearranging this equation yields the value of the thermal strength D :

$$D = \frac{2\alpha k_B T}{V\gamma M_s} \quad (2.39)$$

An alternative approach to obtaining the strength of the fluctuations D is to make use of the fluctuation-dissipation theorem. A discussion of the two approaches can be found in [57].

Although the expression for the distribution in equilibrium is known (the Boltzmann distribution), its value in large systems of interacting magnetic moments (or single domain particles) cannot be computed because of the difficulty of computing Z . The partition function Z requires evaluating the multidimensional integral (equation (2.21)), which becomes prohibitively expensive to approximate with numerical methods. Therefore the equilibrium state systems containing many magnetic moments (particles) must be obtained using MCMC methods as in Section 2.3.6. Using MCMC methods to evaluate the equilibrium distribution of a large system is significantly more efficient than stochastic simulation of the Fokker-Planck equation to equilibrium.

2.4.3 Thermal activation

A magnetic moment in a local energy minima, subjected to thermal fluctuations with low energy in comparison to the energy required to escape the minima, will have a very low but finite chance of escaping the well. These states are *meta-stable* states because they appear to be stable for very long periods of time. Occasionally the stochastic fluctuations will be large enough that the moment will escape from the minima and transition to another minima, a process referred to as thermal activation. There is a large separation of time scales between the escape time (long) and local dynamics in the well (short). Under these conditions, a *rate* description of the system is a valid approximation to the

dynamics [86]. The full dynamics of the magnetic moment are reduced to a Markovian jump process over a discrete set of meta-stable states.

2.4.3.1 Single magnetic nanoparticles

As seen previously, the time evolution of the probability mass function over the discrete set of meta-stable states is given by the master equation (2.19). The transition rates in the master equation are all that is required to simulate the jump process. However, reducing the Landau-Lifshitz-Gilbert dynamics to a single transition rate is a non-trivial problem and is an ongoing subject of study [21, 22, 37, 38, 45, 108, 109, 118, 151].

The simplest model of thermal activation was first proposed by Néel [151] and was based on the work of Arrhenius on reaction rate theory in chemistry [8]: The free energy of a magnetic nanoparticle with a field $h(t)$ at time t applied along its anisotropy axis is described by the Stoner-Wohlfarth model (Section 2.2.2) and is approximated with two possible metastable states $\theta_1(t), \theta_2(t)$ corresponding to the minima of equation (2.9) with $h(t) < 1$. The particle dynamics are modelled as jumps from orientation $\theta_i(t)$ to orientation $\theta_j(t)$ according to the transition rate $\Gamma_{i \rightarrow j}(t)$. The corresponding master equation obtained from equation (2.19):

$$\frac{d}{dt} \begin{bmatrix} p(\theta_1; t) \\ p(\theta_2; t) \end{bmatrix} = \begin{bmatrix} -\Gamma_{1 \rightarrow 2}(t) & \Gamma_{2 \rightarrow 1}(t) \\ \Gamma_{1 \rightarrow 2}(t) & -\Gamma_{2 \rightarrow 1}(t) \end{bmatrix} \begin{bmatrix} p(\theta_1; t) \\ p(\theta_2; t) \end{bmatrix} \quad (2.40)$$

where $p(\theta_1; t) + p(\theta_2; t) = 1$ for any t and the magnetisation is computed from the orientation probabilities $M(t) = M_s(p(\theta_1; t) - p(\theta_2; t))$. The transition rates relate to the relaxation time of the particle magnetisation in a time-independent externally applied field:

$$\tau \approx \Gamma_{1 \rightarrow 2}^{-1} + \Gamma_{2 \rightarrow 1}^{-1} \quad (2.41)$$

The Néel-Arrhenius law relates the transition rate between the two states to their energy barrier (equation (2.9)):

$$\Gamma_{i \rightarrow j}(t) = f_0 e^{-\sigma(1 \pm h(t))^2} \quad (2.42)$$

where f_0 is called the attempt frequency. The attempt frequency is often assumed as a constant value between $10^9 - 10^{11}$ as obtained by experiments [22]. However, Brown [21, 22] later derived an expression for the attempt frequency for a Stoner-Wohlfarth particle using the Kramers theory [118], leading to the Néel-Brown model:

$$\Gamma_{i \rightarrow j}(t) = \frac{2\gamma\alpha k_B T \sigma^{1.5} (1 - h^2(t))}{V M_s \sqrt{\pi} (1 + \alpha^2)} (1 \pm h(t)) e^{-\sigma(1 \pm h(t))^2} \quad (2.43)$$

where $\gamma = 1.76086 \times 10^{11}$ is the magnetogyric ratio. The Néel-Brown model is only valid when the energy barrier is sufficiently large, $\sigma(1 - h)^2 \gg 1$, to maintain the two

well structure and ensure a separation of time scales. Computing transition rates for magnetic particles subjected to an oblique magnetic field requires additional complexity and has been addressed in [37, 71].

Larger particles or particles at low temperature are stable and experience very few thermal activations. The transition rates $\Gamma_{i \rightarrow j}$ are infrequent to negligible and the relaxation time τ is large compared to the observation time of the magnetisation. In this case, the particle magnetisation will appear frozen. However, particles with low stability have many transitions and fast relaxation times compared to the observation time. In this case, the particle magnetisation will fluctuate during the observation window. A large ensemble of fluctuating particles will appear as having net zero magnetisation since the particles are randomly fluctuating. However, the transition rates can be influenced by an applied field h in equation (2.43) and the particles can be stabilised. In this case the ensemble will experience a net magnetisation until the field is removed. Therefore, the ensemble exhibits a similar behaviour to paramagnetic materials as described in Section 2.2.1. The non-hysteretic response of an ensemble of single domain magnetic nanoparticles is termed *superparamagnetism*. Superparamagnetic materials distinct from bulk paramagnetic materials by exhibiting strongly aligned magnetic moments and very large magnetisation values.

2.4.3.2 Weakly interacting magnetic nanoparticles

The thermal activation of a single particle may be extended to weakly interacting ensembles of particles. Two particles are weakly interacting if the dipole-dipole effective field is weak enough such that the two-level approximation of the previous section applies to each particle. These models have been simulated using kinetic Monte-Carlo techniques, which are described in [189].

2.4.3.3 Complex magnetic systems

Magnetic materials with more complex domain structures have also been treated with thermal activation. The multidimensional energy landscape for these materials are so complex that there are currently no expressions of the transition rates. Rather, numerical techniques are first used to find local minima (meta-stable states) by using energy minimisation methods such as gradient descent or micromagnetics [60]. The transition rates between these meta-stable states are computed using the Arrhenius relationship and the energy barriers are obtained using methods such as the nudged elastic band method [88].

2.4.4 Linear response theory

Linear response theory describes the response of systems in equilibrium to small perturbations from an external driving force. For an overview of linear systems in general and linear response calculations in Hamiltonian systems specifically see [205] and [166] respectively. For a thorough review of linear response models for single magnetic particles see [65].

The magnetisation in thermal equilibrium of an ensemble of Stoner-Wohlfarth particles (Section 2.2.2) subjected to a small constant externally applied field approximately follows a linear function of the applied field magnitude:

$$M_\infty = \chi_0 H \quad (2.44)$$

where χ_0 is called the equilibrium (or static) susceptibility. This relation is approximately true in the general case that H is *small enough*. In the particular case of an ensemble of Stoner-Wohlfarth particles with their anisotropy axes aligned with the externally applied field, the static susceptibility may be obtained in two cases. When the reduced energy barrier $\sigma = 0$ then $\chi_0 = \mu_0 M_s^2 V / (3k_B T)$. Alternatively, in the limit where $\sigma \rightarrow \infty$ then $\chi_0 = \mu_0 M_s^2 V / (k_B T)$.

The non-equilibrium magnetisation response of an ensemble of particles to a time-varying applied field is also predicted by linear response theory. A first-order linear system subjected to a harmonic force $H(t) = H_0 \cos(\omega t)$ is also a harmonic of equal frequency but with shifted amplitude and phase [183, 205]:

$$M(t) = \frac{\chi_0}{1 + j\omega\tau} H_0 \cos(\omega t) \quad (2.45)$$

where the gain term χ_0 is the static susceptibility and the time constant τ is the Néel relaxation time in equations (2.42) (2.41) [27, 165]. Recall from Section 2.4.3.1 that the Néel relaxation time additionally requires that $\sigma(1-h)^2 \gg 1$. These results have a small domain of validity but have proved useful in models of heat dissipation (Section 3). For higher-order approximations of nanoparticle dynamics see [65].

2.4.5 Heat and work

Work must be done in order to change the magnetisation of a material. A magnetic material subjected to an external field source $\mathbf{B}_0 = \mu_0 \mathbf{H}_0$, where $\mu_0 = 1/(4\pi) \times 10^{-7} \text{kg s}^{-2} \text{A}^{-2}$ is the permeability of a vacuum, will produce a magnetic field in response $\mathbf{B} = \mu_0(\mathbf{M} + \mathbf{H})$ where $\mathbf{H} = \mathbf{H}_0$ is the magnetic field not due to the presence of the material. Recall from Section 2.2.1 that the magnetisation is a potentially multi-valued nonlinear function of the applied field. Changing the magnetic field produced by the

material \mathbf{B} requires doing work to the system through the applied field \mathbf{H} . The first law of thermodynamics states that for an isolated system:

$$\Delta U = \delta W + \delta Q \quad (2.46)$$

in other words the change in internal energy U is the sum of the work done to the system δW and the heat added to the system Q . As the external field changes the magnetic field of the material, magnetic work must be performed by the underlying circuit driving the magnetic field (i.e. the magnetic field generator), which is defined as [188]:

$$\Delta U = \int dv \int \mathbf{B} \cdot d\mathbf{H} + \delta Q \quad (2.47)$$

where $\int dv$ represents the integral over the volume of the magnetic material. Thus the work done to the system will lead to an increase in internal energy or heat dissipated from the system.

In a cyclical process, the system returns to its initial state such that the internal energy, which is a function only of the state of the system, is unchanged. Therefore, the work done to the system is equal to the negative heat added $\delta W = -\delta Q$. Therefore, the heat added to the system¹:

$$Q = - \oint \mathbf{B} \cdot \mathbf{H} \quad (2.48)$$

$$= -\mu_0 \oint \mathbf{M} \cdot \mathbf{H} = -\mu_0 \oint \mathbf{H} \cdot \mathbf{M} \quad (2.49)$$

where \oint represents an integral on a closed line in the B-H parametric plot and we have used integration by parts. Note that this is exactly the area enclosed by a dynamic hysteresis loop. Therefore, the heat dissipated by an arbitrary magnetic system may be simply determined from the loop area. If a magnetic system is subjected to a continually alternating magnetic field, equation (2.48) represents the heat dissipated during a single cycle of the field. A common metric for hyperthermia is the specific power loss (SPL) measured in Wg^{-1} :

$$\text{SPL} = \frac{f\mu_0}{\rho} \oint \mathbf{H} \cdot \mathbf{M} \quad (2.50)$$

where f is the frequency of the applied alternating field and ρ is the density of the material (in gm^{-3}).

¹Note that we will assume that a closed loop in the M - H plane represents a cyclical process. However the macroscopic variable, magnetisation M , may have degenerate microscopic states (the individual orientations of the magnetic moments in the system)

Chapter 3

Literature Review: Magnetic particle hyperthermia

3.1 Therapeutic applications

Magnetic hyperthermia is a biomedical technique to induce localised heating in tissue mediated by magnetic nanoparticles. Hyperthermia techniques use the application of heat to destroy, palliate, or otherwise treat cancerous tumours of the internal organs and may be applied locally (i.e. directly to the tumour) or systemically. Methods of heating include thermal conduction by the direct application of heat as well as noninvasive methods such as electromagnetic radiation and acoustic waves [43]. An alternative noninvasive approach has been the use of localised magnetic nanoparticles within the tumour, which dissipate heat when exposed to an external alternating magnetic field. The use of magnetic nanoparticles as heat mediators in hyperthermia was proposed in 1957 [74] but has only recently gained significant interest [63, 160].

The fundamental objective of magnetic hyperthermia is to destroy cancerous tissue either by inducing cell necrosis (immediate cell death) or cell apoptosis (an irreversible process of controlled cell death) [33] as an alternative to ablative surgery. The temperature range for apoptosis is around 41–46°C while necrosis requires $\geq 46^\circ\text{C}$ [119]; the precise requirements for heating the tumour is an area of active research [51]. However, magnetic hyperthermia may also enhance the effects of conventional therapies through increased localised heating [80]. General hyperthermia techniques have already been shown to improve the effectiveness of radiotherapy and chemotherapy [43, 91]. The ability of increased temperatures to enhance immunotherapy is another active area of research [191]. Furthermore, the magnetic nanoparticles administered for magnetic hyperthermia have properties that serve multiple biomedical functions [63, 110, 119, 203], such as improved imaging and drug delivery [201]. Combining all of this functionality

into a single magnetic nanoparticle will reduce the exposure of the patient to multiple treatments.

Magnetic hyperthermia has already been studied in clinical trials for the treatment of prostate cancer [101, 103, 190] and glioblastoma (a brain tumour with a particularly dire prognosis) [133, 134] with positive results. Trials combining general hyperthermia with conventional therapies have also been successful [42, 43, 85, 184], for example in treating breast cancer [117].

3.2 Current challenges

Magnetic hyperthermia is a clinical technique still in its infancy and there are many avenues of research for regulators, clinicians, and biomedical engineers. Problems related to the therapeutic procedure itself are the most difficult to overcome, since these are difficult to replicate *in-vitro*. First of all, the magnetic particles must be administered to the tumour and evenly dispersed throughout the tissue in order to ensure an even distribution of heating [51, 106]. Healthy tissue is at danger of being damaged and cancerous tissue may not reach therapeutic temperatures if the heating is heterogeneous [147]. Reliable monitoring of the temperature distribution is required to assess the quality of the therapy and is still an open question [17]. Currently, one or many temperature probes are used to determine the heating [102, 134] but they suffer from limited spatial resolution. Real-time, non-invasive imaging techniques could offer substantial improvement [169, 204]. Multi-physics simulations, modelling *in-vivo* temperature elevation of tissue with magnetic models of particle heat dissipation, would improve current knowledge of the biophysical interaction [149, 160]. Self-regulated hyperthermia [69, 119, 154] can also prevent overheating by restricting the heating output of particles above the Curie temperature.

The most persistent challenges for magnetic hyperthermia are to minimise the dosage of magnetic nanoparticles and the patient's exposure to alternating magnetic fields. These problems were identified in the earliest studies [74] and continue to be highlighted in recent reviews as prosperous avenues for research [30, 132, 160]. The toxicity of magnetic nanoparticles in the human body have not been fully established [107, 125]; therefore it is preferable to minimise their use. Similarly, safe limits for clinical use of alternating magnetic fields are not well defined [51]. Smaller magnetic fields reduce the chance of undesirable physiological responses and interactions with other magnetic objects during the procedure.

The efficiency of converting the magnetic field generation into a temperature rise in the cancerous tissue is at the heart of these challenges. The heat dissipated by magnetic nanoparticles increases monotonically with the size of the dosage and strength of the applied field. Therefore, the challenge for material scientists is to maximise the heat

dissipated per unit volume of nanoparticles subjected to a clinically acceptable alternating magnetic field. The search for the optimum particle properties for hyperthermia has been a leading area of recent study.

3.3 Heat dissipation in magnetic nanoparticles

A theory of heat dissipation in magnetic systems is required to interpret the relationship between magnetic nanoparticle properties and their resulting heat dissipation. The result of Section 2.4.5 showed that the heat dissipated by an arbitrary magnetic system is entirely defined by its magnetisation response. Consequently, any theoretical model for heat dissipation must rely on the stochastic dynamics of magnetic systems, as described in Section 2.4. The nonlinear nature of magnetic systems greatly complicates the elucidation of the effect of system parameters on magnetic hyperthermia efficiency. The following section summarises current knowledge of the effects of particle and applied field properties on the heat dissipated.

3.3.1 Optimising particle and field properties

The optimum performance of magnetic particle hyperthermia is achieved by maximising the heat dissipated per cycle of the alternating magnetic field (AMF) (or equivalently the dynamic hysteresis area¹) with respect to the properties of the system. This presents a multidimensional, constrained optimisation problem, which has received attention from numerous studies focusing upon the field amplitude and frequency [51]; particle anisotropy, saturating magnetisation, size and shape [116]; as well as the concentration and subsequent interaction and clustering of nanoparticles [46, 174, 180, 189]. The complex interplay between these properties does not allow for each parameter to be optimised in isolation [39]. Moreover, the magnetic material must be biologically compatible [167] and the applied magnetic field must remain within empirically determined safe limits [9].

Early experiments that studied the effect of particle properties on SPL systematically measured the heat dissipated by a range of known materials. In one of the earliest studies [105], it was shown that multi-domain particles had lower SPL than single-domain particles. Indeed, multi-domain particles tend to have a lower coercive field, i.e. are easier to reverse [89], and a higher coercive field correlates with increased SPL [50]. This result

¹Hysteresis loops and their characteristic parameters (coercive field, remanent magnetisation, squareness, saturation magnetisation) can describe either static loops (i.e. loops measured in small increments of field once the system has equilibrated) or dynamic loops (i.e. loops measured from continuously varying fields). In fact, a static loop can be characterised as a dynamic loop with a very slow varying field. Since the dynamic hysteresis loop shape depends on the applied field frequency and field, it is nonsensical to characterise particles without knowledge of the field frequency, shape, and amplitude. Not all studies are clear (see e.g. [50, 80, 158]). Similarly, particles may be characterised as superparamagnetic (Section 2.4.3.1) although a superparamagnetic response also depends on the applied field.

has guided research towards the study of single domain particles [80]. Smaller particles that are superparamagnetic at room temperature also ensure colloidal stability. In other words, the particles will not aggregate in the patient [107].

The experiments in [105] showed that $\text{SPL} \propto H^2$ and $\text{SPL} \propto f$ where H and f are the amplitude and frequency respectively of the applied magnetic field, which agreed with early theoretical work. In his seminal work, Rosensweig [171] used equation (2.48) to predict the heat dissipated by a particle in dynamic equilibrium while subjected to a periodic (sinusoidal) external field. The response of the particle was obtained using the linear response theory (Section 2.4.4), leading to the following formula for the heat dissipated:

$$Q = 2\pi\sigma\mu_0M_s h^2 \frac{\omega\tau(\sigma)}{1 + \omega^2\tau^2(\sigma)} \quad (3.1)$$

where $\omega = 2\pi f$, σ is the *reduced energy barrier* or *stability ratio* (Section 2.4.3.1):

$$\sigma = KV/(k_B T) \quad (3.2)$$

$h = H/H_k$ is the reduced applied field amplitude normalised by the anisotropy field H_k :

$$H_k = 2K/(\mu_0 M_s) \quad (3.3)$$

and $\tau(\sigma)$ is the relaxation time as a function of σ . Recall that $\text{SPL} = fQ/\rho$. The linear response theory result in equation (3.1) agreed with previous experiments that SPL scaled with the square of the applied field [92, 105]. The equation also shows that the maximum energy dissipated per cycle is obtained when $\omega = 1/\tau$, independent of H (assuming that the relaxation time is independent of H). Since the relaxation time of the particle is determined by its geometry, size, and material properties, equation 3.1 was useful for designing optimum particles. Equation 3.1 also shows that the heat scales linearly with the saturation magnetisation M_s of the particle. Therefore, the linear response theory (LRT) equation has become a standard method to both design and interpret *in-vitro/vivo* experimental studies [90, 116, 182, 198, 203].

The relaxation time of the particle τ is determined by the magnetic viscosity (Néel relaxation) and Brownian (i.e. physical rotation of the particle) viscosity. The contributions from Brownian relaxation were studied in early experiments [63] and the relative contribution depended on the frequency of the applied field, particle size and material. The effects of viscosity make it difficult to predict the effect of particles *in-vivo* [92, 110]; though progress is being made [26]. However, several studies of particles in the tumour show that they are mechanically frozen such that Brownian losses are negligible [46, 69, 113], making modelling easier.

The validity of the linear response theory is determined by the linear region of the Langevin function, which is obtained analytically from the equilibrium distribution of a magnetic moment subjected to an external field [27]. The region of applicability is limited

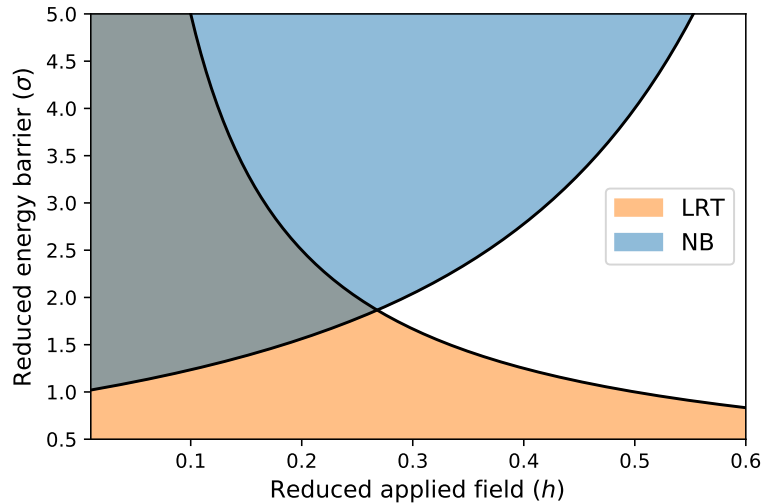


Figure 3.1: The validity of the linear response theory for calculating the SPL of an ensemble of single-domain non-interacting magnetic nanoparticles. The LRT area represents the validity of the linear model for magnetisation response. The NB area represents the validity of the Néel-Brown model for computing the relaxation time of the particle.

to $\mu_0 M_s V H / (k_B T) \ll 1$ such that the field must be small and the particles cannot be too large. Additionally the Néel model for relaxation requires $KV / (k_B T) \ll 1$. Figure 3.1 shows the domain of validity of the LRT in terms of the reduced applied field h/H_K where $H_K = 2K / (\mu_0 M_s)$ and reduced energy barrier $KV / (k_B T)$. The results of LRT should not be extrapolated beyond these limits because it can lead to misinterpreted results, which can be found regularly in the literature [90]. Indeed, the LRT model denies experimental facts: the hysteresis area eventually scales linearly with increasingly large magnetic fields [50] (i.e. geometrically the oblong dynamic hysteresis loop becomes wider) and the optimum particle size is dependent on the applied field. However, there are some generic relationships that can be drawn from [171]: in the limit that the field frequency is very low/high, the particle will exhibit superparamagnetic/frozen behaviour and the hysteresis area will be negligible.

In [27], the thermal activation model for a Stoner-Wohlfarth particle (Section 2.4.3.1) was used to simulate the non-linear magnetisation response. The model is able to simulate particles outside of the LRT limitations but still subject to reasonably large energy barriers for the Néel-Brown approximation (Figure 3.1). The nonlinear model was compared with experiments [139], validating that the coercive field H_c increased with particle size in the single domain region and hysteresis area increases with H_c . In the nonlinear regime, it becomes difficult to extract simple trends between the system parameters and the SPL. For example, the functional form of the relationship between the field strength H and SPL varies substantially with the particle size; in particular between the superparamagnetic and ferromagnetic regimes. A recent analytic work [28] has demonstrated a nonlinear response model capable of extending the availability of analytic solutions.

Unlike the linear response theory, the model replicates many features of nonlinear dynamical simulations.

In reality, particles have heterogeneous properties. The effect of particle size dispersion has received particular attention and has led to conflicting results. On the one hand, increasing dispersity decreases SPL [113, 116] but can also increase SPL [89]. The relationship between dispersion and SPL is further complicated if the distribution incorporates multi-domain particles, which must be accounted for with Rayleigh type losses [89]. Particle synthesis methods have focused on obtaining very narrow size distributions [110] to eliminate the uncertainty. These opposing conclusions were reconciled in [146], which demonstrated that increasing size dispersion will increase the number of larger particles, which will always increase SPL *given that the hysteresis loops are saturated (i.e. major loops)*. The experiments that showed decreasing SPL with increasing dispersity can be explained by their lower field values, leading to minor-loops. The result is that disperse ensembles have a larger proportion of particles with larger coercive fields. Heterogeneity in particle properties also leads to heterogeneity in power loss, which can lead to ineffective hyperthermia. The objective of minimising variation in heating output has only recently been articulated in the works of [146–148]. It was shown in [146] that more dispersed particles contributed to a wider variation in heating.

It is, in general, easier to interpret results when major loops are activated such that $\text{SPL} \propto M_s H_c$ approximately [63], which in non-interacting single particles is similar to $\text{SPL} \propto K$ [136]. Note that the maximum possible hysteresis area for a given H_{\max} is a square shape with saturating magnetisation $A = 4\mu_0 M_s H_c$. If the field is not saturating the particle, then minor loops may come into play and increasing the coercive field has a nonlinear non-monotonic effect on heating.

Conventionally, previous studies have considered the alternating magnetic field as following a sinusoidal shape. Few studies have shown the effect of alternating magnetic field shape on hyperthermia performance. The linear response model [171] was extended to arbitrary waveforms [123] and showed that non-sinusoidal shapes may significantly improve heat dissipation. Similar results were obtained by numerical simulation of the Landau-Lifshitz-Gilbert equation [143, 144], which showed that both square and trapezoidal AMFs may enhance heat dissipation. In common with all of these studies is that the applied field amplitude was sufficiently small such that the magnetisation response could be assumed linear, making their conclusions difficult to generalise to the nonlinear region.

In summary, the effects of particle and field properties on SPL in single particles are highly complex. However, they have been well explored. A typical design would be to increase the coercive field to increase losses, which requires larger applied fields (to ensure reversal) until a clinically acceptable limit is reached. The volume of the particle can then be increased (assuming that it doesn't correlate with the coercive field) to reduce

the effect of thermal fluctuations on reversal. However, increasing volume size increases the contribution of inter-particle interactions [39], which complicates the picture further.

3.3.2 Effect of interactions

Most commonly, existing studies investigating optimum particle and field conditions use the theory of non-interacting particles or experiments without accounting for interaction effects (either they appear in effective parameters or are considered negligible). In reality, particles interact with one another [25], which adjusts the relaxation behaviour and therefore the heating output. In vivo hyperthermia experiments are particularly prone to interaction effects because particles are taken up into lysosomes [189]. However, interactions are hard to control experimentally and complicate theoretical models. Until very recently, the effects of interactions on heat dissipation were controversial [39].

Interacting systems of magnetic nanoparticles may be split into two broad categories:

1. Single particles uniformly distributed through a medium, which interact weakly through long-range dipolar interactions that are approximately seen as a mean-field. The long-range interaction strength depends on the concentration of the particles in the medium.
2. Aggregated clusters of particles, which interact strongly with short-range dipolar interactions, leading to a complex system of interacting particles with non-trivial emergent behaviour.

In the first case, particle concentration is defined as the weight of magnetic particulate matter per weight of suspension medium (e.g. water). The effect of increasing concentration is to increase the contribution of the dipole-dipole interactions (Section 2.1.1.3) in the effective field acting on each particle, since the strength of this interaction is proportional to the cubed distance between particles. By analysing the energy landscape of a large interacting system, it was shown in [34] that in general increasing concentration leads to more disperse energy barriers, however the precise effect on the system depends substantially on the externally applied field and geometry of the system. Two of the earliest studies of interacting particle systems for predicting the effect on heat dissipation, using Metropolis Markov-chain Monte-Carlo (MCMC) [180] and stochastic simulation of the sLLG [84] respectively, concluded that increasing concentration leads to reduced heating output. The negative effect of concentration on SPL was also shown experimentally [175]. A theoretical model, using the mean-field theory, showed that the effect of concentration can be modelled effectively as an additional anisotropy term on each particle [122]. This increase in effective anisotropy matched earlier studies of interactions in nanoparticle solutions not focusing on hyperthermia [29, 81]. The model in [122] showed that magnetically soft particles (i.e. those with low

stability $KV/(k_B T)$) gained from increasing concentration, since the larger energy barriers increased SPL. However, magnetically hard particles (with large stability) could become frozen with energy barriers too large to reverse, decreasing SPL. This effect was confirmed in [146] using kinetic Monte-Carlo (kMC) simulations. Further extensive studies using MCMC [39] and kMC simulations [173, 189] using a wider range of system parameters resolved the conflicting simulated and experimental results. It was shown that increasing the concentration increased the coercivity H_c of the particles and thus the heating output conditional that the external field amplitude H_0 was larger than the coercive field in order to activate major loops. Otherwise if $H_0 < H_c$ the heating output was diminished. Heat dissipation may also be influenced by interactions through varying hysteresis loop shape and [189] and the saturating magnetisation [159].

The dipole-dipole interaction strength between particles dispersed in a solution is relatively weak compared to the interaction effects between particles aggregated into clusters, which leads to a rich landscape of relaxation behaviour for varying cluster shape, size, and interaction strength [95]. In [54], it was shown that increasing the size of particle aggregates in a low concentration solution impacted SPL more than increasing concentration of disaggregated particles. Therefore, an understanding of the heating effects of different nanoparticle aggregations (both size and shape) is needed. The most commonly found aggregates are densely packed clusters and chain-like clusters, which range from just 2 particles to over 200.

Suspended nanoparticles have a tendency to cluster together in an uncontrolled process due to dipolar interactions [20]. Recent breakthroughs in nanoparticle synthesis also allow control over the clustering of particles, which results in stable clusters of a desired size and shape [6, 137]. These synthesis methods could open new avenues for optimising hyperthermia by building structures from smaller building blocks [152]. Early experiments showed that increasingly large clusters can both increase [10] and reduce heat dissipation [20]. MCMC simulations [181] showed a demagnetising effect in clusters of magnetic particles, which form closed magnetic flux loops and reduce their magnetisation. However, it was shown that SPL may actually improve at lower field values. Serantes et.al [181] also compared various aggregation structures and showed similar behaviour for ring, cube, and hexagonal lattice structures but very different behaviour for chains (see below). The demagnetising effects resulting from clustering and the resulting reduced SPL has also been observed experimentally [40, 159]. Interestingly in [40], it was shown that individual particles improve their heat output when influenced by nearby clusters, while the clusters themselves show reduced heat output. A recent combined experimental and computational study using kMC [152] similarly showed that larger clusters of particles (more than four) have lower SPL whereas dimers and trimers (clusters of two and three particles respectively) can show increased heat compared to single particles. In contrast to dispersed particles in a solution, which experience a slight increase in effective anisotropy from nearby particles, interactions in aggregate

structures lead to qualitatively different behaviour and can substantially affect heating rates.

Magnetic nanoparticles are also commonly found in aggregations with chain-like structure. The application of an external magnetic field, leading to alignment of individual particles can cause aggregation into chains but an external field is not required; nanocubes have been observed forming chains without any applied field [136]. Experimental [138] and computational [181] have shown that chains can also enhance or diminish heat output. Chains of nanoparticles have a much larger effective anisotropy than their individual particles [181] and therefore a higher coercivity [136]; heating output is always increased if the alternating field amplitude is large enough to overcome the energy barrier (i.e. major loops). However, in reality there is a danger that chains may become bent and eventually form ring or chain formations [181].

In summary, theoretical and experimental studies have shown that interacting systems of magnetic nanoparticles can have enhanced or diminished heat dissipation compared to non-interacting systems. The enhanced heating rates are achieved by synthesising nanoparticle aggregates with optimum shape and size, which depend on material and field properties. However, the most extensive theoretical studies investigating the effects of shape and size use kMC and MCMC methods to simulate magnetic responses. These Monte-Carlo methods do not yield the correct dynamics, in particular for systems with low damping and small energy barriers [31]. An alternative is to use stochastic simulation of the sLLG equation in order to more accurately simulate the dynamics of interacting aggregates of magnetic nanoparticles. Simulations of the sLLG have already had previous success recovering theoretical [28] and experimental results [196] in magnetic hyperthermia studies.

3.4 Scope of present study

The literature revealed that the heat dissipated by magnetic nanoparticles is influenced by a large number of factors. The open research challenge is to determine the optimum system properties, which maximise the heat dissipated under an alternating magnetic field. The number of possible properties to investigate is too large for a single study and a coherent answer to this challenge will require many years of further research.

In the present study we have chosen to investigate the following two phenomenon:

1. the effect of interactions in dense clusters of single domain magnetic nanoparticles.
2. the effect of non-sinusoidal waveforms on non-interacting single domain magnetic nanoparticles.

We have designed, built and executed computational experiments in place of *in-vitro/vivo* experiments in order gain control over the system properties and investigate their complex effects. The existing literature provided many examples of computational models for heat dissipation: stochastic simulation at the atomistic and macrospin level, kinetic Monte-Carlo, linear systems theory, master equations and phenomenological models. All have been used in an attempt to determine the optimum conditions for effective hyperthermia treatment.

We have used the stochastic Landau-Lifshitz-Gilbert (sLLG) equation to simulate the nonlinear dynamics of systems of interacting single particle macrospins. Although the approach is computationally difficult, it avoids the problems of linear analytic models and of kinetic Monte-Carlo assumptions. In Chapter 4, we discuss methods of simulating the sLLG, which are implemented in an open-source software package described in Chapter 5. We then use the software to simulate interacting ensembles of magnetic nanoparticles in Chapter 6. Finally, we use the master equation to investigate the effects of non-sinusoidal fields in Chapter 7.

Chapter 4

Simulating the Landau-Lifshitz-Gilbert equation

The sLLG equation is a stochastic, multidimensional, nonlinear differential equation, which is hard to solve analytically, except in simplified cases (Section 4.4.3.1). In systems of many particles, the interaction between the magnetic moments complicates the equations further leading to a set of coupled sLLG equations. Therefore, a numerical solver is required to simulate the stochastic Landau-Lifshitz-Gilbert (sLLG). However, the sLLG is a particularly difficult equation to solve, even numerically, due to its mathematical properties (Section 4.1). Numerical methods dedicated to solving the sLLG are an active area of research (see for example [2, 126]). The most sophisticated methods aim to satisfy desirable mathematical properties such as conserving the length of the magnetisation or the total energy of the system but these properties come at the expense of increasing algorithmic complexity. Practical considerations such as: implementing, tuning, and measuring the performance of numerical methods for the sLLG has received less interest. Perhaps due to the popularity and success of zero-temperature micromagnetics [60, 61], there are reviews of numerical methods for the deterministic LLG [13, 36]. In contrast, there are few examples to help choose the appropriate technique for simulating the stochastic LLG.

One recent study [3] compared multiple methods to solve the stochastic Landau-Lifshitz-Gilbert-Slonczewski equation. The paper provides guidance to researchers by simulating relevant test problems and demonstrating the advantages and disadvantages of each approach. Metrics such as the equilibrium distribution accuracy and conservation of the norm of the magnetisation vector \mathbf{m} (Section 2.1.3) are used to measure performance. However, the paper does not quantify the time-dependent errors or strong convergence of the methods, such as in [56]. A detailed study of the performance of numerical methods is required to select the most appropriate algorithm for simulating the sLLG equation; in particular the effects of tuning parameters, noise intensity, and projection methods

on solution accuracy. These results would improve understanding of software packages that already implement these methods as well as to the present work.

In this chapter we first review methods for solving the sLLG¹. We implement and measure the relative performance of the most commonly used methods: the Euler-Maruyama, Heun, and fully implicit scheme, due to their simplicity. We show that conserving the norm of the magnetisation does not necessarily improve the path-wise accuracy of the solution and that the temperature of the system greatly affects the accuracy of the Heun scheme. Finally, we conclude that for moderately-sized interacting systems the implicit midpoint scheme is prohibitively expensive.

Note that we do not consider the case where the magnetisation varies spatially, leading to a stochastic partial differential equation representing the spatio-temporal evolution of the magnetisation. Rather, we consider the temporal evolution of distinct magnetisation vectors interacting through the effective field term. For examples of solving the sLLG for a continuous magnetisation domain (in other words continuous micromagnetics [60, 61] with thermal fluctuations) see [11, 12, 35, 100].

4.1 Equation properties

The sLLG equation is an example of a stiff stochastic differential equation. Stiff systems are loosely defined as possessing dynamics with a large difference between their slowest and fastest timescales [24]. The difficulty with stiff systems is that a numerical method will require very small time steps in order to resolve the fast dynamics but consequently will have to take a very large number of small steps to resolve the slower dynamics. Any numerical method that requires a very large number of steps will accumulate error and in some cases become unstable. In stochastic differential equations, stiffness may appear in the drift and/or diffusion terms. The sLLG possesses a number of timescales relating to the random thermal noise, precessional motion, and relaxational motion, which may have dispersed timescales. Implicit rather than explicit numerical methods are required to handle stiff differential equations effectively (see Section 2.3.6.4).

The structure of the diffusion term $B(\mathbf{x}, t)$ also influences the best choice of numerical method. Systems with *commutative* noise structures (Section 2.3.1) may be solved more efficiently and the numerical solvers may be implemented more simply. Substituting the sLLG diffusion matrix (equation (2.34)) into equation (2.17) shows that the sLLG does not satisfy the commutative condition and therefore the numerical solvers must be implemented in their more general form, which poses a severe restriction on the performance of many methods (e.g. the Milstein scheme).

¹During the preparation of this chapter, a complementary work was published in [55]. This excellent review covers many of the solvers listed below and reinforces our conclusions.

The sLLG equation has the mathematical property that the magnetic moment magnitude is constant $\|\mathbf{m}\|_2 = 1$, as discussed previously. However, numerical approximations of solutions to the sLLG may exhibit varying magnetic moment length due to numerical errors. Methods to ensure the magnetic moment magnitude is conserved are discussed in Section 4.2.

Finally, the parameters of the sLLG (equation (2.32)) have values on vastly different orders of magnitude. For example the particle radius is $\sim 10^{-9}\text{m}$ while the magnetisation is $\sim 10^6\text{Am}^{-1}$. Performing floating point arithmetic on orders of this magnitude can introduce truncation errors, which accumulate over time. Therefore, the sLLG is normalised by introducing *reduced time* ℓ , which simplifies the equations and scales the parameters to around an order of unity.

4.1.1 Reduced equations

The reduced form of the sLLG, effective and thermal field equations are presented here. The details of the normalisation can be found in Appendix B. The reduced sLLG, written in integral form, for a particle i in an interacting system of N particles is:

$$d\mathbf{m}_i = \underline{\mathbf{a}}_i(\mathbf{m}, t)d\ell + B_i(\mathbf{m}_i, t) \circ d\mathbf{W}_i(\ell) \quad (4.1)$$

where $\mathbf{m}_i \in \mathbb{R}^3$ is the magnetic moment of the i^{th} particle, $\mathbf{m} = [\mathbf{m}_1; \mathbf{m}_2; \dots; \mathbf{m}_N] \in \mathbb{R}^{3N}$ is the vector of all magnetic moments, and $\circ d\mathbf{W}_i(\ell) \in \mathbb{R}^3$ is the Wiener process acting on particle i to be interpreted in the Stratonovich sense. The drift and diffusion components for particle i :

$$\underline{\mathbf{a}}_i = -\mathbf{m}_i \times \mathbf{h}_{\text{eff},i}(\mathbf{m}) - \alpha \mathbf{m}_i \times \mathbf{m}_i \times \mathbf{h}_{\text{eff},i}(\mathbf{m}) \quad (4.2)$$

$$B_i = -\left(\sqrt{2D_i}\mathbf{m}_i \times\right) - \left(\alpha\sqrt{2D_i}\mathbf{m}_i \times \mathbf{m}_i \times\right) \quad (4.3)$$

Note that the drift $\underline{\mathbf{a}}_i$ for a single particle i depends on the magnetisation of the entire system \mathbf{m} because of the dipole-dipole interaction term (equation (2.4)) in the effective field $\mathbf{h}_{\text{eff},i}$. The diffusion matrix B_i for particle i is only a function of the particle magnetisation \mathbf{m}_i . The reduced time is defined:

$$\ell \triangleq t \frac{\gamma H_k}{1 + \alpha^2} = t \frac{2\gamma \bar{K}}{M_s (1 + \alpha^2)} \quad (4.4)$$

and the anisotropy field is introduced:

$$H_k \triangleq \frac{2\bar{K}}{\mu_0 M_s} \quad (4.5)$$

which normalises the effective field $\mathbf{h}_{\text{eff},i} = \mathbf{H}_{\text{eff},i}/H_k$. $\bar{K} \triangleq 1/N \sum_i K_i$ is the average anisotropy constant in the system. The reduced effective field term is obtained from

multiplying equation (2.11) by μ_0/H_k :

$$\mathbf{h}_{\text{eff},i} = \mathbf{k}_i (\mathbf{m}_i \cdot \mathbf{k}_i) + \mathbf{h} + \frac{\mu_0 M_s^2}{8\pi \bar{K}} \sum_{j \neq i} \frac{v_j}{\|r_{ij}\|^3} (3(\mathbf{m}_j \cdot \mathbf{r}_{ij}) \mathbf{r}_{ij} - \mathbf{m}_j) \quad (4.6)$$

where $\bar{V} = 1/N \sum_i V_i$ is the average particle volume. The reduced thermal field strength is:

$$D_i = \frac{\alpha k_B T}{2V_i \bar{K} (1 + \alpha^2)} \quad (4.7)$$

The Itô drift correction for the sLLG is obtained from substituting the components of equation (4.1) into equation (2.16) [18, 87]:

$$\mathbf{a}_i(\mathbf{m}, t) = \underline{\mathbf{a}}_i(\mathbf{m}, t) - 2D_i \mathbf{m}_i \quad (4.8)$$

The first term describes preferential alignment of \mathbf{m} with the particle's anisotropy axis, which acts in the unit direction $\mathbf{k}_i \in \mathbb{R}^3$ with magnitude k_i (units Jm^{-3}). The second term describes the effect of an externally applied field $\mathbf{h}_{\text{app}} \in \mathbb{R}^3$. The final term describes the field experienced by particle i through a long-range dipolar-dipolar interaction with a nearby particle j , where V_j is the volume (units m^3) of particle j and $v_j = V_j/\bar{V}$ is the reduced volume; $\bar{V} = 1/N \sum_{n=0}^N V_n$ is the mean volume of all particles in the system; $\bar{K} = 1/N \sum_{n=0}^N k_n$ is the mean anisotropy magnitude; $\mu_0 = 4\pi \times 10^{-7}$ is a constant (units $\text{mkg s}^{-2} \text{A}^{-2}$); M_s is the magnitude of the macrospin (the saturation magnetisation, units Am^{-1}) and $\mathbf{r}_{ij} \in \mathbb{R}^3$ and r_{ij} are the unit vector and magnitude respectively of the reduced distance between particles i and j . The reduced distance is the true distance (units m) divided by $\sqrt[3]{\bar{V}}$ and appears in equation (2.11) because the numerator and denominator of the interaction term are divided by \bar{V} , which has the effect of scaling both values close to unity. The prefactor $\mu_0 M_s^2 / (2\bar{K})$ can be computed in advance and will also evaluate close to unity.

4.2 Review of numerical methods

The most popular numerical method for simulating the sLLG is the Heun scheme, which is a generalisation of the deterministic Heun scheme. The scheme appears in a large proportion of influential studies and software packages [56, 57, 84, 112, 129, 153, 178, 195]. The success of the Heun scheme may be attributed to its relative simplicity, ease of implementation, and consistency with the Stratonovich solution. A rigorous study of the sLLG [66] emphasises the importance of choosing a solver consistent with the Stratonovich and not the Itô solution. The scheme has a strong convergence of order 0.5 in general but increases to 1.0 for SDEs with commutative noise [172]. This detail can

cause confusion because the stochastic LLG at zero-temperature or with a scalar noise term (see Section 4.4.3.1 for results of this simplified sLLG) are commutative².

The Heun scheme does not conserve the length of the magnetisation and as an explicit scheme, it is less robust to stiffness. A number of approaches have also been taken to ensure the conservation of the magnetisation vector magnitude. The first, transforms the sLLG into a 2-dimensional system in spherical coordinates [170]. The resulting equation requires many evaluations of trigonometric functions at each step and also has numerical instabilities around angles of $0, \pi$ rads [4, 18]. Many implementations of the Heun scheme include a renormalisation (or projection) step [56, 127]. The renormalisation simply rescales the length of the magnetisation $\mathbf{m}_{\text{scaled}} = \mathbf{m}/\|\mathbf{m}\|$ after every update to the state. Finally, some authors introduce an additional term in the sLLG that prevents the magnetisation vector from drifting [41]. Conserving the length of the magnetisation is necessary for an accurate numerical solution but not sufficient. It does not necessarily follow that these methods will lead to lower path-wise error. The renormalisation step has been criticised in [53] as a “nonlinear modification of the numerical scheme” and thus invalidating any theoretical guarantees on convergence and stability of the algorithm.

The simplest scheme for solving SDEs is the Euler scheme, which is consistent with the Itô solution [115]. Few studies implement the Itô scheme over the Heun scheme, citing worse stability or the fact that it does not converge to the Stratonovich solution [66]. However, it has been shown in [18] that the Itô induced drift (equation (2.16)) acts parallel to the magnetic moment, thus only affecting the length of \mathbf{m} . An important consequence of this is that if a renormalisation step is used, forcing the length of the magnetisation to be constant, the Itô induced drift cannot take effect. Consequently, when using a renormalisation scheme, the Itô and Stratonovich interpretations give the same solution. The ongoing debate between authors insisting on Stratonovich solvers and those claiming the two approaches are the same is resolved simply by including the Itô induced drift (equation (4.8)) in the sLLG equation when using Itô solvers or discarding it when using Stratonovich solvers.

An implicit midpoint scheme (IMP) was proposed in [52, 142], which is consistent with the Stratonovich solution and conserves the magnetisation length (to within a desired tolerance). Moreover, the method conserves the energy of the system such that an undamped magnetic moment ($\alpha = 0$ in the LLG) will continue to precess [140]. The IMP scheme requires many more evaluations of the sLLG and for a system of N interacting particles requires many evaluations of the effective field (with complexity $\mathcal{O}(N^2)$ equation (2.4)). However, it was shown that in the deterministic case the computational cost of the IMP scheme may be reasonable given that it maintains accuracy and stability at larger time steps [53]. Initial results have been published for the stochastic case [52] and more lately the IMP was shown to outperform the Heun scheme [4]. However, it has been noted qualitatively [140] that solving a $3N$ set of nonlinear equations is very

²See e.g. [87], where the Heun convergence rate is empirically shown as 1.0 in the simplified sLLG

computationally expensive. Moreover, the IMP scheme contains additional tuning parameters (the tolerance of the simultaneous equation solver) and there is no published guidance on their selection or a clear guide on implementation.

The Euler, Heun, and IMP schemes (without renormalisation) have proven convergence and stability properties [115, 142] for an arbitrary SDE. Solvers have also been designed specifically for the sLLG, which have not been treated with the same mathematical rigour but perform well in numerical tests. In [140] two schemes, named Semi-Implicit Scheme A (SIA) and Semi-Implicit Scheme B (SIB), were designed using implicitness in a subset of the sLLG terms in order to ensure conservation of the magnetisation length. These methods avoid implicitness in the effective field calculation, which reduces the implicit problem to a set of N independent 3-dimensional linear systems rather than a single $3N$ -dimensional nonlinear system. Therefore the SIA and SIB methods require far fewer evaluations of the expensive effective field calculation.

More recent work has investigated the use of adaptive time steps [126], which for deterministic ODEs are more efficient and guarantee the error remains with a specified tolerance. Adaptive time steps are difficult to implement for SDEs due to the complexity of ensuring that the Wiener process is statistically correct [99]. The use of higher order (e.g. the Milstein scheme) may also be interesting but solving non-commutative equations with order higher than 0.5 requires the simulation of multiple stochastic integrals, which are very expensive to compute [115, 172, 176]. They also often require many more evaluations of the SDE, which for the sLLG leads to more evaluations of the effective field terms.

Comparing the relative performance of the different numerical schemes is difficult for the same reason that such schemes are required: there is no existing solution to the full sLLG. In place, the equilibrium conditions [44], conservation of energy and magnetic moment length [140], and comparison with experimental results [195] are often used to validate the methods. The most simple performance metric is to measure the error between the true trajectories and the simulated trajectories, which if true is sufficient.

In the subsequent sections, the Euler, Heun, and IMP schemes are implemented and tested. The solutions are compared with analytic solutions to the deterministic LLG and a simplified form of the stochastic LLG. In the case of the full sLLG, we use the local relative error to prove the global convergence order. The equilibrium distribution of an interacting system is compared with the results of an MCMC simulation. Finally, the time taken for each scheme is measured to determine which method offers the best accuracy for the same computational effort. The numerical schemes are summarised in table 4.1.

Table 4.1: Global mean squared error convergence for popular stochastic numerical schemes.

Method	Implicit	Stratonovich	Strong order	Derivatives
Euler-Maruyama	No	No	0.5	None
Heun	No	Yes	0.5*	None
Milstein	No	Yes	1.0	first-order
Implicit midpoint	Yes	Yes	0.5*	first-order#

* 1.0 in the commutative.

approximate or higher order derivatives may be required depending on implementation.

4.3 Technical definition and implementation

For a system of N interacting magnetic moments, the estimated state of the system at time t_n is $\mathbf{m}_n \in \mathbb{R}^n$ where $n = 3N$ for the x, y, z component of magnetisation for each of the N moments. The drift vector is therefore:

$$\underline{\mathbf{a}}(\mathbf{m}_n, t_n) = \begin{bmatrix} \underline{\mathbf{a}}^0(\mathbf{m}_n, t_n) \\ \underline{\mathbf{a}}^1(\mathbf{m}_n, t_n) \\ \vdots \\ \underline{\mathbf{a}}^N(\mathbf{m}_n, t_n) \end{bmatrix} \quad (4.9)$$

where $\underline{\mathbf{a}}^i$ is the drift vector for the i^{th} moment. The diffusion matrix for the system is:

$$B(\mathbf{m}_n, t_n) = \begin{bmatrix} B^1(\mathbf{m}_n^1, t_n) & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & B^2(\mathbf{m}_n^2, t_n) & & \vdots \\ \vdots & & \ddots & \mathbf{0} \\ \mathbf{0} & \dots & \mathbf{0} & B^N(\mathbf{m}_n^N, t_n) \end{bmatrix} \quad (4.10)$$

where \mathbf{m}_n^i is the magnetic moment of the i^{th} particle at time t_n . The diffusion matrix is block diagonal consisting of 3×3 blocks of B^i each of which is the diffusion matrix for particle i (equation (4.2)).

4.3.1 Euler-Maruyama scheme

Taking an Itô-Taylor expansion around the solution to an SDE and discarding all terms except the first, leads to the Euler-Maruyama scheme [115]:

$$\mathbf{m}_{n+1} = \mathbf{m}_n + \mathbf{a}(\mathbf{m}_n, t_n)\Delta t + B(\mathbf{m}, t_n)\Delta \mathbf{W}_n \quad (4.11)$$

The Euler-Maruyama scheme is an explicit method, consistent with the Itô solution, strongly convergent with order 0.5, and weakly convergent with order 1.0. As seen

in Section 2.3.6.4 the scheme is not reliable for solving stiff SDEs. However, the Euler-Maruyama scheme is by far the least computationally expensive and therefore very small time steps may be achieved with fewer resources.

4.3.2 Stochastic Heun Scheme

The Heun scheme is a Runge-Kutta like predictor-corrector method and phenomenological generalisation of the deterministic Heun scheme:

$$\begin{aligned} \mathbf{m}_{n+1} = \mathbf{m}_n + \frac{1}{2} \{ \underline{\mathbf{a}}(\mathbf{m}_n, t_n) + \underline{\mathbf{a}}(\hat{\mathbf{m}}_n, t_n) \} \Delta t \\ + \frac{1}{2} \{ B(\mathbf{m}_n, t_n) + B(\hat{\mathbf{m}}_n, t_n) \} \Delta \mathbf{W}_n \end{aligned} \quad (4.12)$$

where

$$\hat{\mathbf{m}}_n = \mathbf{m}_n + \underline{\mathbf{a}}(\mathbf{m}_n, t_n) \Delta + B(\mathbf{m}_n, t_n) \Delta \mathbf{W}_n \quad (4.13)$$

This scheme is not consistent with Itô calculus unless $B(\cdot) = \text{const.}$ (see [114] p.326) but converges to the Stratonovich solution, which can be shown by example, see [67]. The convergence of the Heun and a family of Runge-Kutta solvers have been discussed in detail in [172]. The Heun scheme has a convergence of 0.5 unless the SDE satisfies the commutativity condition, in which case it achieves an improved order of 1.0 [115]. The algorithm cannot be simplified for commutative equations. The Heun scheme is an explicit method.

4.3.3 Fully implicit midpoint

A fully implicit midpoint method was proposed in [142], which uses a truncated Wiener process to handle the unboundedness of the one-step predictor (Section 2.3.6.4). Let the scalar Wiener process increments be denoted: $\Delta W_n = W_{n+1} - W_n = \xi_n \sqrt{\Delta t}$ where $\xi_n \sim \mathcal{N}\{0, 1\}$. The truncated scalar Wiener process has increments $\Delta \tilde{W}_n = \tilde{W}_{n+1} - \tilde{W}_n = \zeta_n \sqrt{\Delta t}$ where:

$$\zeta_n = \begin{cases} A_{\Delta t} & \xi_n > A_{\Delta t} \\ -A_{\Delta t} & \xi_n < -A_{\Delta t} \\ \xi_n & \text{otherwise} \end{cases} \quad (4.14)$$

and $A_{\Delta t} = \sqrt{2k|\ln \Delta t|}$ for any integer $k > 0$. The case is analogous for a vector of independent Wiener processes \mathbf{W}_n to obtain the truncated processes $\tilde{\mathbf{W}}_n$. The fully implicit midpoint method is then written [142]:

$$\mathbf{m}_{n+1} = \mathbf{m}_n + \underline{\mathbf{a}}\left(\frac{\mathbf{m}_n + \mathbf{m}_{n+1}}{2}, t_n + \frac{\Delta t}{2}\right) \Delta t + B\left(\frac{\mathbf{m}_n + \mathbf{m}_{n+1}}{2}, t_n\right) \Delta \tilde{\mathbf{W}}_n \quad (4.15)$$

Milstein proved that this scheme is strongly convergent with order 0.5 and 1.0 for commutative equations. If the diffusion term is zero, the scheme reduces to the deterministic midpoint method, which has an order of convergence of 2.0. The next state of the system \mathbf{m}_{n+1} appears on both sides of the equation (4.15). Therefore, a nonlinear n dimensional system of equations must be solved at every time step.

In order to solve the system of equations, introduce the following variable $\hat{\mathbf{m}} = (\mathbf{m}_n + \mathbf{m}_{n+1})/2$ such that:

$$2\hat{\mathbf{m}} - \mathbf{m}_n = \mathbf{m}_n + \underline{\mathbf{a}}\left(\hat{\mathbf{m}}, t_n + \frac{\Delta t}{2}\right)\Delta t + B(\hat{\mathbf{m}}, t_n)\Delta\tilde{\mathbf{W}}_n \quad (4.16)$$

Rearranging we define the function $\mathbf{f}(\hat{\mathbf{m}})$:

$$\mathbf{f}(\hat{\mathbf{m}}) = \hat{\mathbf{m}} - \mathbf{m}_n - \frac{1}{2}\underline{\mathbf{a}}\left(\hat{\mathbf{m}}, t_n + \frac{\Delta t}{2}\right)\Delta t - \frac{1}{2}B(\hat{\mathbf{m}}, t_n)\Delta\tilde{\mathbf{W}}_n = 0 \quad (4.17)$$

where we must solve $\mathbf{f}(\hat{\mathbf{m}}) = 0$ in order to obtain $\hat{\mathbf{m}}$. Expanding the drift and diffusion terms with the effective field terms leads to a very complex \mathbf{f} and therefore it must be solved numerically. The Newton-Raphson method [164] is an iterative method to find the roots of \mathbf{f} :

$$\hat{\mathbf{m}}_{\text{next}} = \hat{\mathbf{m}} - J^{-1}(\hat{\mathbf{m}})\mathbf{f}(\hat{\mathbf{m}}) \quad (4.18)$$

where $J = \frac{\partial \mathbf{f}}{\partial \hat{\mathbf{m}}}$ is the Jacobian of the function \mathbf{f} . The iteration is repeatedly applied to the approximate solution of $\hat{\mathbf{m}}$ until the relative error converges such that $\|\hat{\mathbf{m}}_{\text{next}} - \hat{\mathbf{m}}\| < \epsilon\|\hat{\mathbf{m}}\|$ for some tolerance ϵ .

The Newton-Raphson iteration in equation (4.18) requires the inverse of the Jacobian, which as the systems grows becomes very computationally expensive. Rather, the problem is reformulated into the form $A\mathbf{x} = \mathbf{b}$, which represents an n dimensional linear system of algebraic equations. Efficient solutions to this problem are very well studied [164]. The Newton-Raphson iterator was reformulated as the following set of simultaneous equations:

$$J(\hat{\mathbf{m}})[\hat{\mathbf{m}}_{\text{next}} - \hat{\mathbf{m}}] = -\mathbf{f}(\hat{\mathbf{m}}) \quad (4.19)$$

which was solved using the LAPACK routine GAS'S [5].

In the case of the sLLG The Jacobian for $\mathbf{f}(\hat{\mathbf{m}})$ (equation (4.17)) is:

$$J(\hat{\mathbf{m}}) = \frac{\partial \mathbf{f}(\hat{\mathbf{m}})}{\partial \hat{\mathbf{m}}} = I_{n,n} - \frac{1}{2} \frac{\partial \underline{\mathbf{a}}(\hat{\mathbf{m}}, t_n + \Delta t/2)}{\partial \hat{\mathbf{m}}} \Delta t - \frac{1}{2} \frac{\partial B(\hat{\mathbf{m}}, t_n)}{\partial \hat{\mathbf{m}}} \Delta \mathbf{W}_n \quad (4.20)$$

where $I_{n,n}$ is the $n \times n$ identity matrix. The diffusion matrix for a particle i depends only on the magnetic moment of particle i and not on any other particle in the system. Therefore the partial derivative of the diffusion matrix for a particle i with respect to the state of the system (i.e. all magnetic moment components) will be zero everywhere except for the corresponding three components for the particle magnetisation $\frac{\partial B^i}{\partial \hat{\mathbf{m}}} = \frac{\partial B^i}{\partial \hat{\mathbf{m}}^i}$. Consequently the third term in equation (4.20) is a block diagonal matrix. A graphic

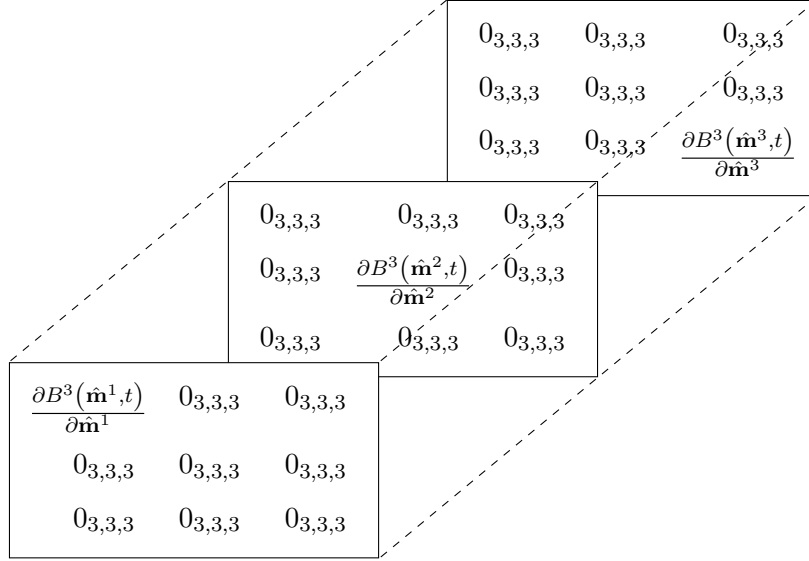


Figure 4.1: Structure of the diffusion Jacobian for a system of 3 magnetic moments. $0_{3,3,3}$ represents the 3-dimensional cube of zeroes of size $3 \times 3 \times 3$. $B^n(\hat{\mathbf{m}}^n, t)$ denotes the diffusion of the n^{th} magnetic moment in the system.

representation of the system block diagonal diffusion Jacobian is shown in Figure 4.1 for a system of three particles.

The drift coefficient for each particle in the ensemble contains the effective field term. If the magnetic moments are interacting with one another, then the effective field on a single particle i is a function of the magnetic moments of all particles in the system. Therefore, the Jacobian of the system drift vector is a dense $n \times n$ matrix:

$$\frac{\partial \underline{\mathbf{a}}(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}} = \begin{bmatrix} \frac{\partial \underline{\mathbf{a}}^1(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}^1} & \frac{\partial \underline{\mathbf{a}}^1(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}^2} & \cdots & \frac{\partial \underline{\mathbf{a}}^1(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}^N} \\ \frac{\partial \underline{\mathbf{a}}^2(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}^1} & \frac{\partial \underline{\mathbf{a}}^2(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}^2} & & \vdots \\ \vdots & & \ddots & \\ \frac{\partial \underline{\mathbf{a}}^N(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}^1} & \frac{\partial \underline{\mathbf{a}}^N(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}^2} & \cdots & \frac{\partial \underline{\mathbf{a}}^N(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}^N} \end{bmatrix} \quad (4.21)$$

where each element of the matrix is a dense 3×3 matrix. However, if the dipole-dipole interaction field is ignored and the effective field term consists only of the anisotropy and Zeeman terms then the effective field vector for particle i is a function of $\hat{\mathbf{m}}^i$ only and the Jacobian reduces:

$$\frac{\partial \tilde{\underline{\mathbf{a}}}(\hat{\mathbf{m}}, t)}{\partial \hat{\mathbf{m}}} = \begin{bmatrix} \frac{\partial \tilde{\underline{\mathbf{a}}}^1(\hat{\mathbf{m}}^1, t)}{\partial \hat{\mathbf{m}}^1} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \frac{\partial \tilde{\underline{\mathbf{a}}}^2(\hat{\mathbf{m}}^2, t)}{\partial \hat{\mathbf{m}}^2} & & \vdots \\ \vdots & & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \frac{\partial \tilde{\underline{\mathbf{a}}}^N(\hat{\mathbf{m}}^N, t)}{\partial \hat{\mathbf{m}}^N} \end{bmatrix} \quad (4.22)$$

where $\tilde{\underline{\mathbf{a}}}$ is the approximate drift term ignoring interactions, which results in a greatly simplified expression.

We define an approximation of the Jacobian $\tilde{J}(\mathbf{m})$, which is identical to the sLLG Jacobian in equation (4.20) but ignores the dipole-dipole interaction term. From the previous discussion, the approximate Jacobian is much easier to compute than the full Jacobian since all of the individual magnetic moments are decoupled. $\tilde{J}(\hat{\mathbf{m}})$ is an $n \times n$ block diagonal matrix with 3×3 :

$$\tilde{J}(\hat{\mathbf{m}}) = \begin{bmatrix} \tilde{J}^1(\hat{\mathbf{m}}) & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \tilde{J}^2(\hat{\mathbf{m}}) & & \vdots \\ \vdots & & \ddots & \mathbf{0} \\ \mathbf{0} & \dots & \mathbf{0} & \tilde{J}^N(\hat{\mathbf{m}}) \end{bmatrix} \quad (4.23)$$

where

$$\tilde{J}^i(\hat{\mathbf{m}}) = I_{3,3} - \frac{1}{2} \frac{\partial \tilde{\mathbf{a}}^i(\hat{\mathbf{m}}^i, t_n + \Delta t/2)}{\partial \hat{\mathbf{m}}^i} \Delta t - \frac{1}{2} \frac{\partial B^i(\hat{\mathbf{m}}^i, t_n)}{\partial \hat{\mathbf{m}}^i} \Delta \mathbf{W}_n^i \quad (4.24)$$

The approximate Jacobian was used in the Newton-Raphson method to solve the implicit formula at each time step. Using an approximate Jacobian to solve the Newton-Raphson problem is referred to as a *Quasi-Newton* method and is often sufficient to find the correct solution.

The Newton-Raphson method is quite robust to approximate Jacobian terms but is very sensitive to the initial guess. The initial guess was taken to be a single Euler step. Therefore to determine the next state \mathbf{y}_{n+1} the fully implicit midpoint method followed the procedure:

1. Begin with the full state of the system at t_n : \mathbf{m}_n
2. Compute an initial guess for \mathbf{m}_{n+1} using the explicit Euler scheme: $\mathbf{m}_{\text{initial}}$
3. Transform the initial guess into an initial guess for the variable $\hat{\mathbf{m}}$ by: $\hat{\mathbf{m}}_{\text{initial}} = (\mathbf{m}_n + \mathbf{m}_{\text{initial}})/2$
4. Formulate the function $F(\hat{\mathbf{m}}) = 0$ and the Jacobian $J(\hat{\mathbf{m}}) = \partial F / \partial \hat{\mathbf{m}}$
5. Set $\hat{\mathbf{m}}_{\text{prev}} = \hat{\mathbf{m}}_{\text{initial}}$
6. Iterate the Newton-Raphson method until convergence:
 - (a) Solve the set of linear equations $J(\hat{\mathbf{m}}_{\text{prev}})[\hat{\mathbf{m}}_{\text{next}} - \hat{\mathbf{m}}_{\text{prev}}] = -F(\hat{\mathbf{m}}_{\text{prev}})$ for $\hat{\mathbf{m}}_{\text{next}}$
 - (b) Check the convergence condition $\|\hat{\mathbf{m}}_{\text{next}} - \hat{\mathbf{m}}_{\text{prev}}\| < \epsilon$
 - (c) If the solution has not converged, set $\hat{\mathbf{m}}_{\text{prev}} = \hat{\mathbf{m}}_{\text{next}}$ and compute another $\hat{\mathbf{m}}_{\text{next}}$ (6a). If it has converged, set $\hat{\mathbf{m}}_{n+1} = \hat{\mathbf{m}}_{\text{next}}$
7. Transform the result into the final estimate for the state \vec{m}_{n+1} at time t_{n+1} by: $\mathbf{m}_{n+1} = 2\hat{\mathbf{m}}_{n+1} - \mathbf{m}_n$

Note that the Newton-Raphson iteration (equation (4.19)) is a $3N$ -dimensional system but since the quasi-Jacobian is block diagonal, the problem can be reduced to solving N 3-dimensional systems, which would reduce memory and computational requirements. Also note that although the quasi-Jacobian ignores the interaction term, the right hand side of the Newton-Raphson iterator $\mathbf{f}(\hat{\mathbf{m}})$ (equation (4.17)) still requires the effective field to be updated at every iteration (i.e. multiple times per step of the integrator $\mathbf{m}_n \rightarrow \mathbf{m}_{n+1}$) an order $\mathcal{O}(N^2)$ problem (equation (4.6)), which is very expensive for large systems. The origin of the complexity can be traced back to equation (4.15), the update step for the implicit integrator, which has implicitness (i.e. \mathbf{m}_{n+1} on the right hand side) in the drift term and therefore in the effective field term.

4.4 Results and validation

4.4.1 Linear test equation

The Euler-Maruyama, Heun and fully implicit schemes were implemented in C++ (see Chapter 5) and applied to the linear test equation (2.24) to compare their stability. The analysis of MS-stability in Section 2.3.6.4 showed that the explicit Euler method is unstable outside of a small region of values for $p = \mu\Delta t$ and $q = \sigma\sqrt{\Delta t}$ while the implicit method is stable for all values given the equation is also MS-stable. The Heun and implicit methods simulated the Stratonovich test equation:

$$dx(t) = \left(\mu - \frac{1}{2}\sigma^2 \right) x(t)dt + \sigma x(t) \circ dW_t \quad (4.25)$$

Following the method in Section 2.3.6.4, the one-step predictor for the Heun scheme was derived for the test equation. The results showed that the Heun scheme has very similar stability properties to the explicit Euler scheme: p should be around -1 and the noise q must be low. The stability of the schemes was visualised by comparing their approximate solutions to the true solution of the test SDE. The analytic solution to the test equation is:

$$x(t) = x_0 e^{(\mu - \frac{1}{2}\sigma^2)t + \sigma W_t}, \quad x_0 = x(0) \quad (4.26)$$

Figure 4.2 shows the true trajectory of the test equation with a given Wiener process, W_t and $\mu = -1, \sigma = 1, \Delta t = 0.01$. The Wiener process W_t was generated by drawing independent and identically distributed increments ΔW_t from a Normal distribution $\Delta W_t \sim \mathcal{N}\{0, \delta t\}$ such that $W_T = \sum_{t < T} \Delta W_t$. The Random values were generated using the Mersenne Twister pseudo-random number generator [70]. Each scheme simulated the test equation from the same initial condition with a reasonably large time step. The Heun and implicit scheme trajectories track the underlying solution closely but the Euler shows slightly worse performance. Stiffness was introduced into equation (2.24) by increasing the noise with coefficients $\mu = -1, \sigma = 4$, the resulting trajectories are shown

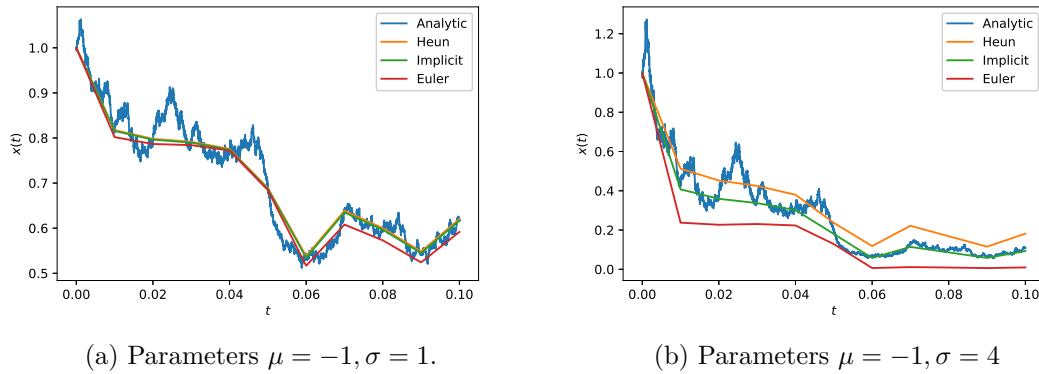


Figure 4.2: Comparison of Heun and fully implicit methods for approximating the linear test equation with different coefficients μ, σ . Setting σ higher introduces stiffness in the diffusion term.

in Figure 4.2. The results show that the explicit schemes quickly diverge from the path. The implicit solution tracks the trajectory closely and even though it diverges slightly the paths recover in equilibrium. A disadvantage of explicit schemes is that their error is not corrected over time.

4.4.2 Deterministic Landau-Lifshitz-Gilbert equation

The deterministic LLG equation is obtained by setting the temperature $T = 0\text{K}$. For a constant applied field in the z -direction and the magnetic moment initially applied along its x -component, the analytic solution to the reduced LLG equation is obtained [87]:

$$\mathbf{m}(t) \begin{bmatrix} m_x(t) \\ m_y(t) \\ m_z(t) \end{bmatrix} = \begin{bmatrix} \text{sech}(\alpha ht) \cos(ht) \\ \text{sech}(\alpha ht) \sin(ht) \\ \tan(\alpha ht) \end{bmatrix} \quad (4.27)$$

An example trajectory of the LLG equation is shown in Figure 4.3 with parameters $h = 0.01, \alpha = 0.1$. The solution shows a precession of the magnetic moment about the field in the z -direction. The oscillations damp over time such that the moment is stationary and aligned with the external field.

The Heun and fully implicit numerical schemes were applied to the LLG equation and compared with the analytic solution. Figure 4.4 shows the path-wise error for each component of the magnetic moment. Overall, both schemes show larger errors during the precessional regime with the errors reducing in the steady state regime. For all time steps, the implicit scheme shows a lower total error than the Heun scheme. At a time step of $\Delta t = 10$ the implicit scheme can be seen to correct itself back to zero error in equilibrium, whereas the explicit Heun scheme remains with a constant error in steady




Figure 4.3: Example solution of the deterministic Landau Lifshitz Gilbert equation for a single spin.

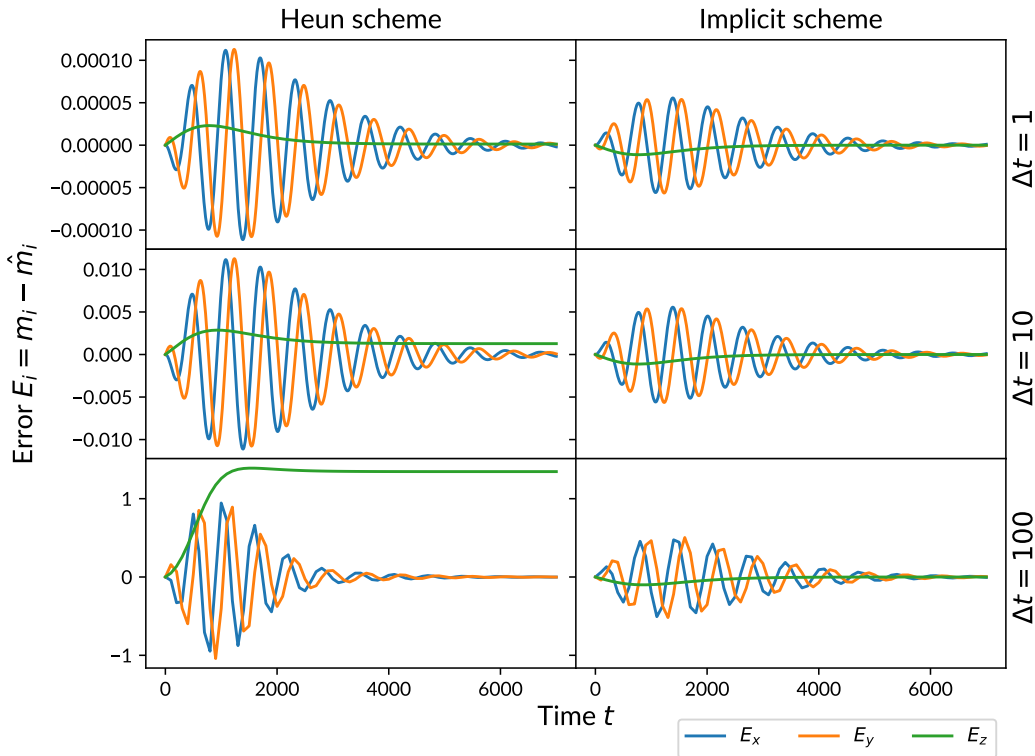


Figure 4.4: Path-wise error of numerical approximations to the deterministic LLG equation using the Heun and fully implicit schemes. The error is shown for three different time steps $\Delta t = 1, 10, 100$.

state. At even larger time steps $\Delta t = 100$ the Heun scheme becomes unstable while the Implicit scheme still achieves very low error in steady state.

For deterministic equations, the Heun and implicit scheme reduce to the deterministic Heun and implicit midpoint respectively. These have a proven strong order of 2.0 [115, 142]. In order to test the implementations of the methods, the convergence was obtained empirically. Each method was applied to the LLG to approximate the global




Figure 4.5: Empirical strong convergence of the Euler, Heun, and implicit ($\epsilon = 10^{-9}$) methods on the deterministic LLG.

solution $\hat{\mathbf{m}}_n$ to the true solution $\mathbf{m}(t_n)$ with a time step Δt . The error was defined as $e = \|\hat{\mathbf{m}}_n - \mathbf{m}(t_n)\|_2$ and from the convergence equation (2.22) we estimate the strong convergence rate γ by fitting a straight line in log-log space of the global error with the time step:

$$C_{2,\gamma} = \arg \min_{C_{2,\gamma}} \sum_{\Delta t} (C_2 + \gamma \log(\Delta t) - \log(e))^2 \quad (4.28)$$

Figure 4.5 shows a log-log plot of the global error e for each scheme for a range of time steps. The convergence plot shows a linear relationship for all schemes but with some divergence at the tails. At large Δt the Heun and Euler schemes become unstable, as was seen for the Heun scheme in the error plots in Figure 4.4. At small Δt the implicit scheme error also diverges and begins to increase. This is due to the tolerance of the internal Quasi-Newton solver (QNS) (see ϵ in Section 4.3.3), which has been set to $\epsilon = 10^{-9}$ in this case. The empirical convergence for the Heun and implicit schemes, obtained by fitting all points except the outliers (hollow circles), was approximately 2.0 as expected (shown in the legend).

The effect of the QNS tolerance on the performance of the fully implicit method was investigated by varying the tolerance and repeating the convergence experiment. Figure 4.6 shows that the implicit scheme behaves identically and converge as expected when the global error is larger than the tolerance of the internal QNS method. However, as the time step reduces, the error from the discrete time approximation becomes close to the tolerance and cannot be reduced further by decreasing the time step. The implicit solver benefits from using the smallest value of ϵ possible. The price to be paid for reducing ϵ is more iterations of the QNS method and therefore more computations of the effective field (Section 4.3.3).

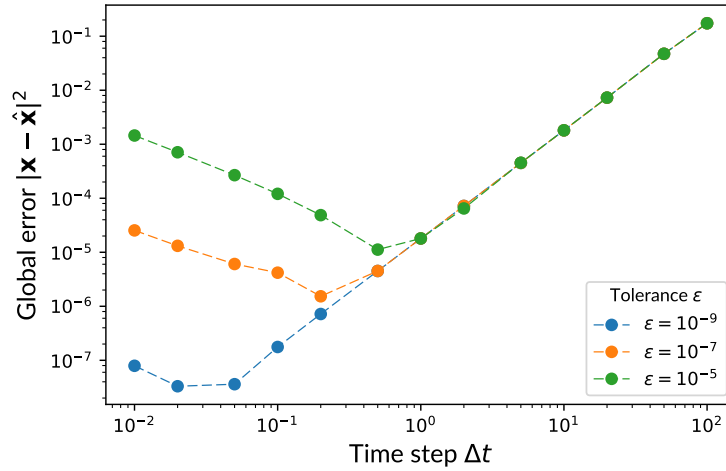


Figure 4.6: Empirical strong convergence of the implicit scheme on the deterministic LLG with varying QNS tolerance ϵ .

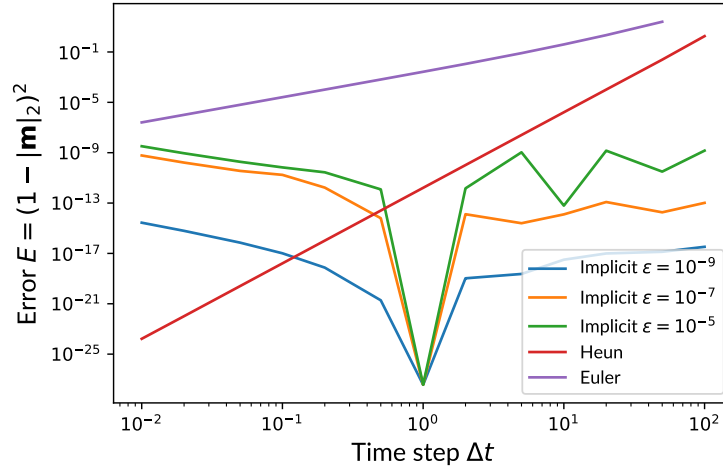


Figure 4.7: Global error in the magnetic moment magnitude as a function of the time step for Euler, Heun, and implicit (with varying tolerance ϵ) methods. The implicit scheme shows an unexpected results at $\Delta t = 1$.

The magnetic moment magnitude $|\mathbf{m}|^2$ should also be conserved throughout the simulation. The global error in the magnitude $1 - |\mathbf{m}|_n^2$ with decreasing time step Δt is shown in Figure 4.7. The Heun scheme shows a very rapid linear convergence in log-log space and can achieve lower error than the implicit scheme with a small time step. The implicit scheme error is around the tolerance level ϵ . The implicit error also showed an unexpectedly low error at $\Delta t = 1$. This phenomenon has no obvious cause and requires investigating further.

At zero temperature, the Heun and implicit methods both showed the same convergence rate and similar error on a simulation of the LLG equation with $h = 0.01, \alpha = 0.1$ for a single magnetic moment. Although the implicit scheme showed approximately 50%

lower error for most cases, the significant increase in complexity would not justify its use over the simpler Heun scheme. In [53] it was suggested that the computational cost of IMP may be reasonable for larger systems of magnetic moments on a spatial array. If guaranteed stability and highly accurate conservation of the magnetisation length at larger time steps are required then the implicit scheme is the recommended choice for single magnetic moments and larger systems. The Euler scheme showed much lower accuracy than the Heun and implicit methods.

4.4.3 Stochastic Landau-Lifshitz-Gilbert equation

The convergence rates and time-varying error of the schemes were analysed easily in the case of the deterministic LLG because an analytic solution was available. In the case of the stochastic LLG, no such solution exists and it is therefore more difficult to quantify the relative performance of each scheme during the dynamic phase of the sLLG. In the following section we first investigate a simplified form of the sLLG with a scalar noise process, which allows an analytic solution, before estimating the convergence of the methods on the full sLLG.

4.4.3.1 Simplified analytic case

If the thermal field in the sLLG is limited to the z -direction only, it is possible to obtain an analytic solution to this simplified sLLG for a constant applied field in the z -direction and the magnetic moment initially aligned along the x -direction. The solution is [87]:

$$\mathbf{m}(t) \begin{bmatrix} m_x(t) \\ m_y(t) \\ m_z(t) \end{bmatrix} = \begin{bmatrix} \operatorname{sech}(\alpha(ht + \sigma W_t)) \cos(ht + \sigma W_t) \\ \operatorname{sech}(\alpha(ht + \sigma W_t)) \sin(ht + \sigma W_t) \\ \tan(\alpha(ht + \sigma W_t)) \end{bmatrix} \quad (4.29)$$

This solution allows the convergence rate to be empirically determined for a simplified case of the stochastic LLG. Care must be taken however, that the results of these experiments cannot be generalised in entirety to the case of the full sLLG equation. Since the noise only occurs in the z direction, noise process is effectively scalar such that:

$$\Delta \mathbf{W}_n = \begin{bmatrix} 0 \\ 0 \\ \Delta W_n^2 \end{bmatrix} \quad (4.30)$$

From the theoretical discussion of the Heun (Section 4.3.2) and implicit (Section 4.3.3) schemes, the convergence rate is expected to be 1.0 in the commutative case but 0.5 in the non-commutative case. Recalling the definition for a commutative noise structure in equation (2.17), the sLLG with scalar noise is commutative whereas the full sLLG is not.




Figure 4.8: A single continuous Wiener process sampled at three different intervals. The Wiener processes used to compute the convergence rates must be identical when using different time steps.

In the stochastic case, each trajectory simulated depends upon a Wiener path where W_t^i for $0 \leq t \leq t_n$ is the i^{th} multidimensional Wiener path. In order to compute the empirical strong convergence rate, the expected error is computed from an ensemble of trajectories. The strong global error

$$e = E \langle \|\hat{\mathbf{m}}_n - \mathbf{m}(t_n)\|_2 \rangle = \frac{1}{N_{\text{runs}}} \sum_{i=1}^{N_{\text{runs}}} \|\hat{\mathbf{m}}_n^i - \mathbf{m}(t_n)\|_2 \quad (4.31)$$

where $\hat{\mathbf{m}}_n^i$ is the approximation of the magnetic moment at time t_n simulated with the i^{th} Wiener path. Each of the N_{runs} Wiener trajectories was generated using a random number generator with a different initial seed. Moreover, for each different time step, the equations must be simulated with the same Wiener processes but with less frequent sampling. An example of a single continuous Wiener process $W(t)$ sampled at different intervals is shown in Figure 4.8. The Wiener increments are first generated for the smallest possible time step simulated. The Wiener increments for subsequent, larger time steps are obtained by summing consecutive time steps along the path.

The Heun and implicit schemes were used to generate 1000 solutions starting from the same initial condition, for a range of time steps. The global error for each time step was computed and the strong convergence γ was obtained using the same fitting principle as previously (equation (4.28)). The results of the convergence experiment are shown




Figure 4.9: Comparison of the Heun and implicit schemes for the simplified stochastic LLG equation. The two schemes were used to estimate the trajectory of the equation with $H = 0.01, \alpha = 0.1, \sigma = 0.05$ (true solution shown in left panel). The empirical strong convergence of the schemes was estimated (middle panel). The global magnetic moment error was computed for each time step Δt (right panel).




Figure 4.10: Comparison of the Heun and implicit schemes for the simplified stochastic LLG equation. The experiment from Figure 4.9 was repeated with the parameters $H = 0.01, \alpha = 0.1, \sigma = 0.25$.

in Figure 4.9. As predicted, both methods converge with strong order 1.0. In contrast to the deterministic LLG the implicit scheme leads to a larger decrease in error. The convergence of the magnitude of the magnetic moment at time t_N is also shown as a function of the time step. The results are analogous to the deterministic case: the error for the Heun scheme monotonically decreases while the error of the implicit scheme increases. This is assumed to be a result of the tolerance of the QNS method, which introduces an error of around ϵ at each time step. A smaller time step requires more steps to reach t_N and therefore more errors accumulate. These results suggest that the tolerance of the implicit solver should be reduced as the time step is reduced.

Introducing a stochastic element into the LLG increased the relative improvement in the implicit scheme over the Heun scheme. Figure 4.10 shows the results of the same convergence test but with a stronger thermal field value $\sigma = 0.25$. The results show that the error of both methods is larger but that the implicit method has increased proportionally less, leading to an even greater increase in performance. As suggested in the initial stability analysis, the implicit scheme appears much more robust to systems with a strongly varying stochastic component.

The Heun scheme is often implemented with a renormalisation step, as discussed earlier




Figure 4.11: Comparison of the implicit scheme and the Heun scheme with an added renormalisation step using the same parameters from Figure 4.9. The missing data points (right panel) correspond to $1 - \|\mathbf{m}\|_2 = 1$ to machine precision.

(Section 4.2). This renormalisation has been shown to help stabilise the Heun scheme at large time steps [140] as well as ensure that the magnetisation length is conserved. The convergence test in Figure 4.9 was repeated but the Heun scheme renormalised the magnetic moment amplitude at each step. The results are shown in Figure 4.11. Interestingly, the results show that even though the renormalisation ensures that the magnetisation length is correct, it actually has very little effect on the path-wise error. From this it may be concluded that a) the renormalisation step does not necessarily improve the performance of a scheme (measured by path-wise global error) and b) that the error in the magnetic moment length does not give information on the true accuracy of the scheme.

4.4.3.2 Non-analytic case

Although there is no analytic solution to the full sLLG equation, it is possible to obtain the strong convergence using the relative error between different time steps [177]:

$$E \left\langle \left\| \hat{\mathbf{m}}_n^{\Delta t} - \hat{\mathbf{m}}^{\Delta t/2} \right\|_2 \right\rangle \leq C \Delta t^\gamma \quad (4.32)$$

Equation (4.32) can be used to estimate the convergence rate but the absolute value of the global error remains unknown, which makes their performance difficult to compare. Computing the relevant convergence was achieved by simulating the sLLG with the Heun and implicit schemes with time steps increasing in powers of two $\Delta t = h_0, 2h_0, 4h_0, 8h_0$ where h_0 is the smallest time step. An ensemble of Wiener processes was simulated as before, ensuring that the same processes were used for both schemes and all time steps. The results, Figure 4.12, agree with theoretical results that both schemes have a convergence of 0.5 on the full sLLG. Note that the convergence results are different from the simplified (commutative sLLG) in Figure 4.9.

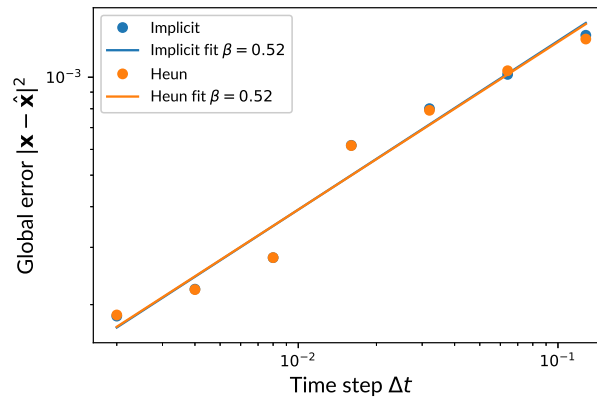


Figure 4.12: Empirical strong convergence of the Heun and implicit methods of the full sLLG equation, estimated using the convergence of the relative global error between subsequent time steps (equation (4.32)).

Table 4.2: Performance results for the fully implicit and Heun scheme solving the sLLG equation for a single magnetic moment. The results are in units of CPU time per unit of simulation time (larger is worse). The speed-up is the relative increase in performance from using the Heun scheme.

$\Delta t \times 10^{-3}$	1	2	4	8	16	32	64	128
Implicit	3555	1751	853	411	201	100	50	25
Heun	190	95	47	23	11	5	3	1
Speed-up	18.6x	18.4x	17.9x	17.3x	17.0x	16.9x	15.9x	17.1x

4.4.4 Performance

The time taken to perform the simulations during the convergence experiment was measured to compare the efficiency of the Heun and implicit schemes. Table 4.2 shows the average time taken to simulate the trajectories for each time step over the 1000 runs with different Wiener process paths. The units are a ratio of CPU time to simulation time. A score of 100 indicates that the method takes on average 100ms of CPU time to simulate the system for 1ms (or equivalently 100s for 1s of simulation time and so on). The results show that the Heun scheme is between 15–20x faster than the implicit scheme for integrating a single magnetic moment. The results of the convergence test of the simplified sLLG with high noise (Figure 4.10) showed that the implicit scheme can achieve the same error with a time step 50x larger than the Heun scheme. In both the linear test equation and sLLG experiments, the implicit scheme performed better at high noise. Therefore, we posit that for the full sLLG — for which we have no absolute error values — it is likely that the implicit midpoint scheme is more efficient than the Heun scheme. However, as the system size increases, the effective field calculation becomes very expensive. This is a particular problem for the implicit scheme, which must evaluate the sLLG many times per step of the simulation. Table 4.3 shows

Table 4.3: Speed-up of the Heun scheme relative to the IMP method for interacting particle chains of increasing length.

Particle chain length	Speed-up
1	11.7x
5	26.5x
10	66.6x
25	180.0x

the relative speed up of the Heun scheme over the implicit method for increasingly large systems of particles. The system simulated was a particle chain, with N particles aligned along their anisotropy axes. The results show that as the system size grows, the implicit method is substantially slower than the Heun method. Therefore, for large systems it is unlikely that the implicit scheme is a viable method for simulating sLLG dynamics.

4.4.5 Equilibrium

The implementation of the implicit scheme was validated by checking the equilibrium conditions of the function. While this has already been shown for the Heun scheme (e.g. [56]), the implicit scheme uses a truncated Wiener process and the subsequent effect on the thermal equilibrium was unknown. In Section 2.3.5, it was shown that the distribution of states in a canonical ensemble follows a Boltzmann distribution. Let $\theta^i = \arccos(m_z^i)$ be the angle between the z -axis and the magnetic moment of the i^{th} magnetic moment in the ensemble. The Boltzmann distribution for the state of the magnetic moment is defined over the solid angles of the magnetisation. Therefore the corresponding distribution for θ is:

$$p(\theta) = \frac{\sin(\theta)e^{E(\theta)/(k_B T)}}{Z} \quad (4.33)$$

where $Z = \int_x \sin(x) \exp[E(x)/(k_B T)] dx$ is the partition function. The energy of a magnetic moment with no applied field and an anisotropy axis aligned with the z -axis is:

$$\frac{E(\theta)}{k_B T} = -\sigma \cos^2(\theta), \quad \sigma = \frac{KV}{k_B T} \quad (4.34)$$

where σ is called the *stability ratio*. A non-interacting ensemble of 10,000 magnetic moments was simulated using the implicit midpoint scheme and compared to the analytic solution. The partition function Z was computed using the trapezoidal numerical integration scheme [164]. The simulation began with all moments initially aligned along the z -direction and dynamically evolved until the empirical distribution of the angles $p(\theta)$ was constant. Figure 4.13 shows the resulting histogram of each value of θ in the ensemble, compared with the analytic solution. The results show a very close match.

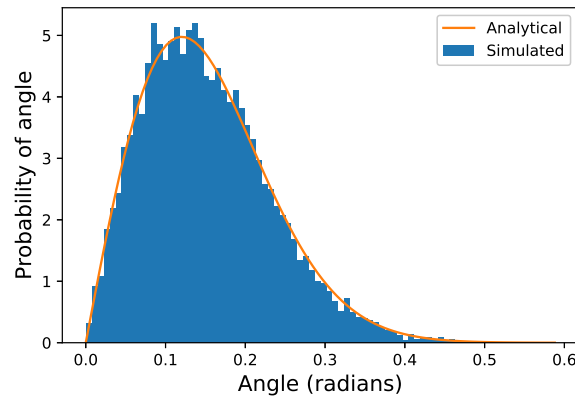


Figure 4.13: Equilibrium distribution of 10,000 non-interacting magnetic moments simulated with the fully implicit method and $\sigma \approx 35$. The normalised histogram resulting from the simulation closely matches the empirical distribution.

The Boltzmann distribution is not only valid for a single spin but for any arbitrary system. Although the single particle equilibrium is regularly tested in previous works, the distribution of larger systems (such as two interacting particles) are not covered. The energy for two particles includes the dipole-dipole interaction term and the energy becomes a function of the elevation and azimuth angle of both particles $E(\theta_1, \phi_1, \theta_2, \phi_2)$. In equation (4.33) the partition function was computed numerically. However, this becomes very computationally expensive with increasing dimensionality of the energy function. Therefore, the true equilibrium distribution was estimated using MCMC methods (Section 2.3.6), which for a sufficient number of samples are guaranteed to approach the true distribution.

Using MCMC methods, allows the thermal equilibrium resulting from simulations of ensembles of more complex systems to be tested. This additional test ensures that the interaction field has been implemented correctly. An interacting two particle system was simulated with both their anisotropy axes aligned in the z -direction. In this case the energy is constant for any ϕ_1, ϕ_2 and only varies with θ_1, θ_2 . Figure 4.14 shows the resulting equilibrium distribution of an ensemble of 50,000 systems compared with the results of the Metropolis MCMC sampler. The results match closely and are seen more clearly in Figure 4.15 that shows the marginal probability distribution of the elevation angle of the first particle $p(\theta_1) \int_{\theta_2} p(\theta_1, \theta_2) d\theta_2$. The marginal distribution for the first magnetic moment was also computed without the interaction term to ensure that the interaction strength was strong enough to shift the equilibrium distribution.

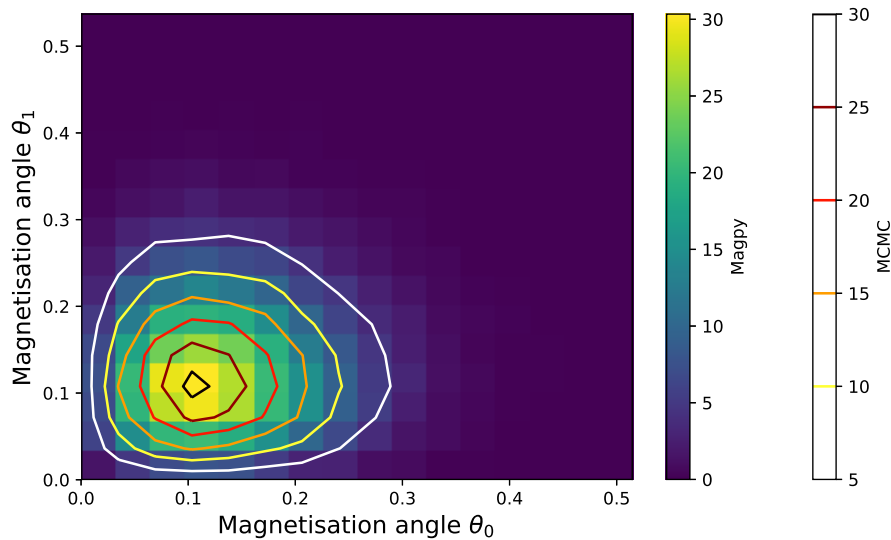


Figure 4.14: Equilibrium distribution for a system of two interacting magnetic moments computed using stochastic simulation of an ensemble of 50,000 systems and the Metropolis MCMC method.

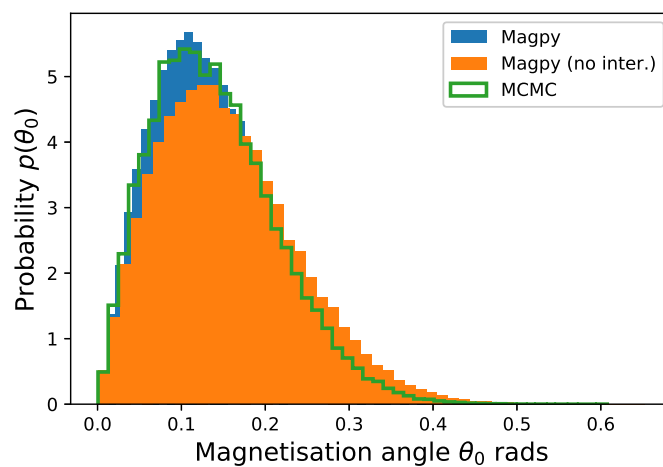


Figure 4.15: The marginal equilibrium distribution for a single spin in a system of two magnetic moments. The results compare the results of stochastic simulation using an ensemble of 50,000 systems with and without the interaction term, against the results of MCMC sampling.

4.5 Conclusions

A range of testing procedures exist for determining the accuracy of numerical methods for the stochastic simulation of the sLLG. The literature revealed a broad range of possible tests such as: energy conservation, thermal equilibrium, hysteresis loops, experimental data, and the conservation of the magnetic moment. However, the fundamental performance metric for numerical solvers is whether the simulated solutions are close to true trajectories of the sLLG equation. The simplified sLLG equation, which has not been utilised often in the literature, allows numerical solutions to be compared with the analytic solution to obtain absolute error values. However, the full suite of possible tests, including measuring the computational cost, should be employed to select the most suitable method.

The implicit midpoint method was implemented and compared against the Heun scheme in a series of tests. The implicit midpoint method consistently achieved lower error, especially on systems with a large noise component. The practical implementation of the implicit midpoint method was discussed in detail, revealing that solving the implicit step required substantially more evaluations of the effective field term. In the end, the implicit midpoint scheme proved computationally unfeasible for large systems.

Additionally, using the analytic solution of the deterministic and simplified sLLG equation, it was shown that the numerical methods were most prone to error in the dynamical region. Studies of the equilibrium properties are not sufficient to determine the performance of numerical schemes if the dynamics are to be simulated accurately. Moreover, it was shown that the error in the magnetic moment length does not necessarily indicate path-wise closeness of the numerical scheme to the true solution. Consequently, introducing a renormalisation step does not significantly improve the error.

There are a number of open questions remaining for the development of efficient numerical methods for the sLLG:

- In other research fields, particularly finance, weak solvers are often employed to obtain the expected solution in a more efficient way; yet their use for solving the sLLG has not been explored until very recently [2]. Many stochastic simulations of the sLLG do not require path-specific information and could be simulated significantly faster.
- Existing studies have presented high order Runge-Kutta solvers for the sLLG [126], which avoid simulating multiple stochastic integrals. Empirical results suggest that these methods generate correct solutions. The mathematical basis for these approaches deserves further investigation.

- Adaptive time step methods are a natural next step for sLLG solvers and are being explored [126]. Interesting examples of adaptive methods in other domains (e.g. [99]) could be attempted in magnetism.
- The SIA and SIB semi-implicit schemes are highly promising solutions to the computational difficulties of the fully implicit method. Further empirical evidence or mathematical proofs of the validity of these heuristic algorithms would be welcomed by the community.
- Advanced computational methods utilising general purpose graphical processing units would be particularly powerful for parallelising the computation of multiple integrals (e.g. for the Milstein scheme) or speeding up methods such as SIA and SIB for very large systems.

Chapter 5

Design of Magpy

The Magpy software package, as described in this chapter, was developed solely by the author during the present research project. The final package is open-source and available online.

5.1 Introduction

Magpy is an open-source C++ and Python package that models nanoparticles and simulates their magnetic state over time as described by the Landau-Lifshitz-Gilbert (LLG) equation [120]. Since the best choice of numerical method to solve the LLG dynamics is an open research question, the numerical solvers in Magpy are implemented in a generic form, independent from the equations of magnetism, which allows the methods to be tuned or replaced easily. The current implementation includes the widely used Heun scheme [66] and the first open-source implementation of the derivative-free fully implicit scheme [142] (Chapter 4). Magpy also includes the thermal activation model for single non-interacting particles (Section 2.4.3), which avoids solving the LLG equation and consequently simulates nanoparticles with significantly less computational effort. Although Magpy was designed for simulating dynamic hysteresis loops for hyperthermia experiments, the software may be used to explain or predict the outcome of magnetic nanoparticle experiments in general.

5.2 A model for magnetic nanoparticle dynamics

Figure 5.1 shows a diagram of a magnetic nanoparticle, which comprises a large number of individual atoms arranged in a regular crystal lattice, each of which possesses a magnetic moment represented by a 3-dimensional vector. The net magnetisation of a material is simply the sum of the individual magnetic moments. For example, if




Figure 5.1: A two-dimensional sketch of a nanoparticle. The atoms of a magnetic material are packed into a regular crystal lattice and each is modelled by a three-dimensional magnetic moment that varies with time. Due to strong interactions between these magnetic moments in small particles, Magpy assumes they rotate coherently and are represented by a single macrospin vector.

the atoms are randomly oriented, the material has zero magnetisation; if they are all aligned, the material has a large magnetisation component. From Section 2.2.2, atomic magnetic moments prefer to align with one another due to the exchange interaction force. This force is strong enough that, in nanoparticles that are particularly small ($< 25\text{nm}$ in diameter), the individual moments are approximately aligned and rotate coherently. Magpy uses this approximation to model the state of a particles atoms by a single 3-dimensional macrospin, rather than simulate the atoms individually. Magpy is able to simulate much longer timescales for the same computational effort compared to simulating the individual atoms due the greatly reduced degrees of freedom. This is particularly relevant for medical experiments that may take seconds or hours but posses magnetic dynamics that must be resolved on a picosecond time scale.

The LLG is implemented in Magpy in the reduced form presented in Section 4.1.1 by introducing the *reduced time* variable ℓ . The equation of motion for a single particle i in an ensemble of N particles was given by equations (4.1) and (4.2). The effective field was computed using equation (4.6) and the thermal field acting on each particles using the strength equation (4.7). The three forces acting on each particle are depicted in Figure 5.2 and include the externally applied field, anisotropy field, and the dipole-dipole interaction field. The equations are solved numerically at discrete time steps, resulting in a simulated trajectory of the system's magnetic state. The simulation outputs, at each discrete time, the value of the applied magnetic field and the x, y, z components of the magnetic state of every particle in the system. These results may be used to obtain

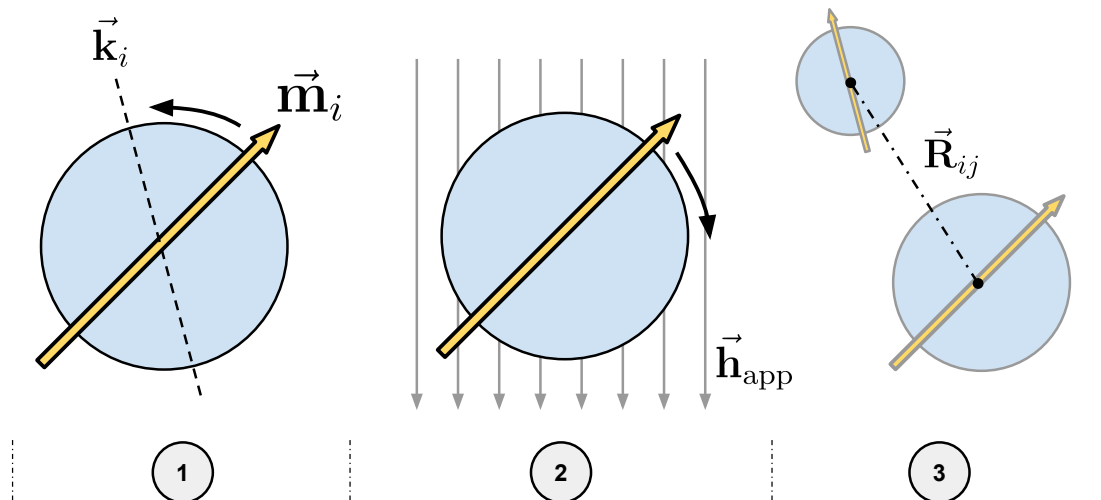


Figure 5.2: The three effective field contributions acting upon macrospin i . (1) The macrospin experiences a force towards alignment with the particle anisotropy axis (dashed line) \mathbf{k}_i in either direction. (2) The macrospin is also forced towards aligning with the externally applied field direction (solid arrows) \mathbf{h}_{app} . (3) Finally, each macrospin is repelled and attracted by nearby macrospins. The force of the dipolar-dipolar interaction diminishes with distance (dash-dotted line) between two particles \mathbf{R}_{ij} .

the total magnetisation of the system $M(t) = M_s \sum_i m_i(t)$, the average magnetisation of an ensemble of systems, static and dynamic hysteresis loops, and the energy dissipated by the system.

Multiple simulations with different random seeds but with identical initial conditions will result in different solutions due to the stochastic nature of the thermal field. For example Figure 5.3, shows the results of five simulations of a 3-particle chain from the same initial condition. The Magpy script used to generate the results is shown in Listing 1. In addition to individual trajectories, the expected system trajectory and higher order statistical moments may be obtained by running a large number of simulations and averaging their results (i.e. stochastic simulation). The number of simulations required to obtain reasonable estimates of these statistical variables is large and depends on the system of interest.

5.3 Thermal activation for single particles

Magpy provides an alternative thermal activation model for simulating non-interacting, anisotropy-dominated particles. A particle is considered anisotropy-dominated if the effective field resulting from anisotropy k is much greater than the thermal fluctuations and the externally applied field such that $\sigma(1-h)^2 \gg 1$ (where $\sigma = KV/(k_B T)$ is termed the reduced energy barrier height). Magpy approximates the full dynamics as a jump process between two discrete states (up and down), which is described mathematically

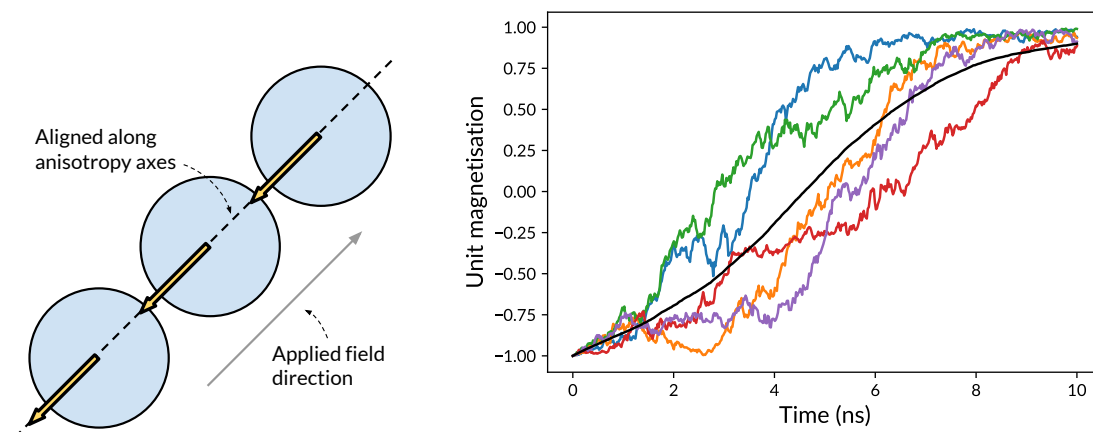


Figure 5.3: Simulating a chain of three particles (created in Magpy using Listing 1). (left) A chain structure of three identical particles are initialised with an external field applied along their anisotropy axis and their magnetisation initially against the applied field. (right) The coloured lines show the total magnetisation in the direction of the magnetisation for 5 simulations from the same initial condition (see left). The black line is the result of averaging 500 simulations from the same initial condition (i.e. the expected or mean trajectory).

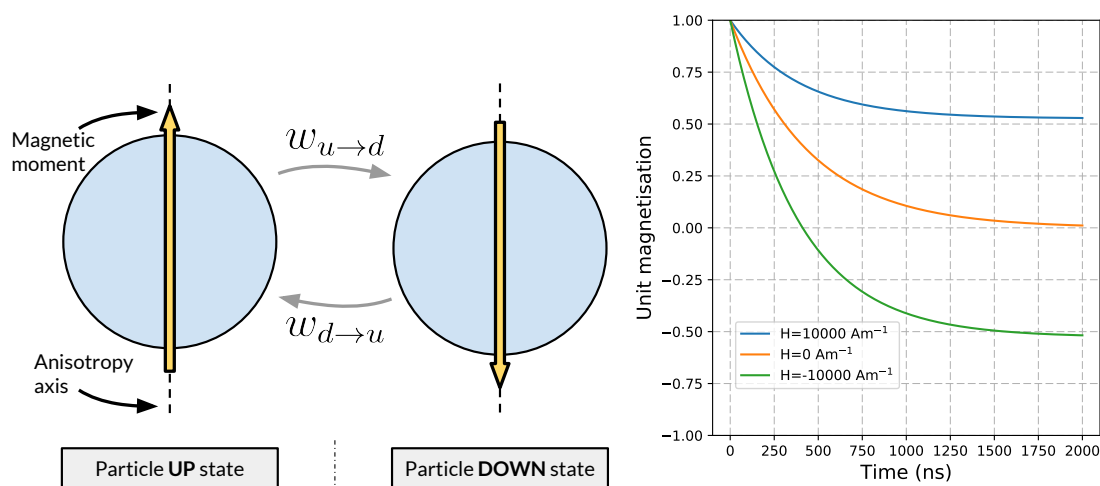


Figure 5.4: In the thermal activation model, a single nanoparticle may occupy one of only two states: *up* or *down*, which each correspond to local energy minima around the anisotropy axis (left). The dynamics of the system are described by a master equation with transition rates $w_{u \rightarrow d}/d \rightarrow u$ between the two states. The solution of the master equation depends on the initial condition, particle properties and applied field. In this example (right) the particle is initialised up with probability 1 and allowed to relax into equilibrium (see Listing 2 for Magpy script). The equilibrium magnetisation depends upon the strength of the constant applied magnetic field.

```

1 import magpy as mp
2
3 chain3_model = mp.Model(
4     radius=[8e-9, 8e-9, 8e-9],
5     anisotropy=[4e3, 4e3, 4e3],
6     anisotropy_axis=[
7         [0., 0., 1.], [0., 0., 1.], [0., 0., 1.]],
8     magnetisation_direction=[
9         [0., 0., -1], [0., 0., -1], [0., 0., -1]],
10    location=[
11        [0., 0., -20e-9], [0., 0., 0.], [0., 0., 20e-9]],
12    magnetisation=400e3,
13    damping=0.1,
14    temperature=300.,
15    field_shape='constant',
16    field_amplitude=30e3)
17
18 chain3_ensemble = mp.EnsembleModel(base_model=chain3_model, N=500)
19 results = ensemble.simulate(
20     time_step=1e-13, end_time=1e-8, max_samples=500, n_jobs=4)
21
22 time = results.time
23 first_run_magnetisation = results.results[0].magnetisation()
24 ensemble_magnetisation = results.ensemble_magnetisation()

```

Listing 1: Simulating an ensemble of five hundred three-particle chains in Magpy (results shown in Figure 5.3). The three-particle chain model is instantiated (lines 3-16) using the `Model` object, which is defined by the properties and locations of each particle in the chain and the applied field. An ensemble of models is created using the `EnsembleModel` object (line 18), which simply represents a collection of individual models and is provided for convenience. The five hundred models are individually simulated (lines 19-20) and the computational work is distributed across four processes by setting `n_jobs=4`. The resulting magnetisation is computed for an individual model (line 23) and the entire ensemble of models (line 24).

by the master equation:

$$\frac{d}{dt} \begin{bmatrix} p_u(t) \\ p_d(t) \end{bmatrix} = \begin{bmatrix} -\Gamma_{u \rightarrow d}(t) & \Gamma_{d \rightarrow u}(t) \\ \Gamma_{u \rightarrow d}(t) & -\Gamma_{d \rightarrow u}(t) \end{bmatrix} \begin{bmatrix} p_u(t) \\ p_d(t) \end{bmatrix} \quad (5.1)$$

where the elements of $\mathbf{p}(t) = [p_u(t), p_d(t)]^T$ are the probability that the system is in the up and down state respectively and $\Gamma_{u \rightarrow d}(t), \Gamma_{d \rightarrow u}(t)$ are the transition rates (units s^{-1}) between the two states. Note that the solution of the master equation is the time-evolution of the *probability mass function* over the discrete state space, whereas the solution of the Landau-Lifshitz-Gilbert equation is the time-evolution of a *single random trajectory* through the state-space \mathbb{R}^3 . The transition rates are computed using the Néel-Brown formula (equation (2.43)), which assumes that the field is applied parallel to the

```

1 import magpy as mp
2
3 Hs = [100e2, 0.0, -100e2]
4 models = [
5     mp.DOModel(
6         radius=5e-9, anisotropy=5e4, damping=0.01,
7         magnetisation=400e3, temperature=300,
8         initial_probabilities=[1.0, 0.0], field_amplitude=H)
9     for H in Hs
10 ]
11 results = [
12     model.simulate(end_time=2e-6, time_step=1e-10, max_samples=1000)
13     for model in models
14 ]
15 expected_magnetisations = [res.magnetisation() for res in results]

```

Listing 2: Simulating three thermal activation models with different applied field properties with Magpy (results plotted in Figure 5.4). The `DOModel` object, representing the thermal activation model, is defined (lines 3-6) by the particle and applied field properties. Three identical particles are modelled each with a different value of the constant applied field amplitude (line 1). Each of the models is simulated (lines 9-12) and the expected magnetisation of each model is computed (line 13).

anisotropy axis direction. The master equation (5.1) is solved numerically from an initial condition $\mathbf{p}(t_0) = [p_u(t_0), p_d(t_0)]^T$ at time t_0 using an adaptive step Runge-Kutta solver (RK45) with Cash-Karp parameters [164]. The total magnetisation at time t for a large ensemble of particles is computed as $M(t) = M_s[p_u(t) - p_d(t)]$.

Figure 5.4 shows an example of a single particle simulated using the Magpy script in Listing 2, with a constant field applied along its anisotropy axis. The initial condition of the system is $\mathbf{p}(t_0) = [1, 0]^T$ and the master equation is solved numerically. As time evolves, the probability that the particle flips into the down state $p_d(t)$ increases and the expected magnetisation reduces. Eventually, the system reaches an equilibrium: in zero field the two states are equally likely and the system has zero magnetisation; for finite applied fields the system favours one state over the other.

5.4 Alternative software

Vinamax[127], implemented in Golang¹ and also motivated by the medical applications of nanoparticles, provides similar functionality to Magpy. A distinguishing feature of Vinamax is its use of a multipole-expansion algorithm, which greatly improves the speed of computing the dipolar-dipolar interaction forces for large systems. Magpy computes

¹<https://golang.org/doc> for more information on the Go programming language.

the interaction field between every pair of particles (equation (2.11)), an operation with complexity $\mathcal{O}(n^2)$; the multipole expansion method uses an approximation, which results in complexity $\mathcal{O}(n \log n)$. Vinamax does not implement an implicit numerical solver.

Magpy is unable to simulate magnetic systems for which the macrospin assumption is not justified, such as for larger particles that exhibit more than a single domain or systems for which surface-to-surface atomistic interactions are significant. In these cases, the magnetic moments of the individual atoms must be modelled. Vampire [56] is an open-source C++ alternative to Magpy for atomistic simulation. Vampire reduces the significant additional computational effort required for simulating individual atoms by leveraging general purpose graphical processing units (GPGPUs). Alternatively, if the effects of temperature can be ignored and the atomistic magnetic moments are closely aligned, the magnetisation of material can be represented as a continuous function resulting in a spatial-temporal partial differential equation. This technique, termed micro-magnetics, is implemented in a range of popular open-source packages: MuMax3 [195], OOMMF [47], fidimag [41], nmag [62].

As discussed, Magpy includes implementations for several numerical methods to compute approximate solutions to stochastic differential equations. Currently, the authors are not aware of a reliable alternative in C++ or Python for the fully implicit method [142]. Though there are mature packages for the solution of ordinary differential equations (e.g. sundials [93]) there are few options for stochastic differential equations. The most mature, SDElab [79] implemented in Matlab, is no longer under development and requires proprietary software. A re-implementation of SDElab using the open-source julia language is currently under development².

5.5 Implementation and architecture

Magpy consists of two components. Firstly, a C++ library implements the core simulation code, which comprises the nanoparticle model and numerical solvers. The second component is a Python interface to the C++ library functionality with additional features for setting up simulations and analysing their results.

The dynamics (Landau-Lifshitz-Gilbert equation), thermal activation model, numerical methods, and the effective field calculations are implemented in a C++ library. C++ was the preferred programming language for implementing the computationally-intensive simulation because of its relatively fast performance and opportunities for optimisation. The Magpy C++ library is optimised for serial execution; uses the BLAS and LAPACK libraries; manages memory manually to minimise allocations and deallocations; and may be compiled with proprietary Intel compilers for enhanced performance on Intel architectures. Furthermore, the C++-11 standard contains features that support a functional

²<https://github.com/tonyshardlow/SDELAB2> for development updates on SDELab2.

programming paradigm (such as closures and partial application), which were used extensively in Magpy to improve testability and modularity of code. The entry point to the Magpy library is through two top-level functions: `simulation::full_dynamics` for the full model and `simulation::dom_ensemble_dynamics` for the thermal activation model. Magpy does not provide a graphical user interface, simulations must be invoked by the user in a C++ program or using the alternative Python interface.

It was the authors' opinion that scripting in C++ was not sufficiently usable because of the low-level syntax, poor availability of plotting tools, and the requirement for compiling scripts, which adds complexity for users. Python, on the other hand, is high-level, interpreted, and has been gaining popularity in the computational science community for the design of user interfaces [15, 59, 131] and as an easy-to-learn tool [58]. Therefore, Python was chosen as the preferred language for scripting and implementing the auxiliary components of Magpy. The interface between Python and C++ was written using Cython [16], which allowed the C++ library functions to be wrapped as Python functions and exposed as a Python package, while retaining the performance benefits of C++. The Python package includes additional features for building models and plotting the simulation results.

The typical workflow for a Magpy experiment consists of running multiple simulations of the same model in order to generate a distribution of possible trajectories, as in Figure 5.3. This motivates an embarrassingly parallel strategy in which each simulation executes concurrently on a single process, since no communication is required between the independent runs. Parallelism is implemented in Python using `joblib`³. A minimum example of how `joblib` is used to execute tasks in parallel is shown in Listing 3. In Magpy, the user creates an ensemble of models (the `EnsembleModel` object in Listing 1 lines 3-18) and begins the simulation (`EnsembleModel.simulate` lines 19-20) utilising the requested number of cores (`n_jobs`). For each model in the ensemble, Magpy creates a new independent python process containing a copy of the model object. Each process then simulates its respective model by calling functions in the C++ library with the model parameters. As many as `n_jobs` simulations may execute concurrently. Once each simulation finishes, the results are returned from the C++ library to the individual python process. Once all processes have completed, the results are gathered back into the python process with which the user was originally interacting. This architecture is displayed graphically in Figure 5.5.

5.6 Quality control

Magpy has been tested to increase confidence in the correctness of the implementation, mathematics, and physics. The lowest level of tests, unit tests, assert that individual

³<https://pythonhosted.org/joblib/> for the `joblib` documentation.

```

1  from joblib import Parallel, delayed
2  import time
3
4  def slow_double(x):
5      time.sleep(1) # 1 second sleep
6      y = 2*x
7      return y
8
9  xs = [2, 6, 12, 24, 40, 72, 126, 240]
10
11 # Serial computation takes approximately 8s
12 ys_serial = [slow_double(x) for x in xs]
13 print(ys_serial)
14 #> [4, 12, 24, 48, 80, 144, 252, 480]
15
16 # Embarrassingly parallel computation takes approximately 2s
17 ys_parallel = Parallel(n_jobs=4)(delayed(slow_double)(x) for x in xs)
18 print(ys_parallel)
19 #> [4, 12, 24, 48, 80, 144, 252, 480]

```

Listing 3: A minimal example of an embarrassingly parallel computation with `joblib`. The function `slow_double` (line 4) doubles a single number and takes approximately one second. The objective is to evaluate the function with eight different arguments (line 9). This is achieved in serial with a `for` loop (line 12) by evaluating the function for each argument in the list singly, taking approximately eight seconds. However, this problem is embarrassingly parallel because all evaluations of `slow_double` may occur concurrently since each function call only depends on its initial argument. Using `joblib`, the eight function calls are evaluated on four processes as shown (line 17) by setting `n_jobs=4`. Two evaluations are distributed to each of the four processes, which execute concurrently, taking approximately two seconds.

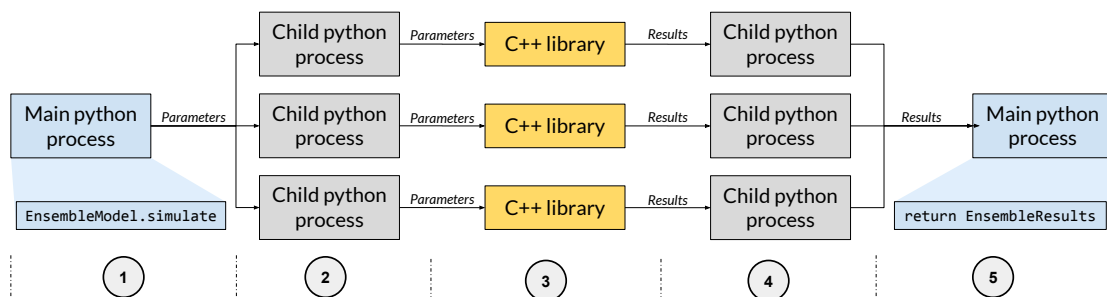


Figure 5.5: The flow of data through an ensemble simulation in Magpy. (1) The user instantiates an `EnsembleModel` object and calls the `EnsembleModel.simulate` function, specifying the number of CPU cores to utilise. (2) The main Python process spawns a new individual process for each model in the ensemble. (3) The individual processes each call the Magpy C++ library using their respective model parameters. (4) The results from the C++ simulation are returned to the individual python process. (5) When all the processes have finished, the results are collected on the main process to be analysed and plotted.

functions return the correct answer given a set of fixed arguments. The unit tests are designed to catch bugs during development and test the installation of the software. Continuous integration, using CircleCI⁴, ensures that tests are automatically executed before changes are committed to the existing code repository on Github. The unit tests are implemented using GoogleTest⁵ for C++ functions and pytest⁶ for Python functions.

Numerical tests are necessary to confirm the stability and robustness of the numerical methods. Magpy includes scripts to evaluate the empirical convergence rates of the SDE solvers and compares them with analytic solutions [115, 142]. The numerical tests should be used during the development of new or existing solvers. Finally, Magpy includes a series of Jupyter notebooks⁷ that present tutorials and examples, including comparisons of simulation results with theoretical solutions from alternative models in physics. These comparisons assert that the simulations, under the appropriate assumptions, correctly approximate the magnetic nanoparticle dynamics.

Magpy was primarily designed for simulating the magnetic dynamics of nanosized particles. The simulation results may be used to compute heat dissipation, relaxation rates, and equilibrium states allowing the software to help predict, explain, or otherwise augment traditional experiments in the laboratory or clinical settings. However, using numerical simulation also allows the exploration of a range of material geometries and properties without expensive equipment or physical limits. Magpy was designed to be accessible to experts and non-experts through the extensive documentation and included examples.

In addition to its uses in physics, the implementation of the numerical solvers for stochastic differential equations may be useful beyond the original purpose of Magpy. The Landau-Lifshitz-Gilbert equation belongs to a class of equations (multi-dimensional, nonlinear, stochastic, non-commutative, stiff) that are challenging to solve numerically. Magpy could also be used for teaching concepts in magnetism as the Python interface will likely be familiar to new students in physics.

A number of additional features remain that have yet to be implemented in Magpy. In particular, the use of a multipole expansion method (inspired by Vinamax) would reduce the time required to compute the interaction fields. It would also be possible to extend Magpy to simulate atomistic-level dynamics by decomposing each macrospin into a lattice of atomistic moments and including the exchange interaction term to the effective field. The thermal activation model currently supports a single particle with the field applied along its anisotropy axis. Allowing arbitrary applied field directions as well as dipolar-dipolar interactions between multiple particles would greatly increase the potential applications of the model.

⁴<https://circleci.com/> for more information.

⁵<https://github.com/google/googletest> for the GoogleTest repository.

⁶<https://docs.pytest.org/en/latest/> for the pytest documentation.

⁷Magpy documentation and examples are hosted at <http://magpy.readthedocs.io>.

5.7 Example: the first passage time distribution

The validity of the thermal activation model has been confirmed by comparing experimental data [200], simulated solutions [179], and analytic solutions [37] to the results of the model. In all three references the relaxation time obtained from the Néel-Brown model (equation (2.43)) was compared to the results to the experimental/simulated/analytic relaxation rate in order to quantify validity. However, the relaxation rate is an ensemble statistic, which has limited descriptive power. The relaxation rate is equivalent to the mean first passage time (MFPT), which describes the average time taken for the particle to switch from its initial state. The first passage time distribution (FPTD) is commonly used to quantify the behaviour of a stochastic process. A more rigorous validation of the approximate dynamics of the thermal activation model would be to compare the FPTD with stochastic simulation of the Landau-Lifshitz-Gilbert equation. In the following section, Magpy is used to simulate the FPTD of a single magnetic nanoparticle to quantify the validity of the thermal activation model for a range of damping regimes and energy barrier sizes.

The first passage time distribution for a 2-state master equation can be obtained analytically [68]:

$$\text{FPT}(t) = \Gamma e^{-t\Gamma}, \quad \Gamma = \Gamma_{12} + \Gamma_{21} \quad (5.2)$$

The FPTD for the thermal activation model is an exponential distribution parameterised by the Néel-Brown transition rates. The mean first passage time is easily obtained from the exponential distribution and is equivalent to the relaxation rate $\tau = 1/\Gamma$: $\text{MFPT} = E\langle \text{FPT}(t) \rangle = \tau$. For the Magpy simulations, the first passage time was defined as the time taken for the magnetisation to reach a specific domain $\mathcal{D} \in \mathbb{R}^3$. In order to compare the solution to the thermal activation model, the two discrete states had to be mapped into \mathbb{R}^3 . A natural choice was that the *up* state was defined as $\{m \in \mathbb{R}^3 : m_z > 0\}$ and the *down* state as $\{m \in \mathbb{R}^3 : m_z \leq 0\}$. However, the magnetisation trajectory may cross the switching condition $m_z = 0$ and return without switching to the other minima [66]. Therefore, after initialising the system in the *up* state, the domain for which the particle was considered to have switched was defined as $\{m \in \mathbb{R}^3 : m_z < 0.5\}$. It should be noted that the FPTD obtained from simulation is conditional on this choice and will vary with different definitions of the switching conditions.

As noted by Kalmykov et.al. [109], it is important to initialise a system in thermal equilibrium distribution in order to obtain the correct first passage time. Initialising a particle in the *up* state thus referred to an initial condition that was Boltzmann distributed in the elevation angle (Section 2.3.5):

$$P(\theta) = \frac{\sin(\theta) \exp\left\{-\frac{E(\theta)}{k_B T}\right\}}{\int_0^{\pi/2} \sin(\theta) \exp\left\{-\frac{E(\theta)}{k_B T}\right\}} \quad (5.3)$$

In order to quantify the validity of the transition rate model for different parameters, a performance metric was required. The dynamics of the transition rate are expected to diverge from the results obtained from Langevin dynamics as the energy barrier is reduced. The qualitative change in the distribution was quantified using the Kolmogorov-Smirnov statistic, which calculates the divergence between an analytic and an empirical distribution. The KS-statistic provided a measure of confidence that the FPTD obtained from numerical simulations was drawn from the analytic distribution from the master equation. The statistic is sensitive to both shape and location mismatch and thus provided a suitable performance indicator of the applicability of the thermal activation model.

The cumulative distribution of first passage times is:

$$F(t) = \int_{-\infty}^t \text{FPT}(\hat{t}) d\hat{t} \quad (5.4)$$

The empirical cumulative distribution function is the set of n ordered observations of first passage times from simulations $\hat{t}_1 < \hat{t}_2 < \dots < \hat{t}_n$:

$$F_n(t) = \frac{1}{n} \sum_{i=1}^n I_{[-\infty, t]}(\hat{t}_i) \quad (5.5)$$

The resulting KS statistic is:

$$\text{KS} = \sup_t |F_n(t) - F(t)| \quad (5.6)$$

The statistic represents the confidence that the empirical distribution differs significantly from the analytic distribution. Therefore the lower the score, the higher the performance.

5.7.1 Results

An ensemble of 500 particles with $\sigma \approx 9$ were simulated using Magpy. The particles were initially distributed in a Boltzmann distribution around the up state. The trajectories were simulated and the first passage time was computed. Figure 5.6 shows the MFPT computed using Magpy compared to the Néel-Brown model. The results show, as expected, that the Néel-Brown model is valid in the intermediate-to-high damping regime (see [38] for a discussion on the damping dependence of the relaxation time).

In addition to the mean first passage time, the distribution of first passage times was computed by taking a histogram over the results of the stochastic simulation. Figure 5.7(a) shows the first passage time obtained by the analytic equation (equation (5.2)) and empirically. In the case of high energy barriers $\sigma \approx 8$ it can be seen that the results of the LLG simulation are also exponentially distributed with the same rate $\tau = x$. Therefore the thermal activation model not only accurately represents the mean relaxation rate of a single magnetic nanoparticle but the jump process is also an accurate representation




Figure 5.6: Mean first passage time for a Stoner-Wohlfarth particle. Comparison of Langevin stochastic simulation and the analytic Néel-Brown formula shows good agreement in the intermediate regime. $\sigma = 8.8$, $\mathbf{h} = 0$, $\hat{\mathbf{u}} = [0 \ 0 \ 1]$, $\hat{\mathbf{m}}_0 = [0 \ 0 \ 1]$.

of the true dynamics for high energy barriers. The simulation was repeated with a lower energy barrier, and the results reveal that the empirical distribution diverges from the true solution. Eventually, the distributions become qualitatively different. For very low σ , Figure 5.7 shows that an inverse Gaussian distribution fits the FPTD closely. The inverse Gaussian distribution is associated with the first passage time distribution for a Wiener process with drift (i.e. a random Brownian walk) [185]. This result supports an intuitive description for the condition of high energy barriers: at low noise the dynamics are dominated by the jump process due to the relatively strong anisotropy leading to a bi-stable energy potential; the thermal effects may be neglected. But at high noise, the thermal forces are so strong that the structure of the energy landscape is hidden by the random noise, leading to a purely random walk of the magnetic moment. In between the limits of a jump process $\sigma \gg 1$ and random noise $\sigma < 1$, it was difficult to determine whether the master equation provided a reasonable approximation.

The FPTD was computed for a range of σ and α using Magpy and compared to the analytic FPTD using the KS-statistic. Figure 5.8 the value of the KS for the simulations. The results show that as the energy barrier decreases, the two distributions diverge. The results show a non-trivial dependence on the two parameters. The figure also shows a decision boundary, which corresponds to 90% confidence. The phase plot can be used to determine whether a given magnetic nanoparticle can be accurately simulated using the thermal activation model. If the material parameters give a σ, α that is below the line then the thermal activation model may be used.

The results of the FPTD experiments showed that, under the correct conditions, the thermal activation model provides a good approximation of the full sLLG dynamics. Moreover, the KS-statistic provided a meaningful method to determine the validity of the model for varying sets of material properties.

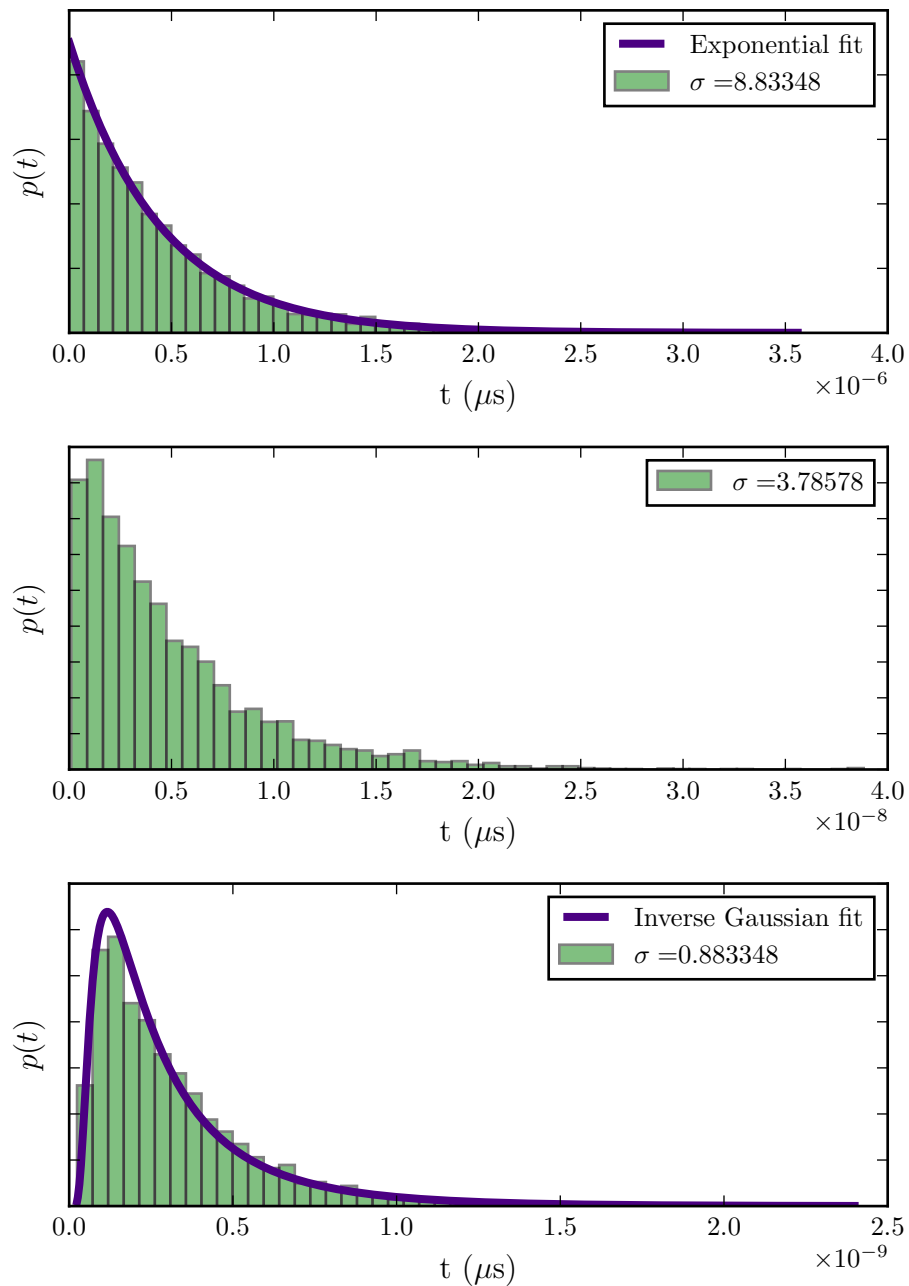


Figure 5.7: The empirical distribution of first passage times from Langevin simulations with decreasing stability ratio σ . Parameters as in Figure 5.6 with $\alpha = 0.1$. The empirical histograms are fitted with parameterised distribution functions by maximising the likelihood.




Figure 5.8: The KS distance measures the difference between the first passage time distribution of the stochastic simulations and the analytic Néel-Brown result. A lower KS distance is desirable. Parameters as in Figure 5.6.

Chapter 6

Nanoparticle cluster geometries and heating effects

Interactions between magnetic nanoparticles alter their behaviour and have the potential to increase or decrease heating output in magnetic hyperthermia experiments. In particular, dense clusters of aggregated particles show qualitatively different behaviour to non-interacting, dispersed nanoparticles. In this chapter, we simulate small clusters of magnetic nanoparticles subjected to alternating fields. The clusters are represented by a system of macrospins coupled by dipole-dipole interactions. We first investigate the dynamic, major hysteresis loops of particle chains and dense clusters. The results show that dense clusters dissipate substantially lower heat due to being magnetically soft. Our results corroborate previous conclusions from kinetic Monte-Carlo experiments in Section 3.3. However, we further show that chains show much greater variation in SPL with orientation to the magnetic field and the strength of the applied field. The results suggest that although chains can reach greater values of SPL they could lead to greater variation of heating in magnetic hyperthermia experiments.

6.1 System definition

Interacting clusters of magnetic nanoparticles were modelled by coupled stochastic Landau-Lifshitz-Gilbert equations and simulated using Magpy. The nanoparticles forming each cluster were identical spheres to isolate the effects of the cluster shape. Unless specified, the parameters used were: anisotropy strength $K = 5000\text{Jm}^{-3}$, particle radius $r = 10\text{nm}$, saturating magnetisation $M_s = 480\text{kAm}^{-1}$, damping $\alpha = 0.1$, density $\rho_{\text{Fe}} = 5.2 \times 10^6\text{gm}^{-3}$, particle spacing (measured from particle midpoints) $R = 25\text{nm}$, temperature $T = 300\text{K}$, alternating field frequency $f = 500\text{kHz}$. The dynamic hysteresis loops of ensembles of clusters were simulated to determine the average heat dissipated.

The simulated system was an ensemble of 50 clusters of N identical magnetic nanoparticles. Inter-cluster interactions were neglected. Unless specified, each member of the ensemble contained three randomly generated properties

1. The constituent particles of each member of the ensemble had a randomly oriented uniaxial anisotropy axis.
2. The geometry (relative positions) of the particles within each member cluster was often identical for all members in the ensemble. However, some geometries have a number of small variations with very similar energies (see Section 6.1.1). In this case, the geometries of each member were randomly chosen.
3. Each member cluster was randomly oriented.

The benefit of randomly generating these properties is that it effectively *marginalises* over the three properties, which allows us to evaluate the effects of size and shape without considering the impact of particle and cluster orientation. We explore the effect of cluster orientation in Section 6.5.2.

6.1.1 Cluster geometries

The simulated clusters had two distinct geometry types:

Particle chains were defined by a chain axis, number of particles N , and the particle-to-particle spacing R (as measured by from the midpoints of the adjacent particles). The particles were placed in a perfect line along the chain axis.

Dense clusters of nanoparticles were defined using the geometries provided in [7], which we will refer to as Arkus clusters. Arkus clusters describe minimally rigid packings of hard spheres subjected to attractive forces, which were obtained analytically by graph theory and provide a useful reference for realistic nanoparticle packings. The Arkus clusters are defined by the number of particles N , particle spacing R , and a unique id since clusters with 5 or more particles have more than one minimal energy packing (see Appendix C for Arkus cluster definitions).

Each ensemble was composed of members of identical shape and number of particles N . However, randomness was introduced into each ensemble by varying the orientation and anisotropy axes of each member in the ensemble. In the case of Arkus clusters larger than 5 particles, each member of the ensemble was drawn from each of possible N -particle packings (the packing id was chosen from a uniform multinomial distribution).

6.1.2 Random anisotropy axes

The anisotropy axes of the particles in each cluster in the ensemble were generated with random orientations using the sphere point picking method, which generates a uniform distribution of points over the surface of a sphere [199]. First, two independent and identically distributed random numbers u_1, u_2 are drawn from a uniform distribution between 0 and 1 $\mathcal{U}[0, 1]$. These numbers generate the azimuth and elevation angle of a point on the sphere:

$$\begin{aligned}\phi &= 2\pi u_1 \\ \theta &= \arccos(1 - 2u_2)\end{aligned}\tag{6.1}$$

where ϕ, θ are the azimuth and elevation angle respectively such that $x = \sin \theta \cos \phi, y = \sin \theta \sin \phi, z = \cos \theta$. A new random point on the sphere can be generated by redrawing u_1 and u_2 .

6.1.3 Random cluster orientation

In addition, the orientation of each cluster in the ensemble was randomly generated. Any rotation in three dimensional space can be represented by a single rotation angle about a rotation axis [111]. The angle and rotation axis are represented as a quaternion. Random quaternions were generated to obtain a uniform distribution of rotations in three dimensional space. A random quaternion was generated as follows:

1. Randomly set the quaternion axis using the uniform sphere point picking algorithm (equation (6.1))
2. Choose a random rotation angle uniformly on $[0, 2\pi]$

6.2 Heat dissipation in the major-loop regime

The literature review in Section 3.3 showed that the effects of interactions depend on whether the particle clusters are in a major- or minor-loop regime. Comparing two systems in different regimes leads to ambiguous and conflicting results. The amplitude of the applied magnetic field was set to $H = 3H_k$ where the anisotropy field $H_k = 2K/(\mu_0 M_s)$ is equivalent to the coercive field for single non-interacting particles. The large field was sufficient to fully reverse clusters in the major-loop regime and therefore simplify interpretation of the results. The heat dissipated in the major-loop regime corresponds to the maximum achievable SPL for a given magnetic system and applied field frequency.

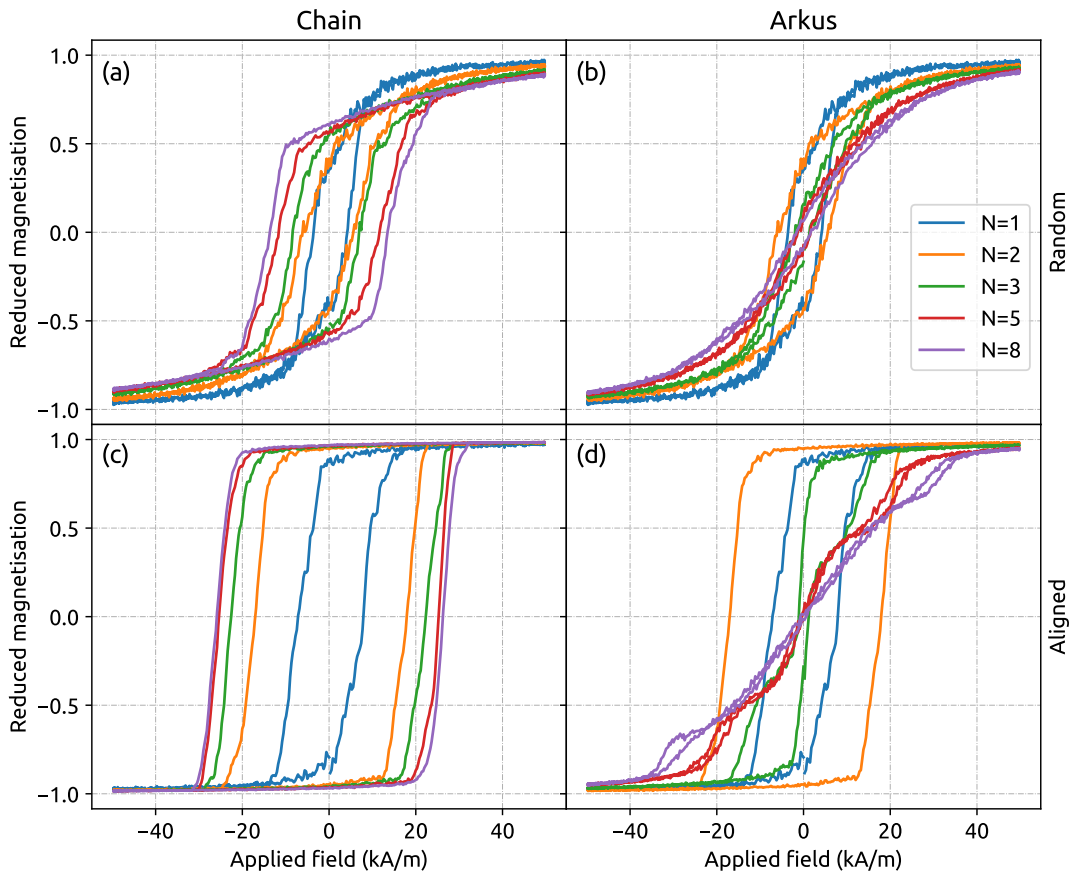


Figure 6.1: Dynamic major hysteresis loops for chain and Arkus cluster geometries with 1-8 particles. Results were obtained by stochastic simulation of coupled sLLG equations using Magpy. The vertical and horizontal axes are shared between all figures. Results are shown for randomly oriented clusters with random anisotropy axes and clusters aligned with the applied field with aligned anisotropy axes.

Chain and Arkus clusters of increasing size were simulated to compare their dynamic hysteresis major-loops and hence heat dissipation. The magnetisation of each particle in the ensemble was randomly initialised with equation (6.1) and simulated for five field cycles (i.e. $5/f$ where $f = 1/T$ is the applied field frequency in Hz and T is the period in seconds) to allow for the system to reach periodic equilibrium such that $M(t + T) = M(t)$ forming a closed loop in the $M - H$ plane.

6.2.1 Randomly oriented clusters

Figure 6.1 shows the dynamic hysteresis loops for the final cycle of the simulation for each ensemble. We can identify the hysteresis loops as major loops by recognising that the applied field value is greater than the field at which the hysteretic behaviour vanishes. Comparing the chain and Arkus clusters in the random case, we find that their hysteresis loops possess qualitatively different shapes (excluding one and two particle clusters, which




Figure 6.2: Dynamic hysteresis loop quantities and SPL of randomly-oriented ensembles of chain of increasing length N . The SPL is computed by numerical integration of the hysteresis loops in figure 6.1(a).




Figure 6.3: Dynamic hysteresis loop quantities and SPL of randomly-oriented ensembles of Arkus clusters of increasing size N . The SPL is computed by numerical integration of the hysteresis loops in Figure 6.1(b).

have identical geometries). Since all material, environmental, and field conditions are equal in both experiments, these differences must be attributed to the geometry of the clusters alone.

The specific power loss (SPL) values of the simulated loops in Figure 6.1 were obtained by numerical integration (see equation (2.50)) using a trapezoidal scheme implemented in the python package *numpy* [155]. Four additional quantities were extracted numerically from the hysteresis loops: the zero-magnetisation dynamic coercive field H_c , zero-field remanent magnetisation M_r , maximum magnetisation $M(H_{\max})$, and squareness M_r/M_{\max} . The hysteresis loop quantities for randomly-oriented chains and Arkus clusters of increasing size are shown in Figures 6.2 and 6.3 respectively.

The randomly oriented chains in Figure 6.1(a) show a larger dynamic coercive field H_c , remanent magnetisation M_r , and squareness (M_r/M_{\max}) compared to the randomly oriented Arkus clusters, which leads to a substantially larger hysteresis area hence SPL. The increasing coercive field suggests that as the chain length increases, larger field amplitudes are required to reverse the chains. The increasing M_r shows that as the chain length grows, the particles are more likely to remain aligned as the field reverses.

Figure 6.2 shows that the SPL monotonically increases with increasing chain length. Increasing the chain length above $N = 3$ leads to diminishing returns; the dipolar interaction field decreases with the cube of the distance between particles (see equation (2.4)).

The Arkus clusters in Figure 6.1(b) show close to superparamagnetic behaviour: leaning loops with low squareness and narrow hysteresis area. In other words, the Arkus clusters are magnetically soft. This agrees with physical intuition that the dipolar interactions cause the magnetic moments to favour end-to-end alignment, which consequently leads to closed flux paths with very low total energy. In contrast to the chains, the Arkus clusters show a non-monotonic dependence of the SPL on the size of the cluster. After $N = 1$ and $N = 2$, which are equivalent to the chain, the symmetry of the cluster shape is broken and the SPL rapidly reduces. The plot of H_c reveals that the Arkus clusters are substantially easier to reverse.

Both the randomly-oriented chains and Arkus clusters have a similar dependence of the maximum magnetisation M_{\max} , which is a consequence of thermal fluctuations causing particles to oppose the applied field even for very strong field amplitudes. The results suggest that in the randomly-oriented ensembles there are relatively low energy barriers that correspond with relaxation times on the same order of the period of the applied field $T = 1/f$.

The gradual decrease of the system magnetisation with the field at the ends of the hysteresis loops suggests a broad range of switching times (time taken for individual particles to reverse direction against the field), which arises from both:

1. *The system energy barrier distribution:* different configurations of particles within the clusters that are separated by lower energy barriers (either by individual or collective rotation of the magnetic moments) require a smaller applied field to transition. The distribution of energy barriers is broadened by the different orientations of the clusters and their constituent particle anisotropies.
2. *Thermal fluctuations:* higher temperature leads to larger thermal fluctuations, increasing the probability of particles spontaneously activating over the energy barriers.

The low squareness of the Arkus clusters therefore suggests that the system has lower and/or more dispersed energy barriers. On the other hand, the hysteresis loops for chains of length $N = 5, 8$ show inflection points at which a large proportion of particles switch simultaneously, suggesting a narrow distribution of energy barriers much greater than $k_B T$, which dominate the reversal process. In appendix D we show results from repeating the experiment at lower temperature $T = 30\text{K}$. The resulting hysteresis loops have a similar shape, which suggests that the distribution of switching times is dominated by the random orientations of clusters and particles rather than by thermal activation.




Figure 6.4: Arkus clusters of anisotropy-aligned particles. All anisotropy directions (red arrows) are aligned in the same direction as the alternating magnetic field. The figure shows rendered clusters of size 1-6. In the aligned simulation, the relative positions of the particles are also fixed in the arrangements shown. The longitudinal axis of cluster (2) and the plane spanned by the particles in (3) are parallel with the applied field.

6.2.2 Anisotropy-aligned clusters

Figure 6.1 shows the hysteresis loops obtained from aligned ensembles, which consist of 50 shapes oriented in the same direction with their anisotropy axes aligned with the external field. Examples of aligned Arkus clusters of size $N = 1 - 6$ are shown in Figure 6.4.

The aligned chains in Figure 6.1(c) loops show a larger coercivity and squareness for all chain lengths compared with the randomly-oriented ensembles. The square loop shape is similar to the response of a single Stoner-Wohlfarth particle, suggesting a single large energy barrier substantially greater than the thermal energy. Overcoming this large energy barrier requires a large reversal field, which leads to a substantial dissipation in heat.

In Figure 6.1(d), the simulated Arkus cluster loops present a complex shape. The $N = 3$ loop is narrow but has high squareness resulting from a large energy barrier to reverse the first particle. The closed $N = 5$ and $N = 8$ loops show distinct inflection points that suggest multiple meta-stable states separated by energy barriers. Determining the configuration of the particles within these meta-stable states requires further investigation of the dynamics of the individual magnetic moments. However, one possible interpretation is that the particles in the clusters tend to switch individually and consecutively. This contrasts with the chains in which all particles rotate simultaneously as evidenced by the near-instantaneous switch in the magnetisation. The aligned $N = 5$ and $N = 8$ loops also show a small hysteresis area as in the randomly-oriented ensembles. This corroborates with the structure of the Arkus clusters in Figure 6.4, which shows that clusters with $N > 3$ particles possess approximate rotational symmetry and therefore will respond consistently to alternating external fields in different orientations.

Comparing the tails of the chain and $N < 5$ Arkus loops to the randomly oriented case reveals that the aligned clusters have flat tails and the randomly-oriented clusters have sloping tails. The flat tails indicate no switching in the initial reversal of the alternating field. In the randomly-oriented case, the broader distribution of switching must be attributed to either the distribution particle anisotropies or cluster geometry rotations. For example: the decreased coercivity in randomly-oriented chains agrees with physical intuition that chains that are perpendicular to the applied field are easier to reverse than those aligned with the field. That the orientation of the clusters substantially affect the hysteresis loop corroborates with the $T = 30K$ simulations in Appendix D, which demonstrated that even with small thermal fluctuations the tails are not flat as in the aligned case.

6.2.3 Discussion

The simulated results of subjecting clusters to large enough fields to activate major-loop dynamic hysteresis agree with previous experimental and computational results comparing chains and dense clusters of particles. The current literature (see Chapter 3 or the recent review [1]) shows that many factors affect the relationship between cluster's geometry/size and resulting SPL [159]. However, chains of particles generally show increased SPL while clusters of particles show reduced SPL. In addition to dense clusters and chains, Serantes et. al. [181] simulated ring, cube, and hexagonal structures that all showed reduce SPL with increasing size compared with chains. The results showed that the SPL correlated with the symmetry of the nanoparticle structures: aligned linear aggregates (one-dimensional) and 3 particle clusters (two-dimensional) showed higher SPL compared to all three-dimensional shapes. Our results also agree exactly with [64], which describes these clusters as having a larger effective anisotropy leading to larger SPL (assuming clusters are exhibiting major-loop hysteresis).

In existing studies with minor-loop activated clusters, the detrimental effect of dense clustering on SPL is also seen [159, 192]. In [159] it is shown that larger clusters of particles have lower SPL due to demagnetising effects but in contrast to our major-loop simulations they also show increased dynamic H_c with increasing cluster size.

Our results for small $N < 9$ perfectly aligned chains also match the behaviour of large approximate chain aggregates simulated at low temperature [181]. In particular, the coercive field, remanent magnetisation and squareness increase with increasing chain size, which leads to an increasing SPL with diminishing returns as the hysteresis loop approaches a square shape. Moreover we found this trend to be true for both perfectly aligned chains and completely randomly-oriented ensembles. It is shown in [181] that $N = 3$ chains perform worse than $N = 2$ chains when oriented perpendicular to the applied field. Thus the gain from increasing the lengths of parallel chains dominates the losses from increasing the lengths of perpendicular chains in a random ensemble.

However, aligned chains always dissipate higher SPL than randomly-oriented chains in the major-loop regime.

6.3 Heat dissipation in the minor-loop regime

In the following simulations we investigate the heat dissipated under smaller applied field amplitudes, while keeping the field frequency constant. As the alternating field amplitude is reduced, the hysteretic behaviour enters the minor-loop regime. In this regime, the applied field amplitude is not sufficient to fully activate the particle cluster and the height of hysteresis loop becomes a function of the applied field amplitude, frequency, and the relaxation time of the system (see Section 3.3).

6.3.1 Effect of increasing field amplitude

Figure 6.5 shows SPL (obtained as in Section 6.2) as a function of the reduced applied field amplitude H/H_k for the chain and Arkus geometries in both random and aligned orientations. The SPL increases monotonically with the applied field amplitude for all cluster geometries. Increasing the amplitude deterministically broadens the hysteresis loop (i.e. along the horizontal H axis). If the larger field does not lead to a response with reduced remanent magnetisation (i.e. along the vertical M axis) then the hysteresis area hence SPL will increase. Thus we find that the remanent magnetisation does not decrease substantially in any of the clusters.

As the field amplitude is increased we expect SPL to saturate as the hysteresis area approaches its theoretical limit $H = 4\mu_0 M_s H_c$. This can be seen in previous kinetic Monte-Carlo simulations at very low temperature [181] and in Appendix D where we repeat the experiments at $T = 30\text{K}$. However, at $T = 300\text{K}$ the simulations show that SPL can still be increased with diminishing returns at field amplitudes much larger than H_k . Figure 6.6 shows the dynamic hysteresis response of an ensemble of anisotropy-aligned single particles (equivalent to Arkus and chain clusters with $N = 1$). Although the ensemble displays major-loops for $H \geq 0.5$ the SPL is increased by further increasing the applied field amplitude, which broadens the hysteresis area. We expect that as large fields increase above $H = 3H_k$ that the SPL will eventually saturate. The field required to reach the maximum SPL for a given shape is therefore a function of the temperature of the system.

6.3.2 Effect of increasing cluster size

Figure 6.2 showed SPL increasing monotonically with the length of the particle chains in the ensemble. However, Figure 6.5(a)(b) shows that for $H < 1.5H_k$ longer particle

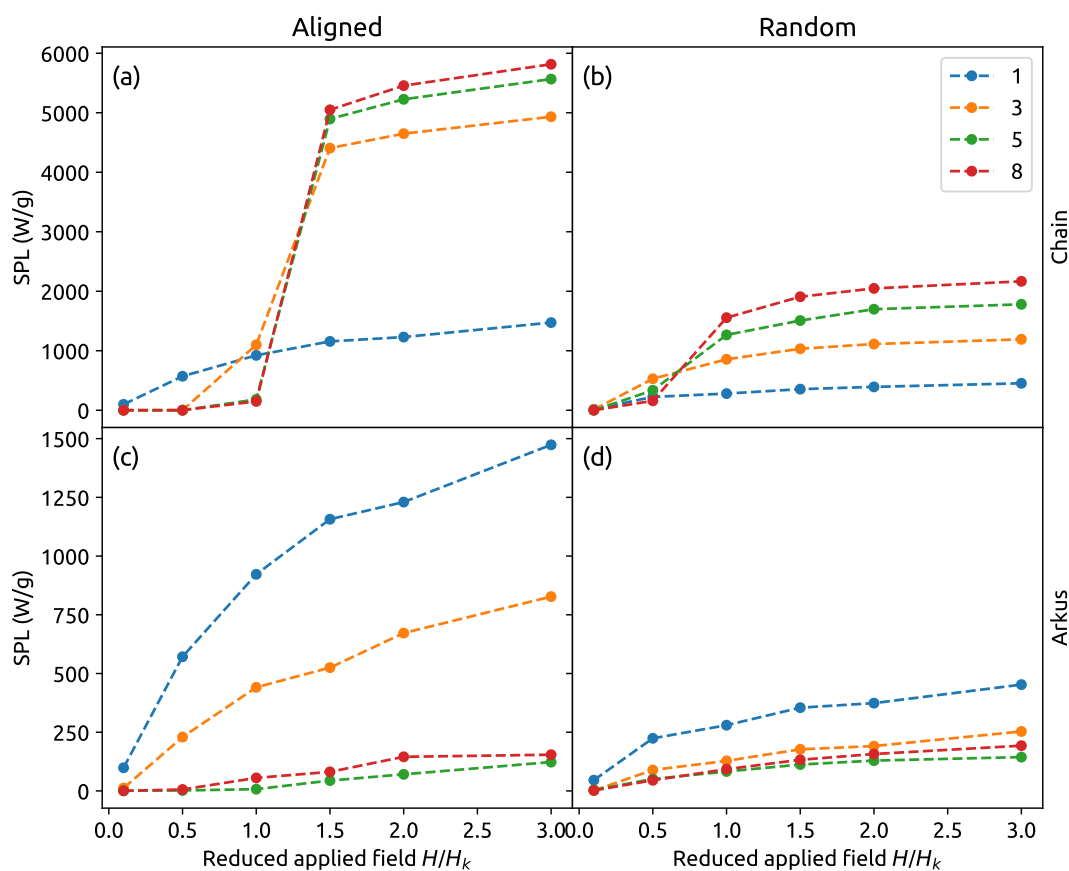


Figure 6.5: SPL obtained from numerical integration of hysteresis loops. The hysteresis loops were simulated for Chain and Arkus geometries of various sizes both randomly-oriented and aligned with the externally applied field. The results are shown as a function of the alternating magnetic field reduced amplitude (normalised by the anisotropy field H_k for a single particle).

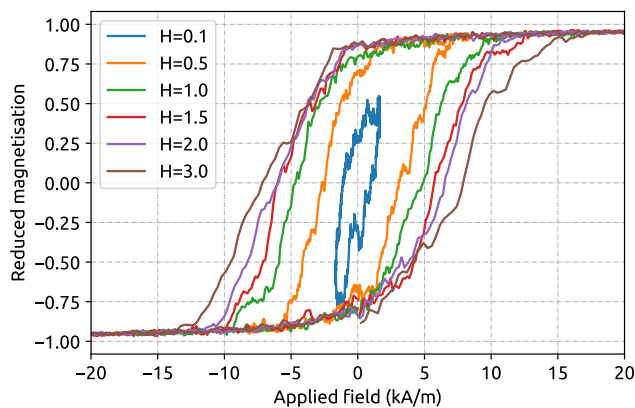


Figure 6.6: Simulated dynamic hysteresis loops for an ensemble of aligned non-interacting nanoparticles. The amplitude of the alternating magnetic field was increased from $0.1H_k$ to $3.0H_k$.


chains have lower SPL than shorter chains. This is also true for large Arkus clusters; $N = 8$ and $N = 5$ clusters have very low heating output in Figures 6.5(c)(d). The heat dissipated under a $H = 0.5H_k$ field is plotted again in Figure 6.7a as a function of cluster size for randomly-oriented Arkus clusters and chains, which shows a similar trend for both cluster shapes. The relationship between H and SPL for chains no longer monotonically increases (as in Figure 6.2) whereas Arkus clusters behave similarly (as in Figure 6.3).

We investigated the origins of this non-monotonic relationship by inspecting the hysteresis loops of the two shapes at a weaker field in Figure 6.7b. The figure shows the evolution of the hysteresis area from $N = 2$ in subfigure (i) (identical for both chain and Arkus clusters) to $N = 5$ in subfigures (ii) and (iii) for chains and Arkus clusters respectively. The loops are qualitatively different and suggests different physical origins of the reduced SPL. For Arkus clusters (ii), as in the case of $H = 3H_k$ in Figure 6.1, the loop is fully activated but close to superparamagnetic. This is due to the low energy barriers, which are easy to overcome and lead to very fast relaxation times. On the other hand, the $N = 5$ chains (iii) are no longer fully activated and have reduced magnetisation. The field is not sufficient to reverse the large chains leading to minor hysteresis loops, as seen in low field kMC simulations in [181]. It has been shown previously that longer chains have longer switching times [124]. The blocked state of the particle chains can be seen in Figure 6.5(a) which shows almost no heat output followed by a large change around $H = H_k$, which is sufficient to fully reverse the particle chains. In contrast, the Arkus clusters show an SPL as a smooth function of H .

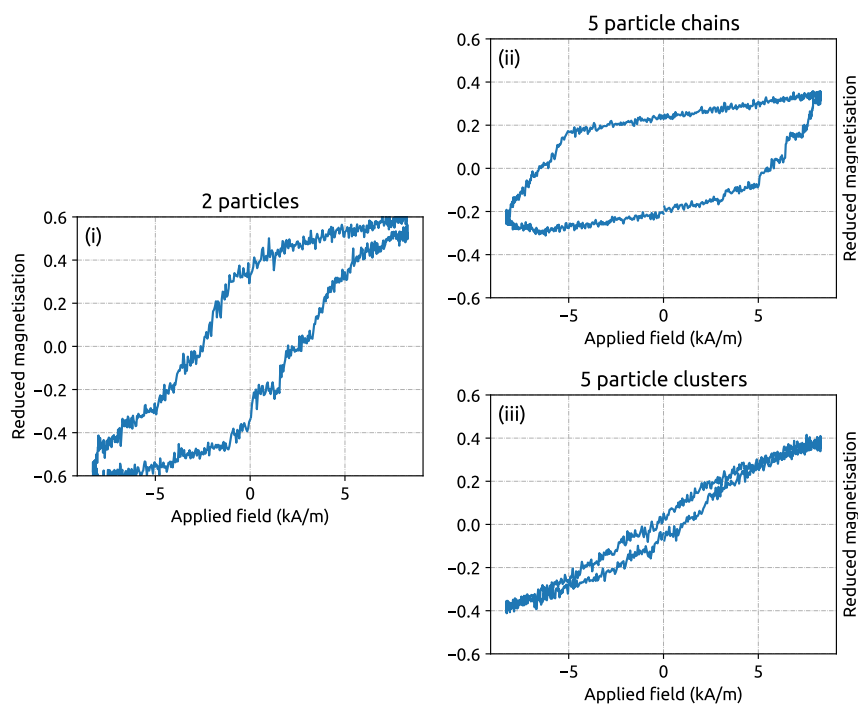
These results revealed that, at low fields, clusters of various shapes can show decreased SPL with increasing size of the cluster. However, the origins of the decreased heating output varies for chains and densely-packed clusters. Particle chains exhibit increased energy barriers that must be overcome; SPL can be increased by increasing the thermal fluctuations, reducing the coercivity of the chains or increasing the applied field amplitude. Particle clusters exhibit decreased energy barriers leading to a paramagnetic-like response; SPL can be increased by reducing thermal fluctuations, increasing the relaxation time of the system or increasing the applied field frequency.

6.3.3 Discussion

The relationship between cluster shape and SPL is related closely to the dimensionality and symmetry of the clusters. In [181] many shapes were simulated and the results suggested that 1-dimensional clusters dissipated more than 2-dimensional and 3-dimensional. Two other recent works [96, 192] used the fractal dimension to describe the geometry of clusters and showed a trend of decreasing SPL with increasing fractal dimension D_f . Our simulations corroborate with this trend and can be seen in Figure 6.5(c)(d) where $N = 1$ is 1D, $N = 3$ is 2D and $N = 5, 8$ are 3D. Furthermore, we have



(a) Results from Figure 6.5 are re-plotted for a single applied field amplitude $H = 0.5H_k$. The SPL is plotted as a function of increasing cluster size for ensembles of randomly-oriented particle chains and Arkus clusters.



(b) Hysteresis loops for three different ensembles of cluster types subjected to a low amplitude $H = 0.5H_K$ field. (i) an ensemble of two particle clusters (identical for chains and Arkus clusters) showing major-loop activation (ii) an ensemble of 5 particle chains showing minor-loop activation (iii) an ensemble of 8 particle Arkus clusters showing major-loop activation close to superparamagnetic behaviour.

Figure 6.7: Heat dissipation (SPL) resulting from subjecting Arkus clusters and chains of increasing size to an alternating magnetic field of low amplitude $H = 0.5H_k$ where H_k is the anisotropy field of the constituent particles. In all figures, the particles and clusters are randomly oriented.

observed that this rule-of-thumb remains true for both aligned and randomly-oriented geometries in the major- and minor-loop regimes. The distribution of SPL with cluster orientation is explored further in Section 6.5.

Long chains of anisotropy-aligned particles achieve the maximum possible SPL, if there are no restrictions on the applied magnetic field amplitude. Specifically, the field amplitude H must be larger than H_k . This corroborates with numerous previous numerical and experimental (see e.g. [181]) studies. However, below the anisotropy field, large chains show SPL substantially lower than a single particle whereas clusters show a smoother function of SPL with the applied field. These discontinuities may cause problems: if the applied field is around H_k then some of the chains may never activate while others may be dissipating substantially. In Section 6.5 we explore the impact of variation of heating between particle clusters in detail.

In all simulations the $N > 3$ Arkus clusters showed a reduced SPL compared with a single, isolated particle. This was true for randomly-oriented and anisotropy-aligned particles across a range of field values. Hysteresis loops revealed that this reduced SPL was due to paramagnetic behaviour leading to small hysteresis areas. This has been observed in experiments [64, 159] as well as in simulated systems [181, 192]. However, there exists examples [159, 181] of clusters improving SPL versus a single particle in cases where the applied field is very low (thus not activating clusters) or the material parameters are different. Furthermore, a study of large, dense clusters (3000 particles placed randomly in a sphere with particle-to-particle spacing of 25nm) showed a low resultant SPL due to increased coercivity leading to minor-loop behaviour (i.e. as in Figure 6.7b(ii)) rather than because of paramagnetic behaviour (i.e. as in Figure 6.7b(iii)). The results in [189] show that particles in the middle of the clusters behave differently to particles at the edge, which could account for this unexplained difference in behaviour. Overall, the comparison with existing results suggests that SPL in dense clusters is strongly dependent on a large number of system properties. In Section 6.4 we replicate the results of a previous experiment [152] to investigate a specific example of $N = 3$ Arkus clusters improving the SPL of single particles.

6.4 Comparison with experimental data

The previous section showed that, for the specified material properties, Arkus clusters showed lower SPL values than single, isolated particles. This result contrasts with a previous combined experimental and kinetic Monte-Carlo (kMC) investigation [152]. The kMC simulations showed that randomly-oriented $N = 3$ clusters (or trimers) dissipated more heat than a single particle. We attempted to reproduce the simulations in [152] by using the same material properties: $M_s = 450\text{kA/m}$, radius $r = 10\text{nm}$, $K = 50\text{kJ/m}^3$, particle spacing $R = 23\text{nm}$, $T = 300\text{K}$, field amplitude $H = 23.8\text{kA/m}$, field frequency




Figure 6.8: SPL of a single magnetic nanoparticle simulated with Magpy using varying values of the damping ratio α . High values of α lead to high longitudinal damping of the magnetic moment. The various values are compared to the SPL simulated using kinetic Monte-Carlo simulations (kMC) in [152]

$f = 300\text{kHz}$. The method to randomly generate the clusters and anisotropy orientations wasn't controlled but the differences were assumed to be negligible.

6.4.1 Tuning the damping ratio

Figure 6.8, shows the SPL predicted by Magpy simulations by numerical integration of the simulated hysteresis loops for a single magnetic nanoparticle. The value of the damping ratio α was varied and compared with the kMC result, which doesn't account for the phenomenological damping of the magnetic moment. The results show that the SPL, hence hysteresis area, is highly sensitive to the damping ratio. The results of the two simulations match very closely for values of $\alpha \approx 0.1$, which is expected for the two-level approximation with Néel-Brown transition rates (as used in the kMC model) as shown in Section 5.7 and [22, 38].

6.4.2 Heat dissipation in dense clusters

We simulated $N = 1, 2, 3, 6$ Arkus clusters and compared them with the results in [152]. The definition for the clusters can be found in Appendix C. The value of K for all particles in the clusters was varied for all particles, which aims to fit the experimental results. Figure 6.9 compares the predicted SPL using Magpy to the result obtained from kMC simulations, where the dimensionless stability ratio $KV/(k_B T)$ was deliberately varied by adjusting the anisotropy constant K .

For the $N = 1$ clusters (single particles), we see that the predicted SPL matches very closely. However, as interactions are introduced, the predicted SPLs vary greatly between the two approaches. Excluding $\sigma = 20.2$, the resulting SPL from Magpy simulations was smaller than that predicted by kMC simulations. In particular, $N = 2$ clusters (dimer chains) are predicted to have a large increase in SPL compared to the $N = 1$

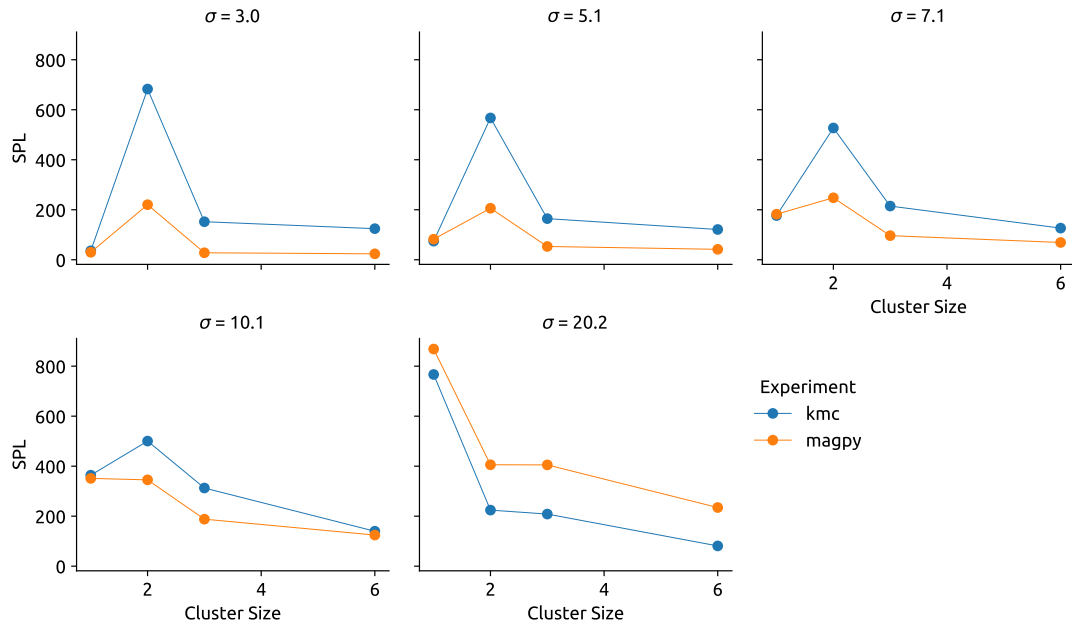


Figure 6.9: SPL values obtained from kinetic Monte Carlo and Magpy simulations of dense clusters of particles. The anisotropy value of the particles was changed in each simulation, which is represented by the dimensionless stability ratio $\sigma = KV/(k_B T)$

cluster. However, although there is a noticeable increase in the SPL for the magpy simulated data, the incremental increase is lower. Furthermore, magpy simulations do not produce SPL values for $N = 3, 6$ clusters that are greater than $N = 1$.

We investigated these differences by inspecting the simulated dynamic hysteresis loops from Magpy and kMC data. The latter obtained from Niculaes et.al. [152]. The $N = 1$ loop looks very similar for both and is a typical shape for a randomly oriented nanoparticle. However $N = 2$ shows a drastically different loop shape to the kMC simulations. In the kMC, the chains are exhibiting minor hysteresis loop due to not being fully activated. This suggests that the clusters show increased coercivity (as expected for chains, see previous section). Moreover, the increased remanent magnetisation suggests a large effective anisotropy and therefore increased thermal stability. Additionally, $N = 3$ and $N = 6$ exhibit minor hysteresis loops in the kMC simulations. Thus the dipolar interactions in the Arkus increase the coercivity compared with single particles. This result agrees with [189] but contradicts with our simulations, which consistently show a reduced coercivity with increasing cluster size. Both the kMC and magpy simulations use the same material properties, therefore the different behaviour must be explained by the method to generate geometries or the simulated dynamics of the system.

Section 6.1 described the generation of uniform distributions of the anisotropy directions and cluster orientations. The clusters used in the kMC simulations were created using a fractal generating algorithm as described in [96] and the anisotropy directions were also uniformly distributed over a sphere. Therefore, it is possible that the cluster shapes have

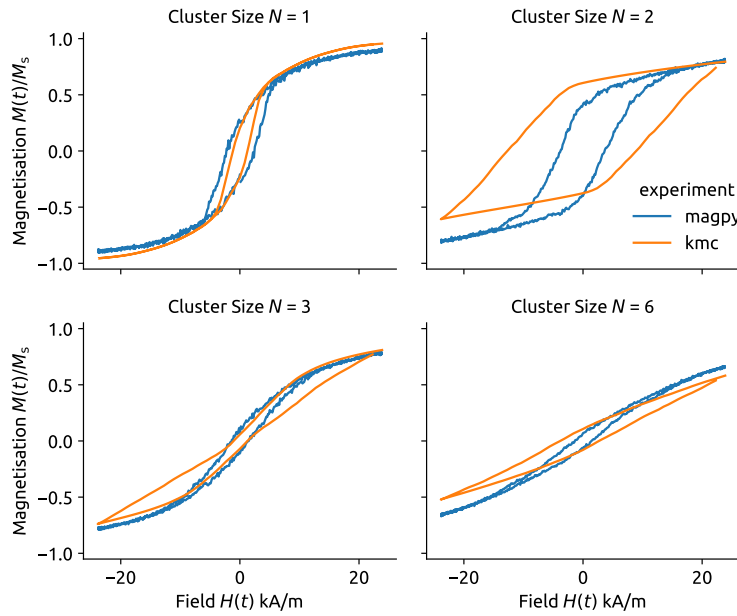


Figure 6.10: Dynamic hysteresis loops simulated using magpy for $\alpha = 0.1$ and $\sigma = 5.1$ ($K = 50\text{kJ/m}$) compared with simulated loops obtained by kinetic Monte Carlo simulation.

impacted the results. In particular, if the fractal generating algorithm is biased towards shapes that align with the external field, the SPL measurements would be larger. The behaviour of these two algorithms should be investigated further. Alternatively, cluster orientations and shapes could be determined from observing particle clustering behaviour in experiments.

A fundamental difference between the Magpy and kMC simulations is the underlying model of the magnetic moment. Magpy simulates the full stochastic Landau-Lifshitz-Gilbert equation whereas the kMC simulations use the Stoner-Wohlfarth model of a two-state magnetic moment such that a system of N particles has 2^N possible configurations. As discussed in Section 2.4.3.1, the energy barriers between these states must be sufficiently large for the approximation to be valid. Given that the anisotropies are randomly varied, it is possible for clusters to be oriented in energy efficient orientations with very low energy barriers between states. These configurations could lead to unrealistic transition rates in the kMC simulations. Indeed, previous studies using kMC simulations have required unphysical anisotropy constants and saturation magnetisation values in order to match experimentally observed results [159].

6.4.3 Heat dissipation in mixed solutions

In the original experiment, a K-value of $K = 50\text{kJ/m}$ was selected from the results of Figure 6.9 in order to explain experimentally observed SPL values in synthesised




Figure 6.11: SPL prefactors for each cluster size for three different solutions synthesised in [152] using different amounts of polymer in order to control the degree of clustering in each solution.

solutions of magnetic nanoparticles. Niculaes et.al. were able to control the degree of clustering in order to synthesise three different solutions. The clustering was controlled by varying the amount of polymer (polystyrene-co-maleic anhydride) used during the particle cluster synthesis process. The three solutions were categorised by the molecules of polymer per nanometer area of magnetic particle surface. The clusters are termed 16.5PScMA (16.5 molecules of polymer per nm^2), 33PScMA, and 66PScMA. The size distribution of the clusters was recorded (e.g. 49% dimers, 20% trimers, etc.) using transmission electron microscopy (TEM). Let n_k be the proportion of clusters in the solution that consist of k -particles. Assuming all particles are the same size and density, we can compute the proportion of particles in the solution that are attached to each of the cluster types (sizes) k :

$$\alpha_k = \frac{kn_k}{\sum_k kn_k} \quad (6.2)$$

α_k should be interpreted as the proportion of particles that belong to a cluster of size k . The α_k values for three solutions obtained in [152], are shown in Figure 6.11. The total SPL dissipated in a given solution may be computed as a weighted sum of the SPL obtained by simulation of the individual clusters. Thus the SPL for solution X is:

$$\text{SPL}_X = \sum_k \alpha_k \text{SPL}_k \quad (6.3)$$

The total heat dissipated in the three solutions was obtained by applying the data from Figure 6.9 to equation (6.3). Figure 6.12 shows the total SPL dissipated in each of the clusters as predicted by the two simulations. Each bar is decomposed into the SPL dissipated by each of the cluster sizes. For example, in Figure 6.11 the 66 solution




Figure 6.12: Cumulative SPL dissipated by solutions of mixed clusters as described in Figure 6.11. The clusters were simulated with $\alpha = 0.1$ and $\sigma = 5.1$. The simulated SPL from the kMC method were a close match to experimental values but the magpy simulations show substantially lower SPL.

contains mostly $N = 6$ clusters and therefore the total SPL in Figure 6.12 consists almost entirely of contributions from $N = 6$ clusters

Figure 6.12 shows that the magpy simulations lead to a significantly lower SPL than obtained in experiments. Figure 6.9 showed that $N = 2$ clusters showed the largest difference in SPL between the two simulations and the 16 and 33 clusters are dominated by SPL contributions from the $N = 2$ cluster. One possible solution to match the data is to use an larger anisotropy value in the magpy simulations, which would increase SPL as shown in Figure 6.9. However, the $N = 3$ and $N = 6$ clusters would still dissipate less heat than a single particles, which we expect to alter the distribution of the SPL across the three mixed solutions in Figure 6.12. Alternatively, many of the material (M_s , α , V) and field (H , f) could be adjusted to find an alternative interpretation of the results.

6.4.4 Anisotropy-aligned clusters

In the simulations of the experiments in [152], it was assumed that the anisotropy orientations and the cluster orientations were randomly distributed. TEM images in [152] confirm that there is no ordering to the rotation of clusters but the anisotropy directions are more difficult to align. In [181], it was suggested that chains of particles tend to be aligned along their anisotropy axes. The anisotropy axes of individual particles in a cluster are difficult to measure. Therefore, we repeated the magpy simulations in Figure 6.9 but the clusters $N = 1, 2, 3, 6$ clusters were aligned as shown in Figure 6.4. Although the axes of the particles in each cluster were aligned, the orientation of the each cluster was still randomly rotated as described in Section 6.1.3.




Figure 6.13: Cumulative SPL dissipated by solutions of mixed clusters as described in Figure 6.11. The kMC simulations used randomly oriented clusters and anisotropy directions. The magpy simulations held the anisotropy axes aligned and randomly oriented the entire clusters. $\alpha = 0.1$ and $\sigma = 5.1$.

Figure 6.13 shows the results of the repeated experiment with aligned clusters. The $N = 2$ chains show much higher SPL as predicted, however the $N = 3$ clusters also show substantially larger SPL. This corroborates with the exploratory simulations of Arkus clusters in Section 6.2. Moreover, the $N = 6$ cluster still shows reduced SPL and inspection of the hysteresis loops show that the clusters are paramagnetic-like major-loops whereas the kMC predicts minor-loops. However, the results are closer to the experimental values and with some further tuning to the material parameters it is possible that the aligned anisotropy can explain the difference between the stochastic LLG simulations and experimental evidence.

6.5 Local heating distribution

In practice, magnetic nanoparticles are heterogeneous in their size, shape, clustering, magnetic properties etc. This heterogeneity not only affects the total heat dissipated by the system, as we have investigated in previous simulations, but also the distribution of heat dissipated in nanoparticles dispersed spatially. The local variation in heat is not discernible from the total SPL and temperature increase, which spatially average the effect. However the heating distribution is of crucial importance: if particles do not contribute to heat then the patient has been exposed to an unnecessary concentration of particles; if the heating is too high then damage could be done to healthy tissue.

The objective of minimising variation in heating output has only recently been articulated in the works of [146–148]. The relationship between polydispersity, particle stability, and particle interactions was explored in [146], which showed that increasing

interactions and polydispersity contributed to a wider variation in heating. In general, the effect of increasing interactions is to increase the distribution of energy barriers in the system [34]. However, the effect may be reduced by decreasing the stability of particles. It was shown again in [148] that choosing particles with lower anisotropy, ensuring all particles are in the major-loop regime, reduces the variation in heat dissipation in an anisotropy dispersed system. [147] presents an example of optimising particle properties, specifically the distribution of particle sizes, to obtain an SPL within a desired range specified by a minimum and maximum.

In this section, we investigate the variation the distribution of heating in the Arkus and Chain nanoparticle systems that were simulated in the previous sections. In particular, we analyse two forms of variation: the variation in heating between different cluster sizes in a mixed solution and the variation in heating from rotating the Arkus and chain shapes.

6.5.1 Variation of SPL between clusters

Consider two fictitious solutions each containing an equal weight of $N = 1, 2, 3, 5, 8$ randomly-oriented Arkus or chain clusters respectively. Figure 6.14 shows SPL as a function of the reduced field amplitude H/H_k for these solutions, which is a simple mean of the data in Figure 6.5. As expected, the SPL increases monotonically with the field and is considerably higher for chains. In comparison to Figure 6.5, the solution of chain ensembles shows a smoother function of the SPL compared to the individual shapes. This would be expected since we effectively introduce a range of coercive fields.

Figure 6.14 also shows the coefficient of variation (defined as the standard deviation divided by the mean) of each of these solutions. Note that the variation is only computed between the ensemble SPL of each of the cluster shapes, thus it does not account for the variation within each of the ensembles in Figure 6.5. The coefficient of variation shows that at low fields, the coefficient of variation tends to be larger. This is because a mixture of shapes have activated whereas a proportion of shapes output no SPL. The result is similar to the effect of disperse anisotropy values in an ensemble [148]. It is better to activate major-loops to decrease local heating variation.

6.5.2 SPL variation with shape orientation

The dynamic hysteresis loops of nanoparticle clusters in the major-loop regime showed a distribution of switching times, which can be attributed to various sources (Section 6.2). One of the sources of randomness comes from the fact that each cluster in the ensemble was randomly oriented. By comparing the variation between the different ensemble SPL values in Figure 6.13, we fail to account for this variation within each of the ensembles,


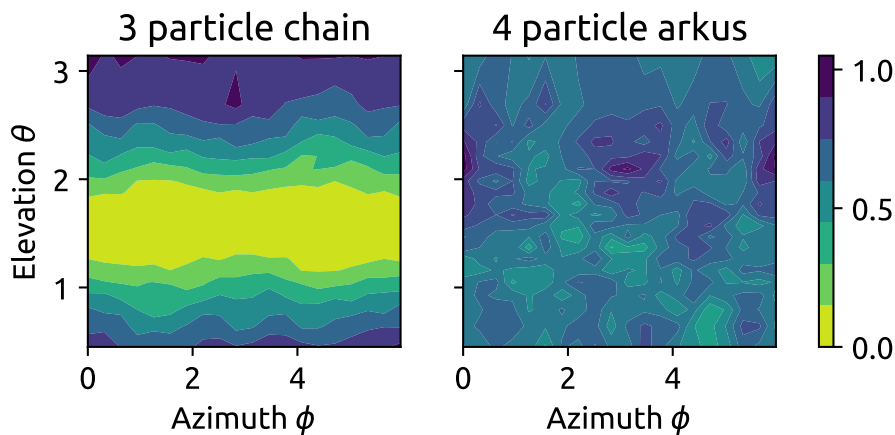


Figure 6.14: Total SPL and coefficient of variation for mixed solutions of Arkus and chain clusters of different sizes as a function of the applied field reduced amplitude H/H_k .


effectively integrating out those differences. This is useful for comparing the total heat dissipated in a large heterogeneous ensemble but not for understanding the internal distribution of SPL.

In order to understand the effects of the shape orientation on SPL, we generated ensembles of aligned $N = 4$ Arkus or $N = 3$ chain clusters with random anisotropy axes drawn for each particle. For each shape, multiple ensembles were created, with each ensemble aligned in a different direction. Using the sphere point picking algorithm, a uniform distribution of directions on the surface of a sphere was generated and used to orient the ensembles. By simulating multiple experiments of ensembles of aligned shapes in different orientations, we constructed a picture of the SPL for the $N = 4$ Arkus and $N = 3$ chain shapes as a function of their orientation with respect to the applied field direction.

Figure 6.15 shows the SPL dissipated by ensembles of aligned $N = 4$ Arkus or $N = 3$ chain clusters as a function of the orientation of the aligned ensemble. The three particle chain, due to rotational symmetry along its axis shows a rotational symmetry in the rotational SPL profile in Figure 6.15. The chain may be rotated about the applied field axis without any change in the heat dissipation. However, the angle between the chain axis and the applied field axis (the elevation angle θ) shows a very strong influence on the SPL. This result agrees with experiments of the relaxation of chains in different orientations [124]. Chains oriented with the field dissipate the maximum energy, whereas chains perpendicular to the field dissipate no SPL. In contrast the $N = 4$ Arkus cluster



(a) Distribution of SPL (normalised by the max SPL achievable) over a range of azimuth ϕ and elevation θ angles for the ensemble orientation vector.



(b) Three-dimensional surface plot with grid points corresponding to a uniform distribution of rotations. The angle of each point on the surface represents the orientation of the cluster. The distance of each point on the surface to the origin represents the normalised SPL of the cluster oriented along that axis. For example, a point far (near) from (to) the origin represents an orientation that dissipates high (low) heat.

Figure 6.15: Distribution of SPL as a function of the cluster orientation in three-dimensional space. Both plots convey the same distribution of SPL in two and three dimensions.

shows a relatively flat rotational SPL profile, with a few small bumps. In other words, the Arkus cluster will possess similar SPL values in all possible orientations.

Comparing the variation of SPL in Figure 6.15 to the variation in SPL between clusters of different size in Figure 6.5, shows that the variation between chains at large field values is smaller than the variation in SPL between different orientations of the same shape. These results suggest that the local heating variation in a solution of randomly oriented chains of various lengths will be dominated by the distribution of cluster orientations. On the other hand, the heating distribution for a solution of Arkus clusters of varying size and orientations will be dominated by the distribution of different cluster sizes. This suggests different design criteria for the two shapes in order to create even spatial heating.

Figure 6.15 also shows that the variance in the SPL as function of cluster orientation is deeply related to the geometric symmetry of the cluster. As discussed previously, the geometric symmetry (or more generally the fractal dimension) also accounts for the increase in shape anisotropy, coercive field, and SPL (subject to major-loop activation). We suggest that the requirements for large SPL values and flat distributions of heating are in direct competition with each other. One-dimensional shapes have the potential to dissipate very large SPL values when aligned with the field but misaligned particles may produce no SPL at all.

6.6 Conclusions

In this chapter we used stochastic simulations of coupled Landau-Lifshitz-Gilbert equations (LLG) to simulate the dynamic hysteresis loops of clusters of magnetic nanoparticles subjected to alternating magnetic fields. We compared the SPL of the clusters as a function of the applied field amplitude, size and orientations of the clusters. We showed that, depending on the strength of the field, all cluster shapes can show a non-monotonic relationship between SPL and the size of the cluster due to entering into minor-loop or paramagnetic-like regimes as the coercivity and relaxation timescales of the systems vary with subtle changes to the cluster configuration.

For large applied fields (i.e. major-loops) we predicted that chains dissipate substantially more heat than a single particle, which dissipates more than Arkus clusters. Although chains of particles aligned along the magnetic field direction are undoubtedly capable of the maximum SPL, whether Arkus clusters can stabilise single particles was not reliably determined. On the one hand, our simulations showed that dense clusters of particles lead to reduced SPL due to paramagnetic-like behaviour. On the other hand, existing experiments and kinetic Monte-Carlo results show an increase in coercivity leading to increased SPL. One possible explanation is that clusters tend to be formed of particles

with approximately aligned (rather than randomly distributed) anisotropy axes. Determining whether stochastic simulation of the LLG or kinetic Monte-Carlo methods are most suitable for simulating the dynamics of real clusters is difficult, since experimental observations of the magnetic state of the clusters and individual particles during a hyperthermia experiment are practically difficult to obtain. Therefore, when the methods disagree, both approaches can offer different possible explanations of experimental evidence.

Finally we investigated the distribution of SPL around the three-dimensional orientations of Arkus and chain clusters. The orientations of the particles within the clusters as well as the orientation of the clusters themselves have a substantial impact on the dissipated heat. Therefore, mixed solutions of multiple cluster types and sizes could lead to large spatial variations in heating, which is undesirable in hyperthermia procedures. The results suggested a direct conflict in optimising cluster shape to maximise SPL while also ensuring an even distribution of SPL over different orientations.

Multiple explanations exist for the substantial difference in SPL predicted by the kinetic Monte-Carlo and LLG simulations for almost identical systems. Investigating these discrepancies in more detail requires additional observations of the system state (for example, the evolution of the total energy during the simulation) in both simulations. Future work in this area would lead to an improved understanding of the limitations of the two approaches to accurately predict SPL in hyperthermia experiments.

Chapter 7

Non-sinusoidal alternating magnetic fields in hyperthermia

The work in this chapter was developed with collaborators: Dr. Robert Woodward, Prof. Tim St. Pierre, Michael McPhail during secondment to the University of Western Australia. I would like to acknowledge Dr. Woodward for his excellent intuition regarding the link between hysteresis area and transition rates in the master equation. I'd also like to acknowledge Michael McPhail for his derivation of a closed-form solution for a single particle subjected to a square-wave.

7.1 Introduction

The response of a magnetic nanoparticle to an alternating magnetic field (AMF) is modified by the amplitude, frequency, and shape of the AMF. The best combination of amplitude and frequency for *sinusoidal* AMFs that maximise heat dissipation in magnetic hyperthermia experiments are well understood. In contrast, the review in Section 3.3 showed that the effect of the AMF shape on heat dissipation has received relatively limited attention. The assumption that the applied AMF is purely sinusoidal overlooks, for example, the effect of higher order harmonic components in the form of thermal noise or novel shapes such as square or trapezoidal waveforms. Whether these additional components enhance or diminish heat dissipation in ensembles of magnetic nanoparticles is an important factor in determining the optimum conditions for hyperthermia.

Existing studies have shown that non-sinusoidal waves could significantly improve heat dissipation [123, 143, 144]. However, these studies use relatively small fields such that the particle response is within the region of linear response theory (LRT), Equation (3.1). The results of the LRT should not be extrapolated beyond their domain of validity, which limits their applicability to clinically relevant field strengths [90].

In this chapter, we investigate the advantage of subjecting non-interacting ensembles of magnetite Fe_3O_4 particles to square alternating magnetic fields rather than sinusoidal fields. Magnetite particles were selected for their biocompatibility [83] and previous use in a wide variety of medical applications, including hyperthermia [167]. We extend previous work by modelling the thermal activation of a Stoner Wohlfarth particle (Section 2.4.3) in the nonlinear region to predict the amplitude dependence of the energy dissipated under the square AMF. We use an analytic formula, derived from the nonlinear model, to quantify the limitations of the linear response theory. Finally, we investigate the effects of size-dispersed ensembles and filtered AMF signals on the heat dissipated.

7.2 Heat dissipation models for arbitrary fields

We considered a non-interacting ensemble of particles subjected to an alternating magnetic field $H(t)$ at time t . The magnetisation of the system $M(t)$ is probed in the same direction as the applied field. The applied field may be any periodic, even, real-valued function with period T and fundamental frequency $\omega_0 = 2\pi/T$, which can be expressed by the Fourier cosine series:

$$H(t) = H_0 \sum_{n=1}^{\infty} a_n \cos(n\omega_0 t) \quad (7.1)$$

where a_n is the amplitude of the n -th harmonic with angular frequency $n\omega_0$. After many cycles of the applied magnetic field, the resulting magnetisation response is considered to be in periodic equilibrium with the field such that $M(t) = M(t+T)$ for all t . The corresponding trajectory through the $M(t) - H(t)$ plane results in a closed-path termed the dynamic hysteresis loop. According to equation (2.48), the area of the hysteresis loop represents the energy dissipated per unit volume during a period T . An important detail is that different $H(t)$ functions have a different intrinsic power P_s as defined by the Fourier components $P_s = \frac{1}{2} \sum_{n=1}^{\infty} a_n^{-2}$. Therefore, a fair comparison of the energy dissipated by two different waveform shapes requires the normalised energy $\Delta U/P_s$.

Each nanoparticle in the ensemble was modelled as a Stoner-Wohlfarth particle (Section 2.2.2), which describes a single-domain particle with a fixed saturation magnetisation M_s (A/m), a uniaxial anisotropy of strength K (J/m³), and volume V (m³). The anisotropy axis, applied magnetic field, and magnetisation probing direction are all aligned. The state of the particle is described by $\theta(t)$ the angle between the anisotropy axis and the macrospin moment of the particle at time t . The energy of the particle is described in equation (2.9). The time-varying magnetisation $M(t) = M_s \cos^{-1}(\theta(t))$ of the particle is then determined by modelling the dynamics $\theta(t)$.

7.2.1 Thermal activation model

We modelled the nonlinear magnetic moment dynamics using the model of thermal activation as described in Section 2.4.3 using the Néel-Arrhenius law (equation (2.42)) for computing the transition rates $\Gamma_{i \rightarrow j}$ for $i, j \in \{1, 2\}$ between the two orientations θ_1, θ_2 . The model of thermal activation is only valid for sufficiently large energy barriers between the two orientations. More specifically we require that $\sigma(1 - h)^2 \gg 1$. Note that the transition rates in equation (2.42) ignore Brownian rotation (i.e. physical rotation) of the particles. Following the work of [46, 63] we consider the contribution from Brownian rotation to the heat dissipated to be negligible *in-vivo*. The model is entirely defined by the master equation (2.40) and the corresponding transition rates (2.42). Since the master equation is not solvable in the general case of time-dependent transition rates, we used Magpy to simulate approximate trajectories of the Stoner-Wohlfarth particle. See Chapter 5 for details on the design of Magpy.

7.2.2 Linear response theory

The linear response theory is described in Section 2.4.4. The limitation of the linear model is that it requires that both $2\sigma h \ll 1$ and $\sigma(1 - h)^2 \gg 1$, as was shown in Figure 3.1. By substituting the magnetisation response $M(t)$ of the linear model (equation (2.45)) into energy dissipation integral (2.48), the formula for heat dissipation under a harmonic AMF is obtained [171]:

$$\Delta U_{\cos}(H, \omega) = 4\pi K h^2 \sigma \frac{\omega \tau}{1 + \omega^2 \tau^2} \quad (7.2)$$

The magnetisation response to an arbitrary AMF may also be computed from the linear model by the application of the superposition principle of linear systems [183] to the Fourier series in equation (7.1). Recall from Section 2.4.4 that a property of linear systems is that the response to a harmonic input is a harmonic output of the same frequency but different amplitude and phase such that if $H(t) = H_0 e^{j\omega t}$ then:

$$M(t) = \chi(j\omega) H_0 e^{j\omega t} \quad (7.3)$$

where $\chi(j\omega)$ is the transfer function, which is termed the generalised susceptibility. If the applied magnetic field is represented by a complex Fourier series:

$$H(t) = H_0 \sum_{n=-\infty}^{\infty} c_n e^{jn\omega t} \quad (7.4)$$

Then, by the superposition principle of linear systems [183]:

$$M(t) = H_0 \sum_{n=-\infty}^{\infty} c_n \chi(jn\omega) e^{jn\omega t} \quad (7.5)$$

$$y(t) = \frac{dM(t)}{dt} = H_0 \sum_{n=-\infty}^{\infty} jn\omega \chi(jn\omega) e^{jn\omega t} \quad (7.6)$$

The Riemann-Stieltjes integral to compute the area of the hysteresis in equation (2.48), may be redefined if $M(t)$ is continuous and differentiable such that:

$$\Delta U = \mu_0 \int_0^T H(t) y(t) dt \quad (7.7)$$

$$= \mu_0 H_0^2 T \sum_{n=-\infty}^{\infty} c_n [c_n jn\omega \chi(jn\omega)]^* \quad (7.8)$$

where z^* represents the complex conjugate of z and we have used the Plancherel-Parseval theorem [141]. Substituting $T = 2\pi/\omega$ and $|c_n|^2 = c_n c_n^*$:

$$\begin{aligned} \Delta U &= 2\pi\mu_0 H_0^2 \sum_{n=-\infty}^{\infty} n |c_n|^2 \text{Im}[\chi(jn\omega)] \\ &\quad - j2\pi\mu_0 H_0^2 \sum_{n=-\infty}^{\infty} n |c_n|^2 \text{Re}[\chi(jn\omega)] \end{aligned} \quad (7.9)$$

The linear model for a single magnetic nanoparticle is a first-order system with time constant τ and gain χ_0 :

$$\chi(j\omega) = \chi_0 \frac{1}{1 + j\omega\tau} \quad (7.10)$$

which after substitution into equation (7.9) gives:

$$\begin{aligned} \Delta U &= 2\pi\mu_0 H_0^2 \chi_0 \sum_{n=-\infty}^{\infty} |c_n|^2 \frac{n^2 \omega \tau}{1 + n^2 \omega^2 \tau^2} \\ &\quad - j2\pi\mu_0 H_0^2 \chi_0 \sum_{n=-\infty}^{\infty} |c_n|^2 \frac{n}{1 + n^2 \omega^2 \tau^2} \end{aligned} \quad (7.11)$$

If the original signal is expressed as a Fourier cosine series with no dc term such that $c_n = \frac{1}{2}a_n = \frac{1}{2}a_{-n}$ and $a_0 = 0$ then we have:

$$\Delta U = \pi\mu_0 H_0^2 \chi_0 \sum_{n=1}^{\infty} a_n^2 \frac{n^2 \omega \tau}{1 + n^2 \omega^2 \tau^2} \quad (7.12)$$

which after comparing to the solution for a single harmonic (equation (7.2)), can be expressed as the sum of energy dissipated by each harmonic (equation (7.13)):

$$\Delta U_{\text{arb}} = \sum_{n=1}^{\infty} n \Delta U_{\text{cos}}(a_n H_0, n\omega_0) \quad (7.13)$$

From equation (7.13), the total energy dissipated can be interpreted as a linear combination of the energy dissipated at each harmonic of the AMF, weighted by its respective amplitude and proportional to its harmonic number n . The preceding n appears because the n th harmonic completes n full cycles in the time that the fundamental period completes one cycle. Consequently, high order harmonics contribute proportionally more energy to the total dissipated.

7.3 Comparison of low-amplitude square and sinusoidal AMFs

Using the linear response theory, we modelled the energy dissipated in ensembles of Stoner-Wohlfarth particles subjected to low-amplitude sinusoidal and square AMFs. The definition of the field function for a square wave over one period T is:

$$H(t) = \begin{cases} H_0 & 0 \leq t < T/2 \\ -H_0 & T/2 \leq t < T \end{cases} \quad (7.14)$$

The square wave has an intrinsic power $P_s = 1$ and Fourier components $|a_n| = \frac{4}{n\pi} \sin\left(\frac{n\pi}{2}\right)$. The heat dissipated by a single particle subjected to a square AMF was obtained by substituting the Fourier coefficients into equation (7.13):

$$\Delta U_{\text{square}}(H_0, \omega) = 16Kh^2\sigma \tanh\left(\frac{\pi}{2\omega\tau}\right) \quad (7.15)$$

Figure 7.1 shows the normalised energy dissipated as a function of $\omega\tau$ for square (equation (7.15)) and sinusoidal (equation (7.2)) waveforms. The normalised energy $\Delta U/(4Kh^2\sigma)$ accounts for the common factors between the two equations, resulting in a function of $\omega\tau$ alone. In other words, the difference in heat dissipated by the two AMF shapes depends only on the product of the particle relaxation time and the AMF frequency and not on the amplitude.

The results of linear response theory show that the square wave dissipates significantly more energy at low values of $\omega\tau$. The energy dissipated by a single harmonic experiences a peak at $\omega\tau$, when the relaxation time of the particle matches the AMF frequency, and tends towards zero for extreme positive and negative values of $\omega\tau$. The energy dissipated by the square AMF also tends towards zero with extreme positive values of $\omega\tau$ but at low values the energy dissipated reaches a maximum plateau. The maximum

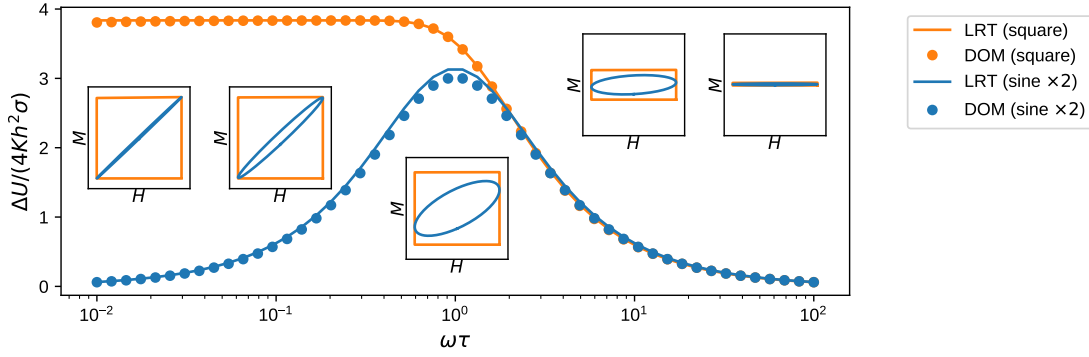


Figure 7.1: Normalised energy dissipated in a single magnetic nanoparticle subjected to a 0.05mT square and sinusoidal alternating magnetic field. The energy has been computed using both the linear response theory (LRT) and numerical simulations (DOM) with $M_s = 10^6 \text{Am}^{-1}$, $K = 10^4 \text{Jm}^{-3}$, $\alpha = 0.1$, $T = 300\text{K}$, $\mu_0 H = 0.05\text{mT}$, $f = 100\text{kHz}$, and r in the range of 8-12nm such that $2\sigma h \ll 1$ for all r . The curves for the two methods match exactly, validating the simulation approach. The results for the sine wave are multiplied by two to be of equivalent power to the square wave. The insets show the two hysteresis loops from the sine (blue) and square (orange) applied magnetic fields resulting from numerical simulations at the indicated value of $\omega\tau$.

ΔU_{square} is also larger than the maximum peak for a single harmonic. This apparent benefit of square waves may be interpreted from equation (7.13), which equates the total energy dissipated from an arbitrary wave to the sum of energy dissipated from each of its harmonics. Consider the square wave with harmonics at $\omega_0, 3\omega_0, 5\omega_0, \dots$ that are multiples of the base frequency ω_0 . When the base frequency of the square wave is too high $\omega_0\tau \gg 1$, the energy dissipated from the fundamental harmonic is very low and lower still for any $n\omega_0\tau$, since the energy dissipated decreases monotonically with $\omega\tau$ for $\omega\tau > 1$. However, when the fundamental frequency is too low $\omega_0\tau \ll 1$ a number of its higher order components will lie near the optimum frequency $n\omega \approx \omega_{\text{opt}}$. The sum total of the contributions from these higher order components are greater than the application of a single harmonic at the optimum frequency. Consequently, the square AMF results in a large ΔU_{square} for any fundamental frequency such that $\omega_0\tau \ll 1$.

The results from the linear response theory demonstrate that applying a square AMF dissipates more energy than a sinusoidal wave of equivalent power for any set of material parameters such that $2\omega h \ll 1$. The relative gain in energy dissipation from a square AMF qualitatively matches results previously obtained by numerical simulations of the Landau-Lifshitz-Gilbert equation [144], where the parameters used give a value $2\sigma h \approx 0.0009 \ll 1$.

At low field values, solutions to the discrete orientation model should match the predictions of the linear response theory. We computed the energy dissipated in an ensemble of identical nanoparticles of magnetite Fe_3O_4 subjected to a sinusoidal and square AMF with numerical simulations of the discrete orientation model. Given an initial condition

$\left[p_0(\theta_0) \quad p_0(\theta_1) \right]^T$, the master equation (2.40) was solved using an adaptive time step RK45 algorithm [164]. The system dynamics were simulated until reaching the periodic equilibrium $M(t) = M(t + T)$ and the area of the hysteresis loop was integrated numerically to determine the heat dissipation (see equation (2.48)). We used the Magpy software package (Chapter 5) to simulate the discrete orientation model using Néel-Brown transition rates. The parameters used for the simulation were $M_s = 10^6 \text{Am}^{-1}$, $K = 10^4 \text{Jm}^{-3}$, $\alpha = 0.1$, $T = 300\text{K}$, $\mu_0 H = 0.05\text{mT}$, $f = 100\text{kHz}$, and r in the range of 8-12nm, resulting in $2h\sigma < 0.087$ and $\sigma(1-h)^2 > 5.15$ for all values of the radius. The resulting energy dissipation was normalised and plotted against the linear response results in Figure 7.1. The inserts in the plot correspond to the simulated dynamic hysteresis loops in the M - H plane after 5 cycles of the external magnetic fields at various values of $\omega\tau$.

Figure 7.1 shows that the discrete orientation model matches the results predicted by the linear response theory, improving confidence in the implementation of the discrete orientation model. The dynamic hysteresis loops also offered an alternative perspective on the relative energy dissipation of the square and sine AMFs. At very high $\omega\tau$, the particle is unable to react fast enough to either of the AMF waveforms and appears frozen with zero magnetisation, which results in the hysteresis loop closing in the vertical direction. As $\omega\tau$ decreases, the harmonic AMF induces an harmonic response in the magnetisation leading to the characteristic elliptical hysteresis [165]. For a square AMF $H(t)$ only takes two values $\pm H_0$ and thus the resulting dynamic hysteresis loop is always oblong. At very low $\omega\tau$ the response of the magnetisation to the sinusoidal wave becomes superparamagnetic, the magnetisation tracks the shape of $H(t)$ without any lag. The magnetisation response of the particle must always be smooth due to magnetic viscosity and therefore cannot follow the discontinuity in the square AMF. This results in a fully open dynamic hysteresis area that cannot close even as $\omega \rightarrow 0$.

It has been previously shown by means of experimental measurements [46] and computational simulations [27] that the results of the linear response theory cannot be generalised to larger fields. The expressions for heat dissipation under the linear response model provide analytic solutions for arbitrary waveforms but are only valid for $2\sigma h \ll 1$.

7.4 Comparison of high-amplitude square and sinusoidal AMFs

A nonlinear model must be used if the heat dissipation resulting from high-amplitude square and sine AMFs are to be compared. For a sinusoidal AMF, we used numerical simulations of the Stoner-Wohlfarth thermal activation to compute ΔU_{\cos} for $2\sigma h > 1$. However, for square waves, we obtained an analytic solution to the master equation, which led to a new formula for ΔU_{square} valid outside of the linear region. The analytic




Figure 7.2: The magnetisation response of a magnetic nanoparticle to a square alternating magnetic field of period T (left). In dynamic equilibrium the magnetisation is assumed to oscillate between $\pm M_p$. The corresponding hysteresis loop in the M - H plane (right) is always an oblong and is defined by the amplitude of the applied field and the peak value of magnetisation M_p .

formula is valid for any σ and h such that $\sigma(1-h)^2 \gg 1$. Comparing the resulting expression to the LRT quantified the error in generalising the results of linear models to larger field amplitudes.

7.4.1 Analytic master-equation dynamics

Figure 7.2 depicts a sketch of the magnetisation response of a single particle subjected to a square wave of period T along its anisotropy axis. The system has reached a periodic equilibrium such that $M(t) = M(t+T)$ and we state without proof that due to the symmetry of the problem $M_p = M(T) = -M(T/2)$. Therefore, the dynamic hysteresis area integral, equation (2.48), is expanded, recalling equation (7.14):

$$\Delta U = -\mu_0 \oint_0^{T/2} H_0 \times dM(t) \quad (7.16)$$

$$- \mu_0 \oint_{T/2}^T -H_0 \times dM(t) \quad (7.17)$$

$$= 4\mu_0 H_0 M_p = 8K h_0 m_p \quad (7.18)$$

where $m_p = M_p/M_s$. Equation (7.16) and Figure 7.2 show that the dynamic hysteresis loop shape for a square AMF is always oblong, as seen during the low-amplitude simulations of the DOM in Figure 7.1. Computing ΔU thus only requires an estimate for the peak magnetisation M_p in periodic equilibrium.

The master equation (2.40) has analytic solutions if the transition rates in the equation are time-invariant [194]. Over the period $0 < t \leq T/2$ the field is constant and therefore the Néel transition rates are constant during the period. Given an initial condition $[p(\theta_1; 0) \ p(\theta_2; 0)]$, the solution of the master equation gives the probability of the

system occupying orientation θ_1 at time t :

$$p(\theta_1; t) = (p(\theta_1; 0) - \Gamma'_{12})e^{-\Gamma t} + \Gamma'_{12} \quad (7.19)$$

where $\Gamma = \Gamma_{12} + \Gamma_{21}$ and $\Gamma'_{12} = \Gamma_{12}/\Gamma$. The time-evolution of the probability of occupying state θ_2 is computed from $p(\theta_1; t) + p(\theta_2; t) = 1$. In order to compute M_p , we applied the boundary condition $M_p = M(T/2) = -M(0)$ or equivalently $p_1(0) = 1 - p_1(T/2)$. Assuming that the magnetic nanoparticle is in periodic equilibrium such that $M(t) = M(t + T) = -M(t + T/2)$ we have that $p(\theta_1; 0) = 1 - p(\theta_1; T/2)$ by using $M(t) = 2p(\theta_1; t) - 1$. Therefore, the solution for $p(\theta_1; T/2)$ given the initial condition $p(\theta_1; 0)$ at time $t = 0$ can be written:

$$p(\theta_1; T/2) = [1 - p(\theta_1; T/2) - \Gamma'_{12}]e^{-\Gamma T/2} + \Gamma'_{12} \quad (7.20)$$

Rearranging equation (7.20) we obtain:

$$p(\theta_1; T/2) = \frac{\Gamma_{12}e^{-\Gamma T/2} + \Gamma_{21}}{\Gamma(1 + e^{-\Gamma T/2})} \quad (7.21)$$

From which the magnetisation may be computed:

$$\frac{M(T/2)}{M_s} = 2p(\theta_1; T/2) - 1 \quad (7.22)$$

$$= 2 \frac{\Gamma_{12}e^{-\Gamma T/2} + \Gamma_{21}}{\Gamma(1 + e^{-\Gamma T/2})} - 1 \quad (7.23)$$

$$= \frac{(\Gamma_{21} - \Gamma_{12})(1 - e^{-\Gamma T/2})}{\Gamma(1 + e^{-\Gamma T/2})} \quad (7.24)$$

$$= \frac{\Gamma_{21} - \Gamma_{12}}{\Gamma} \tanh\left(\frac{\pi}{2\omega\tau}\right) \quad (7.25)$$

where we have used $\tau = \Gamma^{-1}$, $\omega = 2\pi T^{-1}$. If the transition rates are Néel rates with a constant prefactor f_0 :

$$\frac{\Gamma_{21} - \Gamma_{12}}{\Gamma} = \frac{f_0 e^{-\sigma(1-h)^2} - f_0 e^{-\sigma(1+h)^2}}{f_0 e^{-\sigma(1-h)^2} - f_0 e^{-\sigma(1+h)^2}} = \tanh(2\sigma h) \quad (7.26)$$

Note that this is not valid for transition rates using the Néel-Brown prefactor as it is not a constant. Using a constant prefactor, the magnetisation may be written:

$$m_p = \frac{M(T/2)}{M_s} = \tanh(2\sigma h) \tanh(\pi/(2\omega\tau)) \quad (7.27)$$

and substituted into equation (7.16) to obtain the energy dissipated, resulting in:

$$\Delta U_{\text{square}} = 8Kh \tanh(2h\sigma) \tanh\left(\frac{\pi}{2\omega\tau}\right) \quad (7.28)$$

We interpret equation (7.28) as a product of three terms: $8Kh_0$ is proportional to the




Figure 7.3: The equilibrium magnetisation as predicted by the linear (dashed) and nonlinear (solid) models, which is proportional to the normalised energy dissipated $\Delta U/(2\sigma h \tanh(\pi/(2\omega\tau)))$. The linear response model over-predicts ΔU for large field amplitudes $2\sigma h > 1$. Markers (circles) mark $2\sigma h = \frac{1}{3}$.

amplitude of the square AMF and determines the width of the hysteresis loop. From equation (7.16), the remaining two terms in equation (7.28) describe the peak magnetisation m_p and thus the height of the hysteresis loop. The peak magnetisation is therefore proportional to $\tanh(2\sigma h)$, which can be shown to be the equilibrium magnetisation of the particle subjected to a constant external field of amplitude h (see [27] for details on the equilibrium magnetisation under the discrete orientation model). Finally, the peak magnetisation is proportional to $\tanh(\pi/(2\omega\tau))$, which describes the effect of relaxation. Equation (7.28) allows ΔU_{square} to be evaluated for much stronger AMF amplitudes than the LRT solution, without numerical simulation.

From comparing the linear response solution for ΔU_{square} (equation (7.15)) to equation (7.16), we see that the peak magnetisation in the linear model is $m_p = 2\sigma h \tanh(\pi/(2\omega\tau))$. The peak magnetisation of the linear and nonlinear model depend identically on $\omega\tau$, which is expected as both models depend on an exponential relaxation process via the Néel transition rates, but have different values for the equilibrium magnetisation, which results from the assumption that the equilibrium magnetisation in the linear model is described by the constant static susceptibility $\chi_0 = 2\sigma h$. In Section 7.2.2, we stated that the linear model is only valid for small fields such that $2\sigma h \ll 1$. Indeed, we see that the linear and nonlinear expressions for ΔU_{square} are equivalent by the small angle approximation when $2\sigma h \ll 1$. Figure 7.3, shows the equilibrium magnetisation




Figure 7.4: Energy dissipated per cycle in a single magnetic nanoparticle subjected to a 5mT alternating magnetic field of sinusoidal and square shape as a function of the particle radius. The results were obtained using numerical simulations with the same material properties as in Figure 7.1. The insets show the dynamic hysteresis loop obtained from the simulations at the indicated radius.

of the linear and nonlinear models or equivalently the normalised energy dissipated $\Delta U_{\text{square}}/(8Kh_0 \tanh(\pi/(2\omega\tau)))$. Comparing the two solutions allows us to quantify the error in extending the LRT solution beyond its domain of validity to predict energy dissipation. The results show that the LRT solution always over-predicts the energy dissipated by the square wave, substantially so for $2\sigma h > 1/3$. This comparison highlights the danger of generalising the results of the linear response theory and prompts us to investigate again the relative benefit of a square AMF with higher amplitude.

7.4.2 Numerical simulations with high-amplitude fields

We simulated the energy dissipated in magnetite particles subjected to square and sinusoidal AMFs with an amplitude of 5mT using the discrete orientation model. The material parameters used for the simulation were the same as in Section 7.3. Figure 7.4 shows the energy dissipated as a function of the particle radius. The shape of the energy dissipation curve as a function of the relaxation time of the particle is a similar shape to the results with a weak field amplitude in Figure 7.1. In disagreement with the results of the linear model, Figure 7.4 shows that the maximum ΔU_{square} of the plateau is much lower than the maximum value of ΔU_{cos} . For larger particles, with longer relaxation times, the energy dissipation tends towards zero again for both of the field shapes. For smaller particles, the response to a sinusoidal field becomes superparamagnetic. Around the optimum radius for the particle ($R_{\text{opt}} \approx .8\text{nm}$ for magnetite and $f = 100\text{kHz}$) the sinusoidal AMF achieves a higher ΔU . The dynamic hysteresis loops (shown in the inset plots of Figure 7.4) reveal that near R_{opt} , the loop area of the sinusoidal wave is approaching that of the square wave. In the limit that the two hysteresis areas are the same, the sinusoidal wave will achieve twice the effective energy dissipation because it has half the intrinsic power of the square AMF. In other words, the square AMF would




Figure 7.5: The gain in energy dissipated resulting from applying a square alternating magnetic field. Gain is defined as $\log_{10}(\Delta U_{\text{square}}) - \log_{10}(\Delta U_{\text{sine}})$. A particle of magnetite of varying radius was subjected to each of the alternating magnetic field shapes with increasing field amplitude. Material properties as in figure 7.1. The contour line (0.0) represents the h, r combinations for which the energy dissipated is equal under both field shapes. The sinusoidal optimum radius (dash-dotted line) represents $r_{\text{opt}} = \text{argmax}_r \Delta U_{\text{sine}}(r)$ for each value of the applied field h .

require twice as much power to achieve the same energy dissipation in the ensemble. This is expected because a maximum energy dissipation exists that occurs when the particles are fully saturated by the applied field, yielding a maximum value $\Delta U_{\text{max}} = 4\mu_0 H_0 M_s$ for a given H_0 . If both the square and sinusoidal AMFs result in the same maximum dissipated energy then it is twice as efficient to use the sinusoidal AMF. We repeated the experiment for a range of AMF amplitudes h in order to investigate the boundary at which the single harmonic AMF dissipates more energy than the square AMF.

Figure 7.5 shows the relative gain from applying a square AMF over a sinusoidal AMF defined as $G = \log_{10}(\Delta U_{\text{square}}) - \log_{10}(\Delta U_{\text{sine}})$ for various values of the particle radius (relaxation time) and the applied field amplitude. The figure also shows the corresponding optimum radius R_{opt} at each value of the applied field amplitude. The figure shows that the boundary at which $\Delta U_{\text{square}} = 2\Delta U_{\text{cos}}$ is a complex function of the relaxation time of the system (or the AMF frequency), AMF amplitude, and the material properties. For particles with a very fast relaxation time relative to the AMF frequency, the square wave is always the optimum choice because the sinusoidal AMF induces a superparamagnetic behaviour, which dissipates negligible energy. Particles with a long

relaxation time relative to the AMF period are unable to effectively dissipate energy under either of the waves. Around the optimum frequency, the sinusoidal AMF dissipates its maximum energy and approaches the energy dissipation limit ΔU_{\max} as h increases. At this limit, the sinusoidal AMF is twice as efficient as the square AMF. The range of particle radii (relaxation times) for which the single harmonic outperforms the square AMF grows with increasing AMF amplitude h .

The results of the experiments show that the relative gain from the square AMF is a function of the AMF amplitude, frequency, and the relaxation time of the particle. A consistent benefit of the square AMF is that it dissipates significant energy in particles across a broader range of relaxation times. The fact that the harmonic AMF dissipates large ΔU in such a narrow peak has directed research efforts towards synthesising particles with as close to homogeneous properties as possible. An alternative approach would be to explore more robust hyperthermia practices, such as using novel AMF waveforms, that are able to perform effectively under a broad distribution of timescales.

7.5 Polydispersion and filtered AMFs

Magnetic nanoparticles synthesised for biomedical applications exhibit a distribution of sizes, shapes, and material properties leading to a distribution of relaxation timescales [202]. Additionally, experimental equipment does not produce *pure* square waveforms (i.e. with infinite harmonics) but are subjected to filtering effects resulting in attenuated higher harmonics, which we have shown are crucial to dissipate energy in particles with very short relaxation timescales (Section 7.3). These experimental realities violate the assumptions of the previous experiments used to compare the effectiveness of the square AMF. We investigated the effect of these imperfections by simulating the response of polydisperse ensembles of nanoparticles to filtered AMFs.

A practically relevant model for the distribution of particle sizes is the lognormal distribution [146]. The lognormal distribution is characterised by its mean r_0 and variance s :

$$p(r) = \frac{1}{rs\sqrt{2\pi}} \exp\left[-\frac{(\ln r - r_0)^2}{2s}\right] \quad (7.29)$$

We modelled ensembles of 1000 particles, aligned along their anisotropy axes, with polydisperse radii distributed lognormally with 50 values of r_0 between 5nm and 9.5nm and $s = 0.03, 0.12, 0.20$, totalling 150 different particle ensembles. We used the following parameters for the discrete orientation model: $M_s = 10^6 \text{Am}^{-1}$, $K = 2 \times 10^4 \text{Jm}^{-3}$, $\alpha = 0.1$, $T = 300\text{K}$, $\mu_0 H = 5\text{mT}$, $f = 100\text{kHz}$. The ensembles were subjected to four different AMF shapes: sinusoidal, square, approximate square with 10 Fourier components, approximate square with 100 Fourier components. Where we have used a truncated Fourier series as an approximation of the filtering effects of experimental equipment.

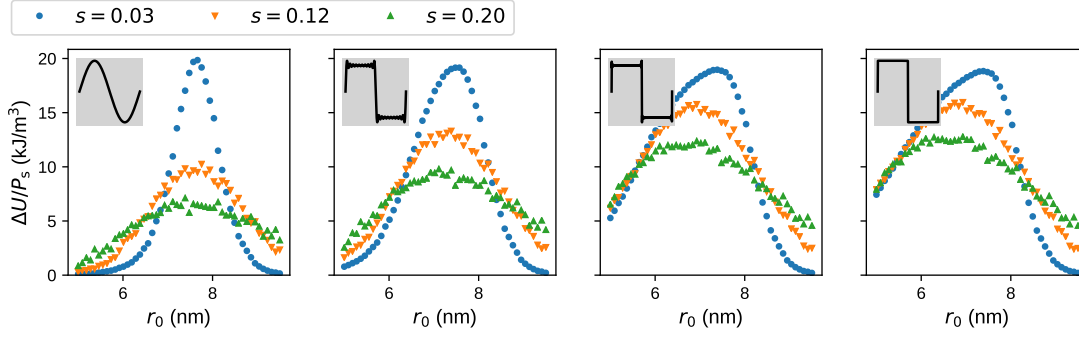


Figure 7.6: Energy dissipated in polydisperse ensembles of magnetic nanoparticles subjected to different applied magnetic field waveforms. The distribution of radii in each ensemble is lognormal with a mean parameter \bar{R} and scale parameter S . Simulations used the following parameters: $M_s = 10^6 \text{ Am}^{-1}$, $K = 2 \times 10^4 \text{ Jm}^{-3}$, $\alpha = 0.1$, $T = 300 \text{ K}$, $\mu_0 H = 5 \text{ mT}$, $f = 100 \text{ kHz}$. The figures (a-d) correspond to simulations using a sinusoidal, square (10 harmonics), square (100 harmonics), and square waveform respectively. The inset plots show a sketch of one period of each applied waveform shape.

Figure 7.6 shows the normalised energy dissipated by each of the 150 ensembles subjected to different AMF waveforms. For the approximately monodisperse ensembles ($s = 0.03$), the heat dissipation ΔU in ensembles with low r_0 was significantly improved by the presence of higher order harmonics. As the Fourier series was truncated, the smaller particles were no longer excited and ΔU dropped away quickly. On the other hand, for ensembles with large r_0 , the energy dissipation was reduced for all AMF shapes. For r_0 close to the optimum value, the width of the maximum delta peak depended on the quality of the signal. As expected, the effect of filtering the square AMF was to decrease ΔU contributions from smaller particles. The increased robustness of the square AMF diminishes with increasingly aggressive filtering.

Figure 7.6 shows that increasing the particle size dispersion also caused a broadening of the maximum energy peak for all signals. The effect of increasing dispersion was to smooth the dependence of ΔU on the mean ensemble radius r_0 . Crucially, ensembles with increased dispersion showed a reduced peak ΔU , which has also been shown in previous studies [171]. However, the energy dissipated by a square AMF showed less sensitivity (more robustness) to increasing size dispersion, the peak ΔU reduced relatively less. This result is due to the higher order harmonics, present in the square AMF, that dissipate heat in small particles. These small particles in the ensemble would not be excited by the standard harmonic wave. In other words, under a square AMF, a larger proportion of the ensemble is dissipating substantial energy. As the square AMF is increasingly filtered, and the higher order harmonics are reduced, fewer particles in the ensemble dissipate energy and the effect of size dispersity becomes more substantial.

We conclude from Figure 7.6 that the square AMF is more *robust* to broadening distributions of the particle properties than the sinusoidal AMF. However, as shown in Section 7.4, subjecting nanoparticles to a square AMF does not necessarily result in




Figure 7.7: Increase in maximum energy dissipation for square waves compared to sine waves at various field values for polydisperse ensembles. The LRT region marks the values of h for which the linear response theory is applicable.

an overall increase in the maximum attainable ΔU . We simulated three polydisperse ensembles with $s = 0.03, 0.1, 0.2$ subjected to both AMF shapes with increasing field amplitude h_0 . For each value of h_0 , we tuned the field frequency to dissipate the maximum energy possible (recalling from figure 7.5 that the optimum applied frequency varies with h_0). Figure 7.7 shows the relative gain in energy dissipated $\Delta U_{\text{square}}/\Delta U_{\text{cos}}$ for the three ensembles. The largest gain is realised for very low field amplitudes, which correspond to the linear regime. Within this region the gain is independent of h_0 . As the field increases, the maximum energy dissipation attainable using a square AMF begins to diminish relative to the sinusoidal AMF. Eventually, the hysteresis area of the two waves approach the $8Kh_0M_s$ limit as the particles fully saturate. We expect that as the field amplitude increases further¹ the gain for all three ensembles will approach 0.5, resulting from both AMF shapes saturating the entire ensemble. The dependence of the gain on h_0 is similar for all three ensembles but the total realised gain increases with increasing dispersity. For the most disperse ensemble ($s = 0.2$) the square AMF still outperforms the sinusoidal AMF for relatively large h .

¹Numerical simulations of the Stoner-Wohlfarth particles were limited by the condition $\sigma(1-h)^2 \ll 1$

7.6 Conclusions

We investigated the increase in energy dissipated in ensembles of magnetic nanoparticles subjected to a square rather than a sinusoidal alternating magnetic field. Using analytic and numerical simulations, we showed that square fields result in large dynamic hysteresis areas due to their fast switching and dissipate energy in particles with low relaxation times, which would otherwise exhibit superparamagnetic behaviour in response to a sinusoidal field. Consequently, total energy dissipation is less sensitive to polydisperse ensembles of particles with a broad distributions of relaxation times. Using square alternating magnetic fields for hyperthermia experiments may lead to improved robustness to variations in the material properties of synthesised nanoparticles.

We have further shown that the relative gain in performance from the square field diminishes with increasing field amplitude. Performance is worsened further from the effects of filtering, which attenuate higher order harmonics that are crucial for activating smaller particles. For applications with approximately monodisperse particles and large fields, the sinusoidal waveform provides the most efficient energy dissipation.

The conclusions of our work point towards new possibilities for using novel alternating magnetic fields to improve the robustness and efficacy of hyperthermia treatments. Future work should validate the results of our calculations with experiments of square alternating magnetic fields applied to magnetic nanoparticles. However, the current safe limits on the applied magnetic field are defined only for sinusoidal waveforms, which prevents clinical testing currently. While there is some limited guidance available [156, 157], more research into the safe limits of the field are needed.

Chapter 8

Conclusions and outlook

Effective magnetic hyperthermia treatments require magnetic nanoparticles with optimum properties to maximise the heat dissipated per weight of administered nanoparticle fluid. Maximising heat dissipation allows for smaller doses of nanoparticles and the use of weaker alternating magnetic fields during the procedure. This thesis detailed the construction, design choices, and testing of numerical methods to simulate the stochastic dynamics of interacting magnetic nanoparticles using coupled Landau-Lifshitz-Gilbert equations. The simulations were used to predict the dynamic hysteresis loops generated by subjecting nanoparticle ensembles to alternating magnetic fields. The simulated hysteresis loops allowed the heat dissipation to be computed and investigated as a function of material, field, and environmental parameters as well as arbitrary configurations of particles in three-dimensional space.

Simulating the stochastic Landau-Lifshitz-Gilbert equation first required a numerical procedure to solve coupled stochastic differential equations. In Chapter 4, the popular Heun scheme was compared with a fully implicit solver. The tests focused on the path-wise error between analytical solutions and the estimated trajectories. The tests showed that conservation of the magnetisation vector does not necessarily reduce the path-wise error and therefore projection methods may not improve the performance of the Heun scheme. The implicit method was shown to be significantly more accurate than the explicit Heun scheme for systems with strong thermal fluctuations (i.e. stochastic component), while conserving the length of the magnetisation vector. However, the complexity of the algorithm implementation and the consequent performance demonstrated that the implicit scheme is computationally unfeasible for large systems.

The numerical methods discussed in Chapter 4 were implemented in an open-source software package Magpy. Magpy was implemented in the python programming language to encourage reuse and accelerated with a C++ library in order to maximise the simulation time per CPU time. The software package was tested with unit tests, numerical tests and physics tests to assure the quality of the implementation of the code, design

of the numerical methods, and the correct behaviour of the simulated physical systems respectively.

Magpy was used to simulate the dynamic hysteresis loops of dense clusters and chain-like clusters of particles in Chapter 6. Intra-cluster interactions were shown to substantially alter the heat dissipated by magnetic nanoparticles, which was highly sensitive to small variations in the cluster configurations. In general, chain-like structures were shown dissipate the most heat per weight of material. However chain structures require large fields in order to activate. In a solution of chains of multiple lengths, it is possible that many structures dissipate large heat while others remain with no heat dissipation due to the applied field being below the coercive field. In experiments this could lead to large spatial variations in heat. Additionally, chain-like structures show a large variation in heat dissipation with the orientation of the structure to the applied field; increasing the spatial heating distribution further. The results suggested a conflict between maximising heat dissipation by configuring particles along the applied field direction and ensuring a uniform distribution by configuring particles in a three-dimensional isotropic shape (such as dense clusters).

Dense clusters of particles were shown, in general, to reduce the SPL due to cycling through low-energy states favouring end-to-end alignment of the magnetic moments. The simulations were used to reproduce an experimental result showing that larger clusters were able to dissipate more effectively than single particles through an increase in coercivity. However, our simulations did not match this result and showed dense clusters in a paramagnetic-like state. The simulations offered an alternative explanation of the experiment that the anisotropy axes of particles within the same cluster are roughly aligned. Magpy was used to replicate a kinetic Monte-Carlo experiment, which failed to reproduce the same results. Comparing the Magpy results to existing simulations was difficult without experimental evidence of the correct dynamic hysteresis loop. The SPL values obtained in experiments can map to many different dynamical behaviours. A detailed comparison of Landau-Lifshitz-Gilbert dynamics (with different numerical methods) and kinetic Monte-Carlo dynamics would be very useful for determining the limitations of both methods. For example, whether thermal activation based models (such as kinetic Monte-Carlo) should be used for systems with substantial dipolar interaction fields.

The effect of applying square-wave shaped alternating magnetic fields to ensembles of non-interacting particles was investigated in Chapter 7. An analytic approach, using the linear response theory, revealed that the square-wave could substantially outperform the sinusoidal-wave since the fast switching times were able to effectively dissipate heat in small particles that would have been paramagnetic otherwise. However, numerical simulations at larger field amplitudes, showed that the benefit of a square-wave diminishes with increasing fields. An ensemble of monodisperse, non-interacting particles are most efficiently activated using a sinusoidal-wave tuned to the correct frequency. Whereas

polydisperse ensembles could dissipate substantial heat with a square-wave. The results suggest that square-waves could be used to increase heat dissipation at low fields and improve the robustness of current procedures. However, practical guidelines for the use of square-wave alternating magnetic fields in clinical applications have not yet been published.

The simulations in this thesis showed that the alternating magnetic field shape and the configuration of particle clusters play an important role in the dynamic hysteresis loops of magnetic nanoparticles. Sections 6.6 and 7.6 describe future routes of investigation for these experiments. These investigations were enabled by developing efficient and easy-to-use computational tools for simulating hyperthermia experiments. However, a better understanding of the appropriate numerical methods and their ability to accurately represent the true dynamics of physical magnetic nanoparticle systems is needed. Future numerical methods (discussed in Section 4.5) will not provide utility to the research community without reliable, well tested and documented implementations, as described in Chapter 5. Reliable, open-source simulation software can reduce the total resource used by the community developing individual solvers to these complex problems. Moreover, modular and documented code can increase the rate at which new developments transition from research into practical use by computational scientists. In turn, this can reduce the time to discovering favourable conditions to maximising heat dissipation in magnetic particles and improving the outcomes of clinical trials.

Appendix A

Stochastic integrals

The stochastic integral defines the integration of a function randomly varying function $b(t)$ with respect to the Wiener process $\mathbf{W}(t)$; it is a generalisation of a Riemann-Stieltjes integral. For a deterministic function $a(t)$, the Riemann-Stieltjes integral can be defined as the summation of a series of N intervals of width $\Delta t = T/N$ such that:

$$\int_0^T a(t)dt = \lim_{N \rightarrow \infty} \sum_{k=0}^{N-1} a(\hat{t}_k)(t_{k+1} - t_k) \quad (\text{A.1})$$

where $a(\hat{t}_k)$ is a constant value within each interval and $\hat{t}_k \in [t_k, t_{k+1}]$ may be freely chosen. In other words, the constant value of the function $a(\hat{t}_k)$ may be chosen at any point along the interval. As the number of intervals tends to infinity and the interval width becomes infinitesimally small the summation converges to the true value of the integral. This result holds true for any choice of \hat{t}_k .

The stochastic analogue of equation (A.1) to integrate $b(t)$ could be written:

$$\lim_{N \rightarrow \infty} \sum_{k=0}^{N-1} b(\hat{t}_k)(W(t_{k+1}) - W(t_k)) \quad (\text{A.2})$$

where $W(t_k)$ is the Wiener process at t_k . Once again, the value of $\hat{t}_k \in [t_k, t_{k+1}]$ must be chosen. However, unlike the deterministic Riemann-Stieltjes integral, the choice of \hat{t}_k changes the result of the series and therefore the result of equation (A.2) is ambiguous.

For example, this can be seen by substituting $b(t)$ for the Wiener process $W(t)$ and computing the expected value of the series:

$$E \left\langle \sum_{k=0}^{N-1} b(\hat{t}_k) (W(t_{k+1}) - W(t_k)) \right\rangle = E \left\langle \sum_{k=0}^{N-1} W(\hat{t}_k) (W(t_{k+1}) - W(t_k)) \right\rangle \quad (\text{A.3})$$

$$= \sum_{k=0}^{N-1} E \langle W(\hat{t}_k) W(t_{k+1}) \rangle - E \langle W(\hat{t}_k) W(t_k) \rangle \quad (\text{A.4})$$

$$= \sum_{k=0}^{N-1} \hat{t}_k - t_k \quad (\text{A.5})$$

where $E \langle W(t_1) W(t_2) \rangle = \min(t_1, t_2)$ is a property of the Wiener process. The final result depends on the choice of \hat{t}_k . This choice will influence the form of $b(\cdot)$ when modelling a stochastic process and as long as the mathematics are consistent, there is no single correct choice.

The most common choices of the interval evaluation point are the Itô stochastic integral $\hat{t}_k = t_k$ and the Stratonovich stochastic integral $\hat{t}_k = \frac{1}{2}(t_{k+1} + t_k)$. These interpretations confer certain mathematical conveniences. For example, the Stratonovich integral is consistent with the chain rule and therefore behaves similarly to deterministic calculus.

In summary, each stochastic differential equation (SDE) may be derived under any interpretation of the stochastic integral (usually the Itô or Stratonovich interpretations). However, the SDE must be interpreted consistently using the same stochastic integral. Different interpretations may be used by transforming the SDE: for example the Itô correction formula for transforming a Stratonovich SDE to an Itô SDE (see Section 2.3.1). See [193] for a more detailed discussion of the Itô-Stratonovich dilemma.

Appendix B

Normalisation

The following calculations are in SI units. A particle i has a magnetic moment $\vec{\mu}$, the magnetisation is defined as the magnetic moment per unit volume $\vec{M} = \vec{\mu}/V_i$. The saturating magnetisation M_s is used to normalise the magnetisation to the unit vector $\hat{m} = \vec{M}/M_s$.

The total energy including contributions from the anisotropy (constant K_i and axis direction \hat{k}_i), applied field (\vec{H}), and dipolar interactions (between two particles i, j separated by the vector $\vec{R}_{ij} = |\vec{R}_{ij}| \hat{r}_{ij}$) is:

$$E_i = K_i V_i \left| \hat{m}_i \times \hat{k}_i \right|^2 - \mu_0 V_i \vec{M}_i \cdot \vec{H} - \sum_{j \neq i} \mu_0 \frac{V_i V_j}{4\pi |\vec{R}_{ij}|^3} \left(3 \left(\vec{M}_i \cdot \hat{r}_{ij} \right) \left(\vec{M}_j \cdot \hat{r}_{ij} \right) - \vec{M}_i \cdot \vec{M}_j \right) \quad (\text{B.1})$$

where μ_0 is the permeability of free space.

Defining the average particle volume and anisotropy constants as $\bar{V} = 1/N \sum_i V_i$ and $\bar{K} = 1/N \sum_i K_i$ respectively, introduce the dimensionless volume and anisotropy constants $v_i = V_i/\bar{V}$ and $k_i = K_i/\bar{K}$. The normalised energy is defined as:

$$\begin{aligned} e &= \frac{e}{2\bar{V}\bar{K}} \\ &= \frac{1}{2} k_i v_i \left| \hat{m}_i \times \hat{k}_i \right|^2 - \mu_0 \frac{v_i}{2\bar{K}} \vec{M}_i \cdot \vec{H} - \sum_{j \neq i} \frac{\mu_0}{2\bar{K}} \frac{v_i V_j}{4\pi |\vec{R}_{ij}|} \left(3 \left(\vec{M}_i \cdot \hat{r}_{ij} \right) \left(\vec{M}_j \cdot \hat{r}_{ij} \right) - \vec{M}_i \cdot \vec{M}_j \right) \\ &= \frac{1}{2} k_i v_i \left| \hat{m}_i \times \hat{k}_i \right|^2 - \mu_0 \frac{v_i M_s}{2\bar{K}} \hat{m}_i \cdot \vec{H} - \sum_{j \neq i} \frac{\mu_0 M_s^2}{2\bar{K}} \frac{v_i V_j}{4\pi |\vec{R}_{ij}|} \left(3 \left(\hat{m}_i \cdot \hat{r}_{ij} \right) \left(\hat{m}_j \cdot \hat{r}_{ij} \right) - \hat{m}_i \cdot \hat{m}_j \right) \end{aligned} \quad (\text{B.2})$$

The anisotropic field is defined as:

$$H_k = \frac{2\bar{K}}{M_s \mu_0} \quad (\text{B.3})$$

so that the dimensionless field can be introduced $\vec{H} = H_k \vec{h}$:

$$e = \frac{1}{2} k_i v_i \left| \hat{m}_i \times \hat{k}_i \right|^2 - v_i \hat{m} \cdot \vec{h} - \sum_{j \neq i} \frac{\mu_0 M_s^2}{2\bar{K}} \frac{v_i V_j}{4\pi \left| \vec{R}_{ij} \right|} \left(3 (\hat{m}_i \cdot \hat{r}_{ij}) (\hat{m}_j \cdot \hat{r}_{ij}) - \hat{m}_i \cdot \hat{m}_j \right) \quad (\text{B.4})$$

The distance between two particles is normalised by a constant term a so that $\vec{R}_{ij} = a \vec{r}_{ij}$. Introducing this normalisation and recalling that $V_i = \bar{V} v_i$:

$$e = \frac{1}{2} k_i v_i \left| \hat{m}_i \times \hat{k}_i \right|^2 - v_i \hat{m} \cdot \vec{h} - \sum_{j \neq i} \frac{\mu_0 M_s^2}{2\bar{K}} \frac{v_i v_j}{4\pi (a^3/\bar{V}) \left| \vec{r}_{ij} \right|} \left(3 (\hat{m}_i \cdot \hat{r}_{ij}) (\hat{m}_j \cdot \hat{r}_{ij}) - \hat{m}_i \cdot \hat{m}_j \right) \quad (\text{B.5})$$

note that the term a^3/\bar{V} is dimensionless. The effective field for a particle i is defined as:

$$\vec{H}_{\text{eff}} = -\frac{1}{\mu_0 V_i} \frac{\partial E}{\partial \vec{M}_i} = -\frac{2\bar{V}\bar{K}}{\mu_0 M_s V_i} \frac{\partial e}{\partial \hat{m}_i} = -H_k \frac{1}{v_i} \frac{\partial e}{\partial \hat{m}_i} \quad (\text{B.6})$$

The dimensionless effective field \vec{h}_{eff} then follows:

$$\therefore \vec{h}_{\text{eff}} = -\frac{1}{v_i} \frac{\partial e}{\partial \hat{m}_i} = -\frac{1}{v_i} \left[-k_i v_i (\hat{m}_i \cdot \hat{k}_i) \hat{k}_i - v_i \vec{h} - \sum_{j \neq i} \frac{\mu_0 M_s^2}{2\bar{K}} \frac{v_i v_j}{4\pi (a^3/\bar{V}) \left| \vec{r}_{ij} \right|^3} \left(3 (\hat{m}_j \cdot \hat{r}_{ij}) \hat{r}_{ij} - \hat{m}_j \right) \right] \quad (\text{B.7})$$

Leading to the final form of the dimensionless effective field for a particle i . All terms are dimensionless and thus well conditioned for numerical computation (far from machine precision):

$$\vec{h}_{\text{eff}} = k_i (\hat{m}_i \cdot \hat{k}_i) \hat{k}_i + \vec{h} + \frac{\mu_0 M_s^2}{8\pi \bar{K}} \sum_{j \neq i} \frac{v_j}{(a^3/\bar{V}) \left| \vec{r}_{ij} \right|^3} \left(3 (\hat{m}_j \cdot \hat{r}_{ij}) \hat{r}_{ij} - \hat{m}_j \right) \quad (\text{B.8})$$

The Landau-Lifshitz-Gilbert equation that governs the dynamical behaviour of the magnetisation \vec{M} of a single domain nanoparticle is:

$$\frac{d\vec{M}}{dt} = -\frac{\gamma}{1+\alpha^2} \vec{M} \times (\vec{B} + \vec{B}_{\text{th}}) - \frac{\gamma\alpha}{(1+\alpha^2) M_s} \vec{M} \times (\vec{M} \times (\vec{B} + \vec{B}_{\text{th}})) \quad (\text{B.9})$$

where M_s is the magnitude of the vector \vec{M} , γ is the magnetogyric ratio, α is the dimensionless damping parameter, and \vec{B} is the effective field in Tesla:

$$\vec{B} = -\frac{1}{V} \frac{\partial E}{\partial \vec{M}} \quad (\text{B.10})$$

\vec{B}_{th} is the thermal field, which is a fluctuating field such that:

$$\langle B_{\text{th}}(t) B_{\text{th}}(t + \delta t) \rangle = \delta(t - t - \delta t) \sigma^2 \delta t \quad (\text{B.11})$$

The dimensionless magnetisation is a unit vector defined $\hat{m} = \vec{M}/M_s$ such that:

$$\frac{d\hat{m}}{dt} = -\frac{\gamma}{1+\alpha^2}\hat{m} \times (\vec{B} + \vec{B}_{\text{th}}) - \frac{\gamma\alpha}{1+\alpha^2}\hat{m} \times (\hat{m} \times (\vec{B} + \vec{B}_{\text{th}})) \quad (\text{B.12})$$

The \vec{H} field has units ampere per meter and is related to the \vec{B} field through $\vec{B} = \mu_0\vec{H}$:

$$\frac{d\hat{m}}{dt} = -\frac{\gamma\mu_0}{1+\alpha^2}\hat{m} \times \left(\vec{H} + \frac{1}{\mu_0}\vec{B}_{\text{th}}\right) - \frac{\gamma\mu_0\alpha}{1+\alpha^2}\hat{m} \times \left(\hat{m} \times \left(\vec{H} + \frac{1}{\mu_0}\vec{B}_{\text{th}}\right)\right) \quad (\text{B.13})$$

The anisotropy field H_k has units of ampere per meter and is used to obtain the dimensionless effective field $\vec{h} = \vec{H}/H_k$:

$$\frac{d\hat{m}}{dt} = -\frac{\gamma\mu_0 H_k}{1+\alpha^2}\hat{m} \times (\vec{h} + \vec{h}_{\text{th}}) - \frac{\gamma\mu_0 H_k \alpha}{1+\alpha^2}\hat{m} \times (\hat{m} \times (\vec{h} + \vec{h}_{\text{th}})) \quad (\text{B.14})$$

where:

$$H_k = \frac{2K}{\mu_0 M_s} \quad (\text{B.15})$$

and

$$\langle h_{\text{th}}(t) h_{\text{th}}(t + \delta t) \rangle = \delta(t - t - \delta t) \frac{\sigma^2}{\mu_0^2 H_k^2} \delta t \quad (\text{B.16})$$

Bringing these terms to the other side, we obtain a dimensionless right-hand-side, the thermal field has also been written to show its explicit time dependence:

$$\frac{1+\alpha^2}{\gamma\mu_0 H_k} \frac{d\hat{m}}{dt} = -\hat{m} \times (\vec{h} + \vec{h}_{\text{th}}(t)) - \alpha\hat{m} \times (\hat{m} \times (\vec{h} + \vec{h}_{\text{th}}(t))) \quad (\text{B.17})$$

Now introduce the reduced time ℓ , using equation (B.15):

$$\ell = t \frac{\gamma\mu_0 H_k}{1+\alpha^2} = t \frac{2K\gamma}{M_s(1+\alpha^2)} \quad (\text{B.18})$$

Substituting ℓ into equation (B.17) yields the dimensionless form of the LLG:

$$\frac{d\hat{m}}{d\ell} = -\hat{m} \times \left(\vec{h} + \sqrt{\frac{\gamma\mu_0 H_k}{1+\alpha^2}} \vec{h}_{\text{th}} \left(\frac{1+\alpha^2}{\gamma\mu_0 H_k} \ell \right) \right) - \alpha\hat{m} \times \left(\hat{m} \times \left(\vec{h} + \sqrt{\frac{\gamma\mu_0 H_k}{1+\alpha^2}} \vec{h}_{\text{th}} \left(\frac{1+\alpha^2}{\gamma\mu_0 H_k} \ell \right) \right) \right) \quad (\text{B.19})$$

where the fluctuating field has been scaled such that the equation is equivalent to:

$$\frac{d\hat{m}}{d\ell} = -\hat{m} \times (\vec{h} + \vec{h}_{\text{th}}(\ell)) - \alpha\hat{m} \times (\hat{m} \times (\vec{h} + \vec{h}_{\text{th}}(\ell))) \quad (\text{B.20})$$

$$\langle h_{\text{th}}(\ell) h_{\text{th}}(\ell + \delta\ell) \rangle = \delta(\ell - \ell - \delta\ell) \frac{\sigma_s^2}{\mu_0^2 H_k^2} \delta\ell \quad (\text{B.21})$$

Appendix C

Arkus cluster definitions

The following code block contains the coordinates of particles that make up the Arkus particle cluster geometries as defined in [7]. The code is written in python and can be imported as a standalone module.

```
1  """
2  Coordinates for Arkus cluster geometries.
3
4  This module contains the coordinates of particles arranged
5  into Arkus cluster geometries of varying size. Arkus clusters are
6  *minimum energy* clusters that represent tightly packed clusters of
7  spherical particles. Their definition (and name) are derived from:
8
9  N. Arkus, V. N. Manoharan, and M. P. Brenner, *Phys. Rev. Lett.* **103**, 118303 (2009).
10 `http://dx.doi.org/10.1103/PhysRevLett.103.118303`
11
12 Clusters of 1-5 particles have only one possible arrangement.
13 Clusters of 6 or more particles have a variety of different
14 configurations, which are accessed using a configuration id. The
15 available configuration ranges are:
16   - 1-5 particles: [0]
17   - 6 particles: [0-1]
18   - 7 particles: [0-5]
19   - 8 particles; [0-12]
20
21 The coordinates of the particles within the cluster are normalised
22 to a unit distance between each pair (where possible). The distance
23 between the particles can be controlled by multiplying by a scaling
24 factor.
25
26 Attributes:
27   ARKUS (dict): a dictionary containing the geometries of Arkus clusters
28   of 1-8 particles. The dictionary is accessed
29   `ARKUS[n_particles][configuration_id]` and returns an `np.ndarray`
30   of shape `(n_particles,3)` containing the coordinates of particles
```

```

31         in the cluster.
32
33     Examples:
34         .. code-block:: python
35
36         >>> ARKUS[1][0]
37         array([[0,0,0]])
38         >>> ARKUS[2][0]
39         array([[0,0,0],
40               [0,0,1]])
41         >>> ARKUS[2][0] * 1e-9 # apply a scaling factor
42         array([[0,0,0],
43               [0,0,1e-9]])
44     """
45
46     import numpy as np
47
48     ARKUS = {
49
50         # 1 PARTICLE
51         1: {
52             0: np.array([
53                 [0,0,0]
54             ])
55         },
56
57         # 2 PARTICLE
58         2: {
59             0: np.array([
60                 [0,0,0],
61                 [0,0,1]
62             ])
63         },
64
65         # 3 PARTICLE
66         3: {
67             0: np.array([
68                 [0,0,0],
69                 [0,0,1],
70                 [0,0.866025403837047,0.5]
71             ])
72         },
73
74         # 4 PARTICLE
75         4: {
76             0: np.array([
77                 [0,0,0],
78                 [0,-1,0],
79                 [0.866025403784439, -0.5, 0],
80                 [0.288675134594813, -0.5, 0.816496580934550]
81             ])

```

```
82     },
83
84     # 5 PARTICLE
85     5: {
86         0: np.array([
87             [0.0, 0.0, 0.0],
88             [0.0, 1.632993161855915, 0.0],
89             [-0.577350269189415, 0.816496580927958, 0.0],
90             [0.288675134595175, 0.816496580927958, -0.499999999999947],
91             [0.288675134595175, 0.816496580927958, 0.500000000000006]
92         ])
93     },
94
95     # 6 PARTICLE
96     6: {
97         0: np.array([
98             [0.0, 0.0, 0.0],
99             [0.0, 1.632993162089509, 0.0],
100            [0.962250448720120, 0.272165527131513, 0.0],
101            [-0.577350269227751, 0.816496581044756, -0.000000000001533],
102            [0.288675134563603, 0.816496581044756, -0.500000000000726],
103            [0.288675134563603, 0.816496581044755, 0.499999999999382]
104        ]),
105        1: np.array([
106            [0.0, 0.0, 0.0],
107            [0.0, -1.414213562373241, 0.0],
108            [-0.707106781186482, -0.707106781186621, 0.0],
109            [0.707106781187920, -0.707106781186621, -0.000000000000000],
110            [0.000000000000126, -0.707106781186621, -0.707106781186493],
111            [0.000000000000126, -0.707106781186621, 0.707106781186493]
112        ])
113     },
114
115     # 7 PARTICLE
116     7: {
117         0: np.array([
118             [0.0, 0.0, 0.0],
119             [0.0, 1.632993164582268, 0.0],
120             [1.026400523384526, 1.270105797971256, 0.0],
121             [-0.577350267595472, 0.816496582291134, 0.000000000023495],
122             [0.962250473328791, 0.272165526445590, 0.000000000034229],
123             [0.288675141480240, 0.816496582291133, 0.499999999969453],
124             [0.288675141493603, 0.816496582291136, -0.500000000035519]
125        ]),
126        1: np.array([
127            [0.0, 0.0, 0.0],
128            [0.0, 1.666666666672721, 0.0],
129            [1.443375673009152, 0.8333333333386875, 0.0],
130            [-0.096225044838686, 0.8333333333336361, -0.544331053982005],
131            [0.769800358948834, 0.333333333331258, -0.544331053982335],
132            [0.769800358945444, 1.333333333372645, -0.544331053954963],
```



```

133         [0.481125224339845, 0.833333333336361, 0.272165526958015]
134     ]),
135     2: np.array([
136         [0.0, 0.0, 0.0],
137         [0.0, 1.632993161864094, 0.0],
138         [-0.962250448673535, 0.272165526969049, 0.0],
139         [0.481125224326110, 1.360827634888840, 0.83333333333549],
140         [-0.288675134603241, 0.816496580932047, -0.500000000000039],
141         [0.577350269184733, 0.816496580932047, -0.00000000000013],
142         [-0.288675134603224, 0.816496580932047, 0.49999999999964]
143     ]),
144     3: np.array([
145         [0.0, 0.0, 0.0],
146         [0.0, 1.618033988749913, 0.0],
147         [0.951056516295152, 0.309016994374951, 0.0],
148         [-0.587785252292505, 0.809016994374957, -0.000000000000003],
149         [0.951056516295155, 1.309016994374955, -0.000000000000003],
150         [0.262865556059561, 0.809016994374957, -0.525731112119131],
151         [0.262865556059561, 0.809016994374957, 0.525731112119127]
152     ]),
153     4: np.array([
154         [0.0, 0.0, 0.0],
155         [0.0, 1.414213562375728, 0.0],
156         [1.000000000039960, -0.000000000007774, 0.0],
157         [-0.500000000011403, 0.707106781187864, 0.499999999996897],
158         [0.50000000002506, 0.707106781187864, -0.49999999999995],
159         [-0.500000000056301, 0.707106781187864, -0.500000000018837],
160         [0.500000000002484, 0.707106781187864, 0.500000000000006]
161     ])
162 },
163
164 # 8 PARTICLE
165 8: {
166     0: np.array([
167         [0.0,0.0,0.0],
168         [0.0,-1.666666685130705,0.0],
169         [-1.443375683706607,-0.833333339154295,0.0],
170         [-0.481125182280860,-0.833333352769395,-1.360827663935889],
171         [-0.481125236072751,-0.833333342565355,0.272165540947502],
172         [0.096225067033021,-0.833333342565357,-0.544331041769898],
173         [-0.769800347378054,-0.333333348343040,-0.544331066284295],
174         [-0.769800346694231,-1.333333354811085,-0.544331083525413]
175     ]),
176     1: np.array([
177         [0.0,0.0,0.0],
178         [0.0,1.632993161855456,0.0],
179         [1.026400478559912,0.362887369300815,0.0],
180         [-0.481125226052722,0.272165525754372,0.833333336324662],
181         [0.962250448649416,1.360827634879555,-0.000000000000007],
182         [-0.577350269189626,0.816496580927727,0.000000000000006],
183         [0.288675134594807,0.816496580927728,-0.500000000000003],

```

```
184         [0.288675134594818,0.816496580927728,0.4999999999999997]
185     ]),
186     2: np.array([
187         [0.0,0.0,0.0],
188         [0.0,1.6666666666667998,0.0],
189         [1.443375673043679,0.8333333333351740,0.0],
190         [0.930175433779948,-0.055555555629319,-0.362887369385129],
191         [0.769800358897570,1.333333333333403,0.544331053988240],
192         [-0.096225044892638,0.833333333333999,0.544331053947160],
193         [0.769800358899221,0.333333333331773,0.544331053992464],
194         [0.481125224341178,0.833333333333999,-0.272165526957415]
195     ]),
196     3: np.array([
197         [0.0,0.0,0.0],
198         [0.0,1.666666667005101,0.0],
199         [-1.443375672934393,0.833333333414405,0.0],
200         [0.032075015577703,-0.055555555417941,-0.997940266790499],
201         [-0.769800358911742,1.33333333345365,-0.544331054028353],
202         [-0.481125224242490,0.833333333502287,0.272165526932724],
203         [0.096225044986524,0.833333333502894,-0.544331053995201],
204         [-0.769800358872755,0.333333333438965,-0.544331053983922],
205     ]),
206     4: np.array([
207         [0.0,0.0,0.0],
208         [0.0,-1.732050817167331,0.0],
209         [1.290994466278070,-1.154700576756221,0.0],
210         [0.129099444946515,-0.288675138370196,0.948683301241860],
211         [-0.387298330204278,-0.866025408583675,0.316227759165464],
212         [0.903696135855355,-0.288675142331924,0.316227744588897],
213         [0.387298347149651,-0.866025408583689,-0.316227783307896],
214         [0.516397782962325,-1.154700544543956,0.632455526960460]
215     ]),
216     5: np.array([
217         [0.0,0.0,0.0],
218         [0.0,1.632993161855452,0.0],
219         [-1.026400478559335,1.270105792554240,0.0],
220         [0.577350269189626,0.816496580927726,0.000000000000000],
221         [-0.962250448649376,0.272165526975909,0.000000000000000],
222         [-0.607737126005969,1.718940171910572,-0.789473686405863],
223         [-0.288675134594813,0.816496580927726,0.500000000000000],
224         [-0.288675134594813,0.816496580927726,-0.500000000000000]
225     ]),
226     6: np.array([
227         [0.0,0.0,0.0],
228         [0.0,-1.677941309409369,0.0],
229         [-1.417542541531493,-0.780132765613074,0.0],
230         [0.108997432281229,-0.838970654704685,-0.533148947579020],
231         [-0.943453589462634,0.036363815680489,0.329504171048663],
232         [-0.767092038848638,-1.321119219647785,-0.533148947570053],
233         [-0.745915762131459,-0.322572875164488,-0.582714695238914],
234         [-0.461720532609713,-0.838970654704685,0.287997205406323]
```

```
235     ]),
236     7: np.array([
237         [0.0,0.0,0.0],
238         [0.0,-1.732050808690269,0.0],
239         [-1.632993162893537,-0.577350266978568,0.0],
240         [-0.000000000618892,-0.866025404345123,-0.50000000029555],
241         [-0.816496581442531,-0.288675133672735,0.50000000071310],
242         [-0.00000000014527,-0.866025404345117,0.500000000127100],
243         [-0.816496581096403,-0.288675133863199,-0.500000000169129],
244         [-0.816496582770094,-1.154700537873727,0.000000000104982]
245     ]),
246     8: np.array([
247         [0.0,0.0,0.0],
248         [0.0,1.732050807568878,0.0],
249         [-0.957427107756341,0.288675134594811,0.0],
250         [-0.029012942640569,1.635825762763941,0.994936676436827],
251         [-0.435194139889244,0.288675134594813,0.852802865422442],
252         [-0.261116483933549,0.866025403784439,-0.426401432711222],
253         [0.261116483933547,0.866025403784439,0.426401432711221],
254         [-0.696310623822793,1.154700538379251,0.426401432711220]
255     ]),
256     9: np.array([
257         [0.0,0.0,0.0],
258         [0.0,1.632993161855453,0.0],
259         [-0.962250448649376,0.272165526975909,0.0],
260         [0.481125224324696,1.360827634879546,0.833333333333344],
261         [-0.545275254236216,-0.090721842330956,0.833333333341596],
262         [0.577350269189626,0.816496580927726,-0.000000000000000],
263         [-0.288675134594813,0.816496580927726,-0.500000000000000],
264         [-0.288675134594813,0.816496580927726,0.500000000000000]
265     ]),
266     10: np.array([
267         [0.0,0.0,0.0],
268         [0.0,1.632993163699450,0.0],
269         [-0.962250449360077,0.272165525396529,0.0],
270         [0.481125224750552,1.360827638529060,-0.833333334082967],
271         [-0.545275266042522,-0.090721868947916,0.833333356462651],
272         [0.577350269193707,0.816496581849730,0.00000000005621],
273         [-0.288675134590899,0.816496581849730,-0.49999999995033],
274         [-0.288675134590901,0.816496581849730,0.500000000005347]
275     ]),
276     11: np.array([
277         [0.0,0.0,0.0],
278         [0.0,-1.414213562373095,0.0],
279         [1.000000000000631,0.00000000000422,0.0],
280         [-1.000000000000002,-1.414213562373096,-0.000000000000001],
281         [0.499999999999999,-0.707106781186547,0.500000000000001],
282         [-0.499999999999999,-0.707106781186547,-0.500000000000002],
283         [0.500000000000001,-0.707106781186547,-0.499999999999999],
284         [-0.500000000000001,-0.707106781186547,0.499999999999999]
285     ]),
```

```
286     12: np.array([
287         [0.0,0.0,0.0],
288         [0.0,1.719939178003135,0.0],
289         [0.956811833812873,0.290707954304311,0.0],
290         [0.088325741592312,1.429231223651817,-0.952726324390386],
291         [0.261284393834782,0.859969589001574,0.438386554924369],
292         [0.375151161705916,0.485198037994381,-0.789838255574746],
293         [-0.412394850520668,0.859969589001576,-0.300637311912409],
294         [0.821096870451597,1.234741139895464,-0.300637311926518]
295     ]),
296 },
297 }
```

Appendix D

Low temperature simulations from Chapter 6

The experiments in Section 6.2 and 6.3 were repeated with $T = 30K$ to understand the effects of the thermal energy on the hysteresis loop shapes. Figure D.1 shows the dynamic hysteresis loops for different clusters and the SPL is plotted in Figure D.2 as a function of the applied field amplitude.




Figure D.1: Dynamic major hysteresis loops for chain and Arkus cluster geometries with 1-8 particles at temperature $T = 30\text{K}$. Results were obtained by stochastic simulation of coupled sLLG equations using Magpy. The vertical and horizontal axes are shared between all figures. Results are shown for randomly oriented clusters with random anisotropy axes and clusters aligned with the applied field with aligned anisotropy axes.

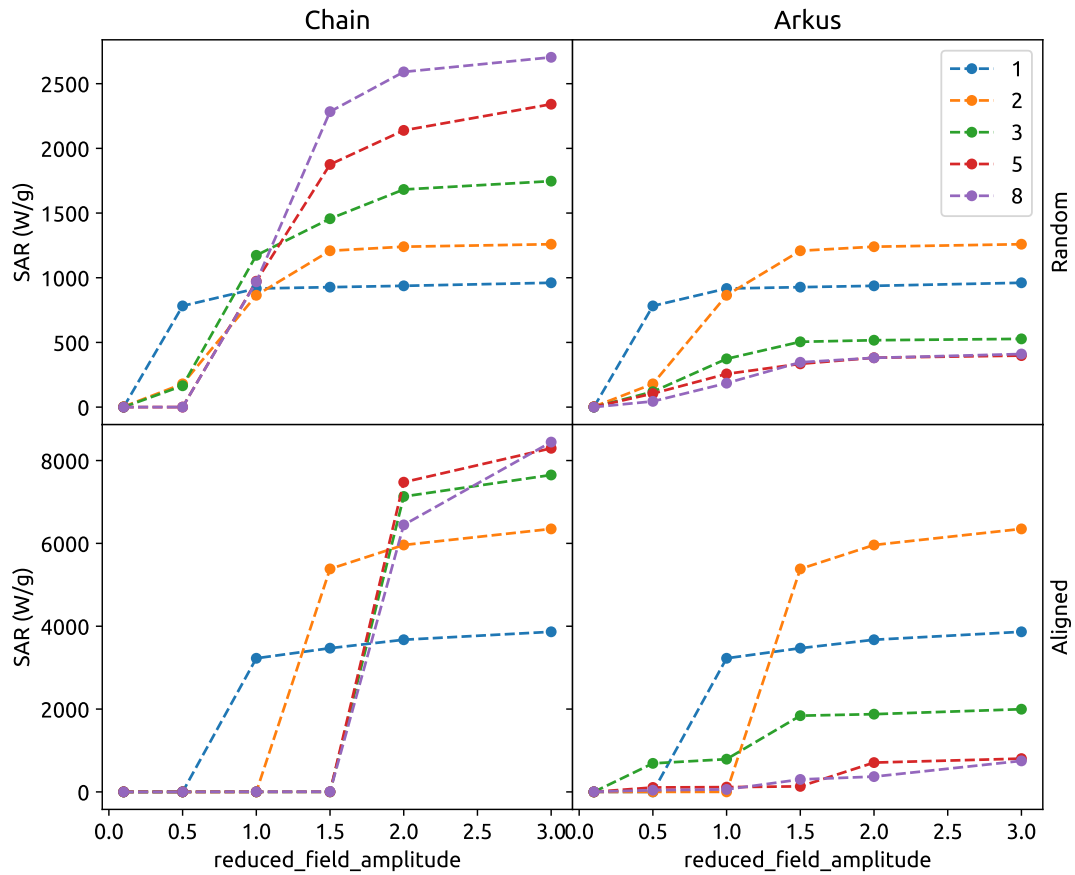


Figure D.2: SPL obtained from numerical integration of hysteresis loops at temperature $T = 30\text{K}$. The hysteresis loops were simulated for Chain and Arkus geometries of various sizes both randomly-oriented and aligned with the externally applied field. The results are shown as a function of the alternating magnetic field reduced amplitude (normalised by the anisotropy field H_k for a single particle).

References

- [1] Eric C Abenojar, Sameera Wickramasinghe, Jesbaniris Bas-Concepcion, and Anna Cristina S Samia. Structural effects on the magnetic hyperthermia properties of iron oxide nanoparticles. *Progress in Natural Science: Materials International*, 26(5):440–448, 2016.
- [2] M Ableidinger and E Buckwar. Weak stochastic Runge–Kutta Munthe-Kaas methods for finite spin ensembles. *Applied Numerical Mathematics*, 118:50–63, 2017.
- [3] Sebastian Ament, Nikhil Rangarajan, Arun Parthasarathy, and Shaloo Rakheja. Solving the stochastic Landau-Lifshitz-Gilbert-Slonczewski equation for monodomain nanomagnets: A survey and analysis of numerical techniques. *arXiv preprint arXiv:1607.04596*, 2016.
- [4] Sebastian Ament, Nikhil Rangarajan, and Shaloo Rakheja. A practical guide to solving the stochastic Landau-Lifshitz-Gilbert-Slonczewski equation for macrospin dynamics. *arXiv preprint arXiv:1607.04596*, 2016.
- [5] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999. ISBN 0-89871-447-8 (paperback).
- [6] Irene Andreu, Eva Natividad, Laura Solozábal, and Olivier Roubeau. Nano-objects for addressing the control of nanoparticle arrangement and performance in magnetic hyperthermia. *ACS nano*, 9(2):1408–1419, 2015.
- [7] Natalie Arkus, Vinothan N Manoharan, and Michael P Brenner. Minimal energy clusters of hard spheres with short range attractions. *Physical review letters*, 103(11):118303, 2009.
- [8] Svante Arrhenius. Über die dissociationswärme und den einfluss der temperatur auf den dissociationsgrad der elektrolyte. *Zeitschrift für physikalische Chemie*, 4(1):96–116, 1889.
- [9] William J Atkinson, Ivan A Brezovich, and Dev P Chakraborty. Usable frequencies in hyperthermia with thermal seeds. *IEEE Transactions on Biomedical Engineering*, (1):70–75, 1984.

- [10] Ki Hyun Bae, Mihyun Park, Min Jae Do, Nohyun Lee, Ji Hyun Ryu, Gun Woo Kim, CheolGi Kim, Tae Gwan Park, and Taeghwan Hyeon. Chitosan oligosaccharide-stabilized ferrimagnetic iron oxide nanocubes for magnetically modulated cancer hyperthermia. *ACS nano*, 6(6):5266–5273, 2012.
- [11] Lubomir Banas, Zdzislaw Brzezniak, Mikhail Neklyudov, and Andreas Prohl. *Stochastic ferromagnetism: analysis and numerics*, volume 58. Walter de Gruyter, 2013.
- [12] Lubomir Banas, Zdzislaw Brzezniak, and Andreas Prohl. Computational studies for the stochastic Landau–Lifshitz–Gilbert equation. *SIAM Journal on Scientific Computing*, 35(1):B62–B81, 2013.
- [13] L’ubomír Bañas. Numerical methods for the Landau-Lifshitz-Gilbert equation. In *International Conference on Numerical Analysis and Its Applications*, pages 158–165. Springer, 2004.
- [14] Subhankar Bedanta and Wolfgang Kleemann. Supermagnetism. *Journal of Physics D: Applied Physics*, 42(1):013001, 2008.
- [15] Marijan Beg, Ryan A Pepper, and Hans Fangohr. User interfaces for computational science: A domain specific language for OOMMF embedded in python. *AIP Advances*, 7(5):056025, 2017.
- [16] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D.S. Seljebotn, and K. Smith. Cython: The best of both worlds. *Computing in Science Engineering*, 13(2): 31–39, 2011. ISSN 1521-9615.
- [17] Jaber Beik, Ziaeddin Abed, Fatemeh S Ghoreishi, Samira Hosseini-Nami, Saeed Mehrzadi, Ali Shakeri-Zadeh, and S Kamran Kamrava. Nanotechnology in hyperthermia cancer therapy: From fundamental principles to advanced applications. *Journal of Controlled Release*, 235:205–221, 2016.
- [18] D V Berkov and N L Gorn. **Thermally activated processes in magnetic systems consisting of rigid dipoles: equivalence of the Itô and Stratonovich stochastic calculus.** *Journal of Physics: Condensed Matter*, 14(13):L281, 2002.
- [19] DV Berkov and NL Gorn. Susceptibility of the disordered system of fine magnetic particles: a Langevin-dynamics study. *Journal of Physics: Condensed Matter*, 13(41):9369, 2001.
- [20] Luis C Branquinho, Marcus S Carrião, Anderson S Costa, Nicholas Zufelato, Marcelo H Sousa, Ronei Miotto, Robert Ivkov, and Andris F Bakuzis. Effect of magnetic dipolar interactions on nanoparticle heating efficiency: Implications for cancer hyperthermia. *Scientific reports*, 3:2887, 2013.

- [21] Jr. Brown, W. Thermal fluctuation of fine ferromagnetic particles. *Magnetics, IEEE Transactions on*, 15(5):1196–1208, Sep 1979. ISSN 0018-9464.
- [22] Willam Fuller Brown. Thermal fluctuations of a single-domain particle. *Physical Review*, 130(5):1677–1686, June 1963.
- [23] Kevin Burrage and PM Burrage. Order conditions of stochastic Runge–Kutta methods by b-series. *SIAM Journal on Numerical Analysis*, 38(5):1626–1646, 2000.
- [24] Kevin Burrage, PM Burrage, and Tianhai Tian. Numerical methods for strong solutions of stochastic differential equations: an overview. In *Proceedings of The Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 460, pages 373–402. The Royal Society, 2004.
- [25] Karen Butter, PHH Bomans, PM Frederik, GJ Vroege, and AP Philipse. Direct observation of dipolar chains in iron ferrofluids by cryogenic electron microscopy. *Nature materials*, 2(2):88–91, 2003.
- [26] D Cabrera, A Lak, T Yoshida, ME Materia, D Ortega, F Ludwig, P Guardia, A Sathya, T Pellegrino, and FJ Teran. Unraveling viscosity effects on the hysteresis losses of magnetic nanocubes. *Nanoscale*, 9(16):5094–5101, 2017.
- [27] Julian Carrey, Boubker Mehdaoui, and Marc Respaud. Simple models for dynamic hysteresis loop calculations of magnetic single-domain nanoparticles: Application to magnetic hyperthermia optimization. *Journal of Applied Physics*, 109(8):083921, 2011.
- [28] Marcus S Carrião, Victor RR Aquino, Gabriel T Landi, Ediron L Verde, Marcelo H Sousa, and Andris F Bakuzis. Giant-spin nonlinear response theory of magnetic nanoparticle hyperthermia: A field dependence study. *Journal of Applied Physics*, 121(17):173901, 2017.
- [29] Daniela Caruntu, Gabriel Caruntu, and Charles J O’Connor. Magnetic properties of variable-sized Fe₃O₄ nanoparticles synthesized from non-aqueous homogeneous solutions of polyols. *Journal of Physics D: Applied Physics*, 40(19):5801, 2007.
- [30] David Chang, May Lim, Jeroen ACM Goos, Ruirui Qiao, Yun Yee Ng, Friederike M Mansfeld, Michael Jackson, Thomas P Davis, and Maria Kavallaris. Biologically targeted magnetic hyperthermia: Potential and limitations. *Frontiers in pharmacology*, 9, 2018.
- [31] X. Z. Cheng, M. B. A. Jalil, Hwee Kuan Lee, and Yutaka Okabe. [Mapping the Monte Carlo scheme to Langevin dynamics: A Fokker-Planck approach.](#) *Phys. Rev. Lett.*, 96:067208, Feb 2006.

- [32] Soshin Chikazumi and Chad D Graham. *Physics of Ferromagnetism 2e*, volume 94. Oxford University Press on Demand, 2009.
- [33] Katrina F Chu and Damian E Dupuy. Thermal ablation of tumours: biological mechanisms and advances in therapy. *Nature Reviews Cancer*, 14(3):199–208, 2014.
- [34] O Chubykalo-Fesenko and RW Chantrell. Multidimensional energy barrier distributions of interacting magnetic particles evaluated at different magnetization states. *Journal of applied physics*, 97(10):10J315, 2005.
- [35] Olga Chugreeva and Christof Melcher. Strong solvability of regularized stochastic Landau–Lifshitz–Gilbert equation. *IMA Journal of Applied Mathematics*, 83(2):261–282, 2018.
- [36] Ivan Cimrak. A survey on the numerics and computations for the Landau-Lifshitz equation of micromagnetism. *Archives of Computational Methods in Engineering*, 15(3):1–37, 2007.
- [37] W T Coffey, D S F Crothers, J L Dormann, L J Geoghegan, E C Kennedy, and W Wernsdorfer. **Range of validity of Kramers escape rates for non-axially symmetric problems in superparamagnetic relaxation.** *Journal of Physics: Condensed Matter*, 10(40):9093, 1998.
- [38] William T. Coffey and Yuri P. Kalmykov. **Thermal fluctuations of magnetic nanoparticles: Fifty years after Brown.** *Journal of Applied Physics*, 112(12):–, 2012.
- [39] Ivan Conde-Leboran, Daniel Baldomir, Carlos Martinez-Boubeta, Oksana Chubykalo-Fesenko, Marıa del Puerto Morales, Gorika Salas, David Cabrera, Julio Camarero, Francisco J Teran, and David Serantes. A single picture explains diversity of hyperthermia response of magnetic nanoparticles. *The Journal of Physical Chemistry C*, 119(27):15698–15706, 2015.
- [40] Diego F Coral, Pedro Mendoza Zelis, Marzia Marciello, Marıa del Puerto Morales, Aldo Craievich, Francisco H Sanchez, and Marcela B Fernandez van Raap. Effect of nanoclustering and dipolar interactions in heat generation for magnetic hyperthermia. *Langmuir*, 32(5):1201–1213, 2016.
- [41] David Cortes-Ortuno, Weiwei Wang, Ryan Pepper, Marc-Antonio Bisotti, Thomas Kluyver, Mark Vousden, and Hans Fangohr. **Fidimag v2.0**, 2016.
- [42] NR Datta, AK Bose, HK Kapoor, and S Gupta. Head and neck cancers: results of thermoradiotherapy versus radiotherapy. *International Journal of Hyperthermia*, 6(3):479–486, 1990.
- [43] NR Datta, S Gomez Ordonez, US Gaipl, MM Paulides, Hans Crezee, J Gellermann, D Marder, E Puric, and S Bodis. Local hyperthermia combined with radiotherapy

- and/or chemotherapy: Recent advances and promises for the future. *Cancer treatment reviews*, 41(9):742–753, 2015.
- [44] Alison E Deatsch and Benjamin A Evans. Heating efficiency in magnetic nanoparticle hyperthermia. *Journal of Magnetism and Magnetic Materials*, 354:163–172, 2014.
- [45] PM Déjardin, DSF Crothers, WT Coffey, and DJ McCarthy. Interpolation formula between very low and intermediate-to-high damping Kramers escape rates for single-domain ferromagnetic particles. *Physical Review E*, 63(2):021102, 2001.
- [46] Riccardo Di Corato, Ana Espinosa, Lenaic Lartigue, Mickael Tharaud, Sophie Chat, Teresa Pellegrino, Christine Ménager, Florence Gazeau, and Claire Wilhelm. Magnetic hyperthermia efficiency in the cellular environment for different nanoparticle designs. *Biomaterials*, 35(24):6400–6411, 2014.
- [47] Michael J Donahue. OOMMF user’s guide, version 1.0. -6376, 1999.
- [48] Jean-Louis Dormann, Dino Fiorani, and Elisabeth Tronc. Magnetic relaxation in fine-particle systems. *Advances in chemical physics*, 98:283–494, 1997.
- [49] JL Dormann, D Fiorani, and E Tronc. On the models for interparticle interactions in nanoparticle assemblies: comparison with experimental results. *Journal of Magnetism and Magnetic Materials*, 202(1):251–267, 1999.
- [50] S Dutz, R Hergt, J Mürbe, R Müller, M Zeisberger, W Andrä, J Töpfer, and ME Bellemann. Hysteresis losses of magnetic nanoparticle powders in the single domain size range. *Journal of Magnetism and Magnetic Materials*, 308(2):305–312, 2007.
- [51] Silvio Dutz and Rudolf Hergt. Magnetic particle hyperthermia—a promising tumour therapy? *Nanotechnology*, 25(45):452001, 2014.
- [52] M d’Aquino, C Serpico, G Coppola, ID Mayergoyz, and G Bertotti. Midpoint numerical technique for stochastic Landau-Lifshitz-Gilbert dynamics. *Journal of applied physics*, 99(8):08B905, 2006.
- [53] Massimiliano d’Aquino, Claudio Serpico, and Giovanni Miano. **Geometrical integration of Landau-Lifshitz-Gilbert equation based on the mid-point rule.** *Journal of Computational Physics*, 209(2):730 – 753, 2005. ISSN 0021-9991.
- [54] Alexander S Eggeman, Sara A Majetich, Dorothy Farrell, and Quentin A Pankhurst. Size and concentration effects on high frequency hysteresis of iron oxide nanoparticles. *IEEE transactions on magnetics*, 43(6):2451–2453, 2007.
- [55] Olle Eriksson, Anders Bergman, Lars Bergqvist, and Johan Hellsvik. *Atomistic spin dynamics: foundations and applications.* Oxford university press, 2017.

- [56] Richard FL Evans, Weijia J Fan, Phanwadee Chureemart, Thomas A Ostler, Matthew OA Ellis, and Roy W Chantrell. Atomistic spin model simulations of magnetic nanomaterials. *Journal of Physics: Condensed Matter*, 26(10):103202, 2014.
- [57] Richard Francis L Evans, Denise Hinzke, Unai Atxitia, Ulrich Nowak, Roy W Chantrell, and Oksana Chubykalo-Fesenko. Stochastic form of the Landau-Lifshitz-Bloch equation. *Physical Review B*, 85(1):014433, 2012.
- [58] Hans Fangohr. A comparison of C, MATLAB, and Python as teaching languages in engineering. *Computational Science-ICCS 2004*, pages 1210–1217, 2004.
- [59] Hans Fangohr, Maximilian Albert, and Matteo Franchin. Nmag micromagnetic simulation tool-software engineering lessons learned. In *Software Engineering for Science (SE4Science), IEEE/ACM International Workshop on*, pages 1–7. IEEE, 2016.
- [60] J Fidler, RW Chantrell, T Schrefl, MA Wongsam, and J Fidler. Micromagnetics: basic principles. In *Encyclopedia of Materials: Science and Technology*, page 5642. Elsevier Science Ltd., 2001.
- [61] Josef Fidler and Thomas Schrefl. Micromagnetic modelling-the current state of the art. *Journal of Physics D: Applied Physics*, 33(15):R135, 2000.
- [62] Thomas Fischbacher, Matteo Franchin, Giuliano Bordignon, and Hans Fangohr. A systematic approach to multiphysics extensions of finite-element-based micromagnetic simulations: Nmag. *Magnetics, IEEE Transactions on*, 43(6):2896–2898, 2007.
- [63] Jean-Paul Fortin, Claire Wilhelm, Jacques Servais, Christine Ménager, Jean-Claude Bacri, and Florence Gazeau. Size-sorted anionic iron oxide nanomagnets as colloidal mediators for magnetic hyperthermia. *Journal of the American Chemical Society*, 129(9):2628–2635, 2007.
- [64] Rong Fu, Yuying Yan, Clive Roberts, Zeyu Liu, and Yiyi Chen. The role of dipole interactions in hyperthermia heating colloidal clusters of densely-packed superparamagnetic nanoparticles. *Scientific reports*, 8(1):4704, 2018.
- [65] JL Garcia-Palacios. On the statics and dynamics of magnetoanisotropic nanoparticles. *Advances in Chemical Physics*, 112:1–210, 2000.
- [66] José Luis García-Palacios and Francisco J Lázaro. Langevin-dynamics study of the dynamical properties of small magnetic particles. *Physical Review B*, 58(22):14937, 1998.
- [67] Thomas C Gard. *Introduction to stochastic differential equations*. M. Dekker, 1988.

- [68] Crispin Gardiner. *Stochastic Methods: A Handbook for the Natural and Social Sciences*, volume 13 of *Springer Series in Synergetics*. Springer, 4th edition, 2009.
- [69] Florence Gazeau, Michael Lévy, and Claire Wilhelm. Optimizing magnetic nanoparticle design for nanothermotherapy. 2008.
- [70] James E Gentle, Wolfgang Karl Härdle, and Yuichi Mori. *Handbook of computational statistics: concepts and methods*. Springer Science & Business Media, 2012.
- [71] LJ Geoghegan, WT Coffey, and B Mulligan. **Differential recurrence relations for non-axially symmetric rotational Fokker-Planck equations**. *Advances in Chemical Physics*, 100:475–641, 1997.
- [72] J Willard Gibbs. *Elementary principles in statistical mechanics*. Courier Corporation, 2014.
- [73] T.L. Gilbert. A phenomenological theory of damping in ferromagnetic materials. *Magnetics, IEEE Transactions on*, 40(6):3443–3449, Nov 2004. ISSN 0018-9464.
- [74] RK Gilchrist, Richard Medal, William D Shorey, Russell C Hanselman, John C Parrott, and C Bruce Taylor. Selective inductive heating of lymph nodes. *Annals of surgery*, 146(4):596, 1957.
- [75] Daniel T Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of computational physics*, 22(4):403–434, 1976.
- [76] Daniel T Gillespie. Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry*, 81(25):2340–2361, 1977.
- [77] Daniel T Gillespie. Monte Carlo simulation of random walks with residence time dependent transition probability rates. *Journal of Computational Physics*, 28(3):395–407, 1978.
- [78] Daniel T Gillespie. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications*, 188(1):404–425, 1992.
- [79] Hagen Gilsing and Tony Shardlow. Sdelab: A package for solving stochastic differential equations in matlab. *Journal of Computational and Applied Mathematics*, 205(2):1002–1018, 2007.
- [80] Andrew J Giustini, Alicia A Petryk, Shiraz M Cassim, Jennifer A Tate, Ian Baker, and P Jack Hoopes. Magnetic nanoparticle hyperthermia in cancer treatment. *Nano Life*, 1(01n02):17–32, 2010.

- [81] GF Goya, TS Berquo, FC Fonseca, and MP Morales. Static and dynamic magnetic properties of spherical magnetite nanoparticles. *Journal of Applied Physics*, 94(5): 3520–3528, 2003.
- [82] Carl Graham and Denis Talay. *Stochastic simulation and Monte Carlo methods: mathematical foundations of stochastic simulation*, volume 68. Springer Science & Business Media, 2013.
- [83] Ajay Kumar Gupta and Mona Gupta. Synthesis and surface engineering of iron oxide nanoparticles for biomedical applications. *Biomaterials*, 26(18):3995–4021, 2005.
- [84] Christian Haase and Ulrich Nowak. Role of dipole-dipole interactions for hyperthermia heating of magnetic nanoparticle ensembles. *Physical Review B*, 85(4): 045435, 2012.
- [85] L Hafström, CM Rudenstam, Erik Blomquist, C Ingvar, PE Jönsson, B Lagerlöf, C Lindholm, U Ringborg, G Westman, and L Ostrup. Regional hyperthermic perfusion with melphalan after surgery for recurrent malignant melanoma of the extremities. Swedish melanoma study group. *Journal of clinical oncology*, 9(12): 2091–2094, 1991.
- [86] Peter Hänggi, Peter Talkner, and Michal Borkovec. **Reaction-rate theory: fifty years after Kramers**. *Rev. Mod. Phys.*, 62:251–341, Apr 1990.
- [87] Jonathan David Hannay. *Computational Simulations of Thermally Activated Magnetisation Dynamics at High Frequencies*. PhD thesis, School of Informatics, University of Wales, University of Wales, Bangor, Gwynedd, UK, July 2001.
- [88] Graeme Henkelman, Blas P Uberuaga, and Hannes Jónsson. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *The Journal of chemical physics*, 113(22):9901–9904, 2000.
- [89] Rudolf Hergt, Silvio Dutz, and Michael Röder. Effects of size distribution on hysteresis losses of magnetic nanoparticles for hyperthermia. *Journal of Physics: Condensed Matter*, 20(38):385214, 2008.
- [90] Rudolf Hergt, Silvio Dutz, and Matthias Zeisberger. Validity limits of the Néel relaxation model of magnetic nanoparticles for hyperthermia. *Nanotechnology*, 21(1):015706, 2009.
- [91] Aziliz Hervault and Nguyễn Thị Kim Thanh. Magnetic nanoparticle-based therapeutic agents for thermo-chemotherapy treatment of cancer. *Nanoscale*, 6(20): 11553–11573, 2014.
- [92] R Hiergeist, W Andrä, N Buske, R Hergt, I Hilger, U Richter, and W Kaiser. Application of magnetite ferrofluids for hyperthermia. *Journal of Magnetism and Magnetic Materials*, 201(1-3):420–422, 1999.

- [93] Alan C Hindmarsh, Peter N Brown, Keith E Grant, Steven L Lee, Radu Serban, Dan E Shumaker, and Carol S Woodward. SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Transactions on Mathematical Software (TOMS)*, 31(3):363–396, 2005.
- [94] Matthew D Hoffman and Andrew Gelman. The no-u-turn sampler: adaptively setting path lengths in hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- [95] O. Hovorka, J. Barker, G. Friedman, and R. W. Chantrell. **Role of geometrical symmetry in thermally activated processes in clusters of interacting dipolar moments.** *Phys. Rev. B*, 89:104410, Mar 2014.
- [96] Ondrej Hovorka. Thermal activation in statistical clusters of magnetic nanoparticles. *Journal of Physics D: Applied Physics*, 50(4):044004, 2017.
- [97] Alex Hubert and Rudolf Schäfer. *Magnetic domains: the analysis of magnetic microstructures*. Springer Science & Business Media, 2008.
- [98] Òscar Iglesias and Amílcar Labarta. Magnetic relaxation in terms of microscopic energy barriers in a model of dipolar interacting nanoparticles. *Physical Review B*, 70(14):144401, 2004.
- [99] Silvana Ilie and Alexandra Teslya. An adaptive stepsize method for the chemical Langevin equation. *The Journal of chemical physics*, 136(18):184101, 2012.
- [100] Song Jiang, Qiangchang Ju, and Huaqiao Wang. Martingale weak solutions of the stochastic Landau–Lifshitz–Bloch equation. *Journal of Differential Equations*, 2018.
- [101] M Johannsen, U Gneveckow, L Eckelt, A Feussner, N Waldöfner, R Scholz, S Deger, P Wust, SA Loening, and A Jordan. Clinical hyperthermia of prostate cancer using magnetic nanoparticles: presentation of a new interstitial technique. *International journal of hyperthermia*, 21(7):637–647, 2005.
- [102] Manfred Johannsen, Uwe Gneveckow, Burghard Thiesen, Kasra Taymoorian, Chie Hee Cho, Norbert Waldöfner, Regina Scholz, Andreas Jordan, Stefan A Loening, and Peter Wust. Thermotherapy of prostate cancer using magnetic nanoparticles: feasibility, imaging, and three-dimensional temperature distribution. *European urology*, 52(6):1653–1662, 2007.
- [103] Manfred Johannsen, Burghard Thiesen, Peter Wust, and Andreas Jordan. Magnetic nanoparticle hyperthermia for prostate cancer. *International Journal of Hyperthermia*, 26(8):790–795, 2010.
- [104] Petra E Jonsson. Superparamagnetism and spin glass dynamics of interacting magnetic nanoparticle systems. *Advances in Chemical Physics*, 128:191–248, 2004.

- [105] A Jordan, P Wust, H Föhlin, W John, A Hinz, and R Felix. Inductive heating of ferrimagnetic particles and magnetic fluids: physical evaluation of their potential for hyperthermia. *International Journal of Hyperthermia*, 9(1):51–68, 1993.
- [106] Andreas Jordan, Regina Scholz, Peter Wust, Horst Föhling, and Roland Felix. Magnetic fluid hyperthermia (mfh): Cancer treatment with ac magnetic field induced excitation of biocompatible superparamagnetic nanoparticles. *Journal of Magnetism and Magnetic Materials*, 201(1):413–419, 1999.
- [107] Lina Kafrouni and Oumarou Savadogo. Recent progress on magnetic nanoparticles for magnetic hyperthermia. *Progress in biomaterials*, 5(3-4):147–160, 2016.
- [108] Yu. P. Kalmykov. **Evaluation of the smallest nonvanishing eigenvalue of the Fokker-Planck equation for the Brownian motion in a potential. ii. the matrix continued fraction approach.** *Phys. Rev. E*, 62:227–236, Jul 2000.
- [109] Yuri P Kalmykov, William T Coffey, Unai Atxitia, Oksana Chubykalo-Fesenko, Pierre-Michel Déjardin, and Roy W Chantrell. Damping dependence of the reversal time of the magnetization of single-domain ferromagnetic particles for the Néel-Brown model: Langevin dynamics simulations versus analytic results. *Physical Review B*, 82(2):024412, 2010.
- [110] Ganeshlenin Kandasamy and Dipak Maity. Recent advances in superparamagnetic iron oxide nanoparticles (spions) for in vitro and in vivo cancer nanotheranostics. *International journal of pharmaceuticals*, 496(2):191–218, 2015.
- [111] Charles FF Karney. Quaternions in molecular modeling. *Journal of Molecular Graphics and Modelling*, 25(5):595–604, 2007.
- [112] Natalia Kazantseva, Denise Hinzke, Ulrich Nowak, Roy W Chantrell, Unai Atxitia, and Oksana Chubykalo-Fesenko. Towards multiscale modeling of magnetic materials: Simulations of FePt. *Physical Review B*, 77(18):184428, 2008.
- [113] Amit P Khandhar, R Matthew Ferguson, and Kannan M Krishnan. Monodispersed magnetite nanoparticles optimized for magnetic fluid hyperthermia: Implications in biological systems. *Journal of applied physics*, 109(7):07B310, 2011.
- [114] Peter E Kloeden and Eckhard Platen. *Higher-order implicit strong numerical schemes for stochastic differential equations*, volume 66. Springer, 1992.
- [115] Peter E. Kloeden and Eckhard Platen. *Numerical Solution of Stochastic Differential Equations*, volume 23 of *Stochastic Modelling and Applied Probability*. Springer-Verlag Berlin Heidelberg, 1992.
- [116] Arati G Kolhatkar, Andrew C Jamison, Dmitri Litvinov, Richard C Willson, and T Randall Lee. Tuning the magnetic properties of nanoparticles. *International journal of molecular sciences*, 14(8):15977–16009, 2013.

- [117] Susanne Kossatz, Julia Grandke, Pierre Couleaud, Alfonso Latorre, Antonio Aires, Kieran Crosbie-Staunton, Robert Ludwig, Heidi Dähring, Volker Ettelt, Ana Lazaro-Carrillo, et al. Efficient treatment of breast cancer xenografts with multifunctionalized iron oxide nanoparticles combining magnetic hyperthermia and anti-cancer drug delivery. *Breast Cancer Research*, 17(1):66, 2015.
- [118] H.A. Kramers. **Brownian motion in a field of force and the diffusion model of chemical reactions**. *Physica*, 7(4):284 – 304, 1940. ISSN 0031-8914.
- [119] Challa SSR Kumar and Faruq Mohammad. Magnetic nanomaterials for hyperthermia-based therapy and controlled drug delivery. *Advanced drug delivery reviews*, 63(9):789–808, 2011.
- [120] M Lakshmanan. The fascinating world of the Landau–Lifshitz–Gilbert equation: an overview. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 369(1939):1280–1300, 2011.
- [121] Lale Landau and Es Lifshitz. On the theory of the dispersion of magnetic permeability in ferromagnetic bodies. *Phys. Z. Sowjetunion*, 8(153):101–114, 1935.
- [122] Gabriel T Landi. Role of dipolar interaction in magnetic hyperthermia. *Physical Review B*, 89(1):014403, 2014.
- [123] Gabriel T Landi and Mario J de Oliveira. Linear response theory for arbitrary periodic signals. *arXiv preprint arXiv:1208.3995*, 2012.
- [124] O Laslett, S Ruta, RW Chantrell, J Barker, G Friedman, and O Hovorka. Consistent energy barrier distributions in magnetic particle chains. *Physica B: Condensed Matter*, 2015.
- [125] Nohyun Lee, Dongwon Yoo, Daishun Ling, Mi Hyeon Cho, Taeghwan Hyeon, and Jinwoo Cheon. Iron oxide based nanoparticles for multimodal imaging and magnetoresponse therapy. *Chemical reviews*, 115(19):10637–10689, 2015.
- [126] Jonathan Leliaert, Jeroen Mulkers, Jonas De Clercq, Annelies Coene, M Dvornik, and Bartel Van Waeyenberge. Adaptively time stepping the stochastic Landau-Lifshitz-Gilbert equation at nonzero temperature: Implementation and validation in MuMax3. *AIP Advances*, 7(12):125010, 2017.
- [127] Jonathan Leliaert, Arne Vansteenkiste, Annelies Coene, Luc Dupré, and Bartel Van Waeyenberge. Vinamax: a macrospin simulation tool for magnetic nanoparticles. *Medical & biological engineering & computing*, 53(4):309–317, 2015.
- [128] Randall J LeVeque. *Finite difference methods for ordinary and partial differential equations: steady-state and time-dependent problems*. SIAM, 2007.
- [129] Z Li and S Zhang. Thermally assisted magnetization reversal in the presence of a spin-transfer torque. *Physical Review B*, 69(13):134416, 2004.

- [130] Anders Logg, Kent-Andre Mardal, and Garth Wells. *Automated solution of differential equations by the finite element method: The FEniCS book*, volume 84. Springer Science & Business Media, 2012.
- [131] Anders Logg, Garth N Wells, and Johan Hake. Dofin: A C++/python finite element library. *Automated Solution of Differential Equations by the Finite Element Method*, pages 173–225, 2012.
- [132] Keon Mahmoudi, Alexandros Bouras, Dominique Bozec, Robert Ivkov, and Constantinos Hadjipanayis. Magnetic hyperthermia therapy for the treatment of glioblastoma: a review of the therapy’s history, efficacy and application in humans. *International Journal of Hyperthermia*, pages 1–13, 2018.
- [133] Klaus Maier-Hauff, Ronny Rothe, Regina Scholz, Uwe Gneveckow, Peter Wust, Burghard Thiesen, Annelie Feussner, Andreas von Deimling, Norbert Waldoefner, Roland Felix, et al. Intracranial thermotherapy using magnetic nanoparticles combined with external beam radiotherapy: results of a feasibility study on patients with glioblastoma multiforme. *Journal of neuro-oncology*, 81(1):53–60, 2007.
- [134] Klaus Maier-Hauff, Frank Ulrich, Dirk Nestler, Hendrik Niehoff, Peter Wust, Burghard Thiesen, Helmut Orawa, Volker Budach, and Andreas Jordan. Efficacy and safety of intratumoral thermotherapy using magnetic iron-oxide nanoparticles combined with external beam radiotherapy on patients with recurrent glioblastoma multiforme. *Journal of neuro-oncology*, 103(2):317–324, 2011.
- [135] SA Majetich and M Sachan. Magnetostatic interactions in magnetic nanoparticle assemblies: energy, time and length scales. *Journal of Physics D: Applied Physics*, 39(21):R407, 2006.
- [136] Carlos Martinez-Boubeta, Konstantinos Simeonidis, Antonios Makridis, Makis Angelakeris, Oscar Iglesias, Pablo Guardia, Andreu Cabot, Lluís Yedra, Sonia Estradé, Francesca Peiró, et al. Learning from nature to improve the heat generation of iron-oxide nanoparticles for magnetic hyperthermia applications. *arXiv preprint arXiv:1304.1298*, 2013.
- [137] Maria Elena Materia, Pablo Guardia, Ayyappan Sathya, Manuel Pernia Leal, Roberto Marotta, Riccardo Di Corato, and Teresa Pellegrino. Mesoscale assemblies of iron oxide nanocubes as heat mediators and image contrast agents. *Langmuir*, 31(2):808–816, 2015.
- [138] B Mehdaoui, RP Tan, A Meffre, J Carrey, S Lachaize, B Chaudret, and M Respaud. Increase of magnetic hyperthermia efficiency due to dipolar interactions in low-anisotropy magnetic nanoparticles: Theoretical and experimental results. *Physical Review B*, 87(17):174419, 2013.

- [139] Boubker Mehdaoui, Anca Meffre, Julian Carrey, Sébastien Lachaize, Lise-Marie Lacroix, Michel Gougeon, Bruno Chaudret, and Marc Respaud. Optimal size of nanoparticles for magnetic hyperthermia: a combined theoretical and experimental study. *Advanced Functional Materials*, 21(23):4573–4581, 2011.
- [140] JH Mentink, MV Tretyakov, A Fasolino, MI Katsnelson, and Th Rasing. Stable and fast semi-implicit integration of the stochastic Landau–Lifshitz equation. *Journal of Physics: Condensed Matter*, 22(17):176001, 2010.
- [141] Kent G. Merryfield. [Notes for Fourier analysis. course notes for math 472.](#) 2009.
- [142] Grigori N Milstein, Yu M Repin, and Michael V Tretyakov. Numerical methods for stochastic systems preserving symplectic structure. *SIAM Journal on Numerical Analysis*, 40(4):1583–1604, 2002.
- [143] Sean M Morgan, Hweerin Sohn, and RH Victora. Use of trapezoidal waves and complementary static fields incident on magnetic nanoparticles to induce magnetic hyperthermia for therapeutic cancer treatment. *Journal of Applied Physics*, 109(7):07B305, 2011.
- [144] Sean M Morgan and RH Victora. Use of square waves incident on magnetic nanoparticles to induce magnetic hyperthermia for therapeutic cancer treatment. *Applied Physics Letters*, 97(9):093705, 2010.
- [145] Steen Mørup, Mikkel Fougth Hansen, and Cathrine Frandsen. Magnetic interactions between nanoparticles. *Beilstein journal of nanotechnology*, 1:182, 2010.
- [146] Cristina Munoz-Menendez, Ivan Conde-Leboran, Daniel Baldomir, Oksana Chubykalo-Fesenko, and David Serantes. The role of size polydispersity in magnetic fluid hyperthermia: average vs. local infra/over-heating effects. *Physical Chemistry Chemical Physics*, 17(41):27812–27820, 2015.
- [147] Cristina Munoz-Menendez, Ivan Conde-Leboran, David Serantes, Roy Chantrell, Oksana Chubykalo-Fesenko, and Daniel Baldomir. Distinguishing between heating power and hyperthermic cell-treatment efficacy in magnetic fluid hyperthermia. *Soft matter*, 12(43):8815–8818, 2016.
- [148] Cristina Munoz-Menendez, David Serantes, Juan M Ruso, and Daniel Baldomir. Towards improved magnetic fluid hyperthermia: major-loops to diminish variations in local heating. *Physical Chemistry Chemical Physics*, 19(22):14527–14532, 2017.
- [149] M Nabil, P Decuzzi, and P Zunino. Modelling mass and heat transfer in nano-based cancer hyperthermia. *Royal Society open science*, 2(10):150447, 2015.
- [150] Radford M Neal et al. Mcmc using hamiltonian dynamics. *Handbook of Markov Chain Monte Carlo*, 2(11), 2011.

- [151] Louis Néel. Théorie du traînage magnétique des ferromagnétiques en grains fins avec applications aux terres cuites. *Ann. géophys*, 5(2):99–136, 1949.
- [152] Dina Niculaes, Aidin Lak, George C Anyfantis, Sergio Marras, Oliver Laslett, Sahitya K Avugadda, Marco Cassani, David Serantes, Ondrej Hovorka, Roy Chantrell, et al. Asymmetric assembling of iron oxide nanocubes for improving magnetic hyperthermia performance. *ACS nano*, 11(12):12121–12133, 2017.
- [153] Ulrich Nowak, Roy W Chantrell, and EC Kennedy. Monte Carlo simulation with time step quantification in terms of Langevin dynamics. *Physical review letters*, 84(1):163, 2000.
- [154] Ihab M Obaidat, Bashar Issa, and Yousef Haik. Magnetic properties of magnetic nanoparticles for efficient hyperthermia. *Nanomaterials*, 5(1):63–89, 2015.
- [155] Travis E Oliphant. *A guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.
- [156] International Commission on Non-Ionizing Radiation Protection et al. Guidance on determining compliance of exposure to pulsed and complex non-sinusoidal waveforms below 100 khz with ICNIRP guidelines. *Health physics*, 84(3):383–387, 2003.
- [157] International Commission on Non-Ionizing Radiation Protection et al. ICNIRP statement on the “guidelines for limiting exposure to time-varying electric, magnetic, and electromagnetic fields (up to 300 ghz)”. *Health Physics*, 97(3):257–258, 2009.
- [158] Daniel Ortega and Quentin A Pankhurst. Magnetic hyperthermia. *Nanoscience*, 1(60):e88, 2013.
- [159] Jesús G Ovejero, David Cabrera, Julian Carrey, T Valdivielso, Goroka Salas, and Francisco J Teran. Effects of inter-and intra-aggregate magnetic dipolar interactions on the magnetic heating efficiency of iron oxide nanoparticles. *Physical Chemistry Chemical Physics*, 18(16):10954–10963, 2016.
- [160] Elio Alberto Périgo, Gauvin Hemery, Olivier Sandre, Daniel Ortega, Eneko Garaio, Fernando Plazaola, and Francisco Javier Teran. Fundamentals and advances in magnetic hyperthermia. *Applied Physics Reviews*, 2(4):041302, 2015.
- [161] ML Plumer, J van Lierop, BW Southern, and JP Whitehead. Micromagnetic simulations of interacting dipoles on an fcc lattice: application to nanoparticle assemblies. *Journal of Physics: Condensed Matter*, 22(29):296007, 2010.
- [162] Pankaj Poddar, Tamar Telem-Shafir, Tcipi Fried, and Gil Markovich. Dipolar interactions in two-and three-dimensional magnetic nanoparticle arrays. *Physical Review B*, 66(6):060403, 2002.
- [163] Paolo Politi and Maria Gloria Pini. Dipolar interaction between two-dimensional magnetic particles. *Physical Review B*, 66(21):214414, 2002.

- [164] William H Press. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007.
- [165] Yu L Raikher, VI Stepanov, and R Perzynski. Dynamic hysteresis of a superparamagnetic nanoparticle. *Physica B: Condensed Matter*, 343(1):262–266, 2004.
- [166] Karl Johan Åström and Richard M Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.
- [167] L Harivardhan Reddy, José L Arias, Julien Nicolas, and Patrick Couvreur. Magnetic nanoparticles: design and characterization, toxicity and biocompatibility, pharmaceutical and biomedical applications. *Chemical reviews*, 112(11):5818–5878, 2012.
- [168] Christian P Robert. *Monte carlo methods*. Wiley Online Library, 2004.
- [169] Harley F Rodrigues, Gustavo Capistrano, Francielli M Mello, Nicholas Zufelato, Elisângela Silveira-Lacerda, and Andris F Bakuzis. Precise determination of the heat delivery during in vivo magnetic nanoparticle hyperthermia with infrared thermography. *Physics in Medicine & Biology*, 62(10):4062, 2017.
- [170] Federico Romá, Leticia F. Cugliandolo, and Gustavo S. Lozano. **Numerical integration of the stochastic Landau-Lifshitz-Gilbert equation in generic time-discretization schemes**. *Phys. Rev. E*, 90:023203, Aug 2014.
- [171] Ronald E Rosensweig. Heating magnetic fluid with alternating magnetic field. *Journal of magnetism and magnetic materials*, 252:370–374, 2002.
- [172] Werner Rümelin. Numerical treatment of stochastic differential equations. *SIAM Journal on Numerical Analysis*, pages 604–613, 1982.
- [173] S Ruta, R Chantrell, and O Hovorka. Unified model of hyperthermia via hysteresis heating in systems of interacting magnetic nanoparticles. *Scientific reports*, 5, 2015.
- [174] S Ruta, O Hovorka, and R Chantrell. Unified model of hyperthermia via hysteresis heating in systems of interacting magnetic nanoparticles. *arXiv preprint arXiv:1412.3814*, 2014.
- [175] ME Sadat, Ronak Patel, Jason Sookoor, Sergey L Bud’ko, Rodney C Ewing, Jiaming Zhang, Hong Xu, Yilong Wang, Giovanni M Pauletti, David B Mast, et al. Effect of spatial confinement on magnetic hyperthermia via dipolar interactions in Fe₃O₄ nanoparticles for biomedical applications. *Materials Science and Engineering: C*, 42:52–63, 2014.
- [176] Klaus Schmitz Abe. *Pricing exotic options using improved strong convergence*. PhD thesis, University of Oxford, 2008.

- [177] Klaus Schmitz Abe and William T Shaw. Measure order of convergence without an exact solution, Euler vs Milstein scheme. *International Journal of Pure and Applied Mathematics*, 24(3):365–381, 2005.
- [178] Werner Scholz, Thomas Schrefl, and Josef Fidler. Micromagnetic simulation of thermally activated switching in fine particles. *Journal of Magnetism and Magnetic Materials*, 233(3):296–304, 2001.
- [179] J. Schratzberger, J. Lee, M. Fuger, J. Fidler, G. Fiedler, T. Schrefl, and D. Suess. **Validation of the transition state theory with Langevin-dynamics simulations.** *Journal of Applied Physics*, 108(3):–, 2010.
- [180] D Serantes, D Baldomir, C Martinez-Boubeta, K Simeonidis, M Angelakeris, E Natividad, M Castro, A Mediano, D-X Chen, A Sanchez, et al. Influence of dipolar interactions on hyperthermia properties of ferromagnetic particles. *Journal of Applied Physics*, 108(7):073918, 2010.
- [181] David Serantes, Konstantinos Simeonidis, Makis Angelakeris, Oksana Chubykalo-Fesenko, Marzia Marciello, Maria Del Puerto Morales, Daniel Baldomir, and Carlos Martinez-Boubeta. Multiplying magnetic hyperthermia response by nanoparticle assembling. *The Journal of Physical Chemistry C*, 118(11):5927–5934, 2014.
- [182] Rhythm R Shah, Todd P Davis, Amanda L Glover, David E Nikles, and Christopher S Brazel. Impact of magnetic field parameters and iron oxide nanoparticle properties on heat generation for use in magnetic hyperthermia. *Journal of magnetism and magnetic materials*, 387:96–106, 2015.
- [183] Sigurd Skogestad and Ian Postlethwaite. *Multivariable feedback control: analysis and design*, volume 2. Wiley New York, 2007.
- [184] Penny K Sneed, Paul R Stauffer, Michael W McDermott, Chris J Diederich, Kathleen R Lamborn, Michael D Prados, Susan Chang, Keith A Weaver, Laura Spry, Mary K Malec, et al. Survival benefit of hyperthermia in a prospective randomized trial of brachytherapy boost±hyperthermia for glioblastoma multiforme. *International Journal of Radiation Oncology* Biology* Physics*, 40(2):287–295, 1998.
- [185] G Grimmet D Stirzaker and D Grimmett. Probability and random processes. *Oxford Science Publications, ISBN 0, 19(853665):8*, 2001.
- [186] Joachim Stöhr and Hans Christoph Siegmann. *Magnetism: from fundamentals to nanoscale dynamics*, volume 152. Springer Science & Business Media, 2007.
- [187] E. C. Stoner and E. P. Wohlfarth. A mechanism of magnetic hysteresis in heterogeneous alloys. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 240(826):599–642, 1948.

- [188] Julius Adams Stratton. *Electromagnetic theory*. John Wiley & Sons, 2007.
- [189] RP Tan, J Carrey, and M Respaud. Magnetic hyperthermia properties of nanoparticles inside lysosomes using kinetic Monte Carlo simulations: Influence of key parameters and dipolar interactions, and evidence for strong spatial variation of heating power. *Physical Review B*, 90(21):214421, 2014.
- [190] Burghard Thiesen and Andreas Jordan. Clinical applications of magnetic nanoparticles for hyperthermia. *International journal of hyperthermia*, 24(6):467–474, 2008.
- [191] Seiko Toraya-Brown and Steven Fiering. Local tumour hyperthermia as immunotherapy for metastatic cancer. *International Journal of Hyperthermia*, 30(8):531–539, 2014.
- [192] N. A. Usov, O. N. Serebryakova, and V. P. Tarasov. **Interaction effects in assembly of magnetic nanoparticles**. *Nanoscale Research Letters*, 12(1):489, Aug 2017. ISSN 1556-276X.
- [193] NG Van Kampen. Itô versus stratonovich. *Journal of Statistical Physics*, 24(1):175–187, 1981.
- [194] Nicolaas Godfried Van Kampen. *Stochastic processes in physics and chemistry*, volume 1. Elsevier, 2007.
- [195] Arne Vansteenkiste, Jonathan Leliaert, Mykola Dvornik, Mathias Helsen, Felipe Garcia-Sanchez, and Bartel Van Waeyenberge. The design and verification of MuMax3. *Aip Advances*, 4(10):107133, 2014.
- [196] Ediron Lima Verde, Gabriel Teixeira Landi, Juliano de Andrade Gomes, Marcelo Henrique Sousa, and Andris Figueiroa Bakuzis. Magnetic hyperthermia investigation of cobalt ferrite nanoparticles: Comparison between experiment, linear response theory, and dynamic hysteresis simulations. *Journal of Applied Physics*, 111(12):123902, 2012.
- [197] Arthur F Voter. Introduction to the kinetic Monte Carlo method. *Radiation effects in solids*, pages 1–23, 2007.
- [198] Xuman Wang, Hongchen Gu, and Zhengqiang Yang. The heating effect of magnetic fluids in an alternating magnetic field. *Journal of Magnetism and Magnetic Materials*, 293(1):334–340, 2005.
- [199] Eric W. Weisstein. **Sphere point picking**. *MathWorld—A Wolfram Web Resource*, 2016.
- [200] W Wernsdorfer, E Bonet Orozco, K Hasselbach, A Benoit, B Barbara, N Demoncey, A Loiseau, H Pascard, and D Mailly. Experimental evidence of the Néel-Brown model of magnetization reversal. *Physical Review Letters*, 78(9):1791, 1997.

-
- [201] Frank Wiekhorst, Uwe Steinhoff, Dietmar Eberbeck, and Lutz Trahms. Magnetorelaxometry assisting biomedical applications of magnetic nanoparticles. *Pharmaceutical research*, 29(5):1189–1202, 2012.
- [202] Wei Wu, Zhaohui Wu, Taekyung Yu, Changzhong Jiang, and Woo-Sik Kim. Recent progress on magnetic iron oxide nanoparticles: synthesis, surface functional strategies and biomedical applications. *Science and technology of advanced materials*, 16(2):023501, 2015.
- [203] Dongwon Yoo, Jae-Hyun Lee, Tae-Hyun Shin, and Jinwoo Cheon. Theranostic magnetic nanoparticles. *Accounts of chemical research*, 44(10):863–874, 2011.
- [204] Jing Zhong, Wenzhong Liu, Li Kong, and Paulo Cesar Morais. A new approach for highly accurate, remote temperature probing using magnetic nanoparticles. *Scientific reports*, 4:6338, 2014.
- [205] Robert Zwanzig. *Nonequilibrium statistical mechanics*. Oxford University Press, 2001.