

TOPICAL REVIEW • OPEN ACCESS

# Atomistic spin model simulations of magnetic nanomaterials

To cite this article: R F L Evans *et al* 2014 *J. Phys.: Condens. Matter* **26** 103202

View the [article online](#) for updates and enhancements.

## You may also like

- [Gaussian diffusion interrupted by Lévy walk](#)  
Piotr Weber and Piotr Pepowski
- [Generalized hydrodynamics and approach to generalized Gibbs equilibrium for a classical harmonic chain](#)  
Saurav Pandey, Abhishek Dhar and Anupam Kundu
- [Isotropic finite-difference discretization of stochastic conservation laws preserving detailed balance](#)  
Mahan Raj Banerjee, Sauro Succi, Santosh Ansumali et al.

## Topical Review

# Atomistic spin model simulations of magnetic nanomaterials

R F L Evans<sup>1</sup>, W J Fan, P Chureemart, T A Ostler<sup>2</sup>, M O A Ellis and R W Chantrell

Department of Physics, The University of York, York YO10 5DD, UK

E-mail: [richard.evans@york.ac.uk](mailto:richard.evans@york.ac.uk)

Received 3 December 2013, revised 2 January 2014

Accepted for publication 7 January 2014

Published 19 February 2014

## Abstract

Atomistic modelling of magnetic materials provides unprecedented detail about the underlying physical processes that govern their macroscopic properties, and allows the simulation of complex effects such as surface anisotropy, ultrafast laser-induced spin dynamics, exchange bias, and microstructural effects. Here we present the key methods used in atomistic spin models which are then applied to a range of magnetic problems. We detail the parallelization strategies used which enable the routine simulation of extended systems with full atomistic resolution.

Keywords: magnetism, atomistic model, classical spin model, Monte Carlo, spin dynamics

(Some figures may appear in colour only in the online journal)

## Contents

1. Introduction	2	5. Test simulations	10
2. The atomistic spin model	2	5.1. Angular variation of the coercivity	11
2.1. The classical spin Hamiltonian	3	5.2. Boltzmann distribution for a single spin	11
2.2. A note on magnetic units	4	5.3. Curie temperature	12
3. System parameterization and generation	4	5.4. Demagnetizing fields	12
3.1. Atomistic parameters from <i>ab initio</i> calculations	4	6. Parallel implementation and scaling	15
3.2. Atomistic parameters from macroscopic properties	5	7. Conclusions and perspectives	16
3.3. Atomistic system generation	6	Acknowledgments	17
4. Integration methods	6	Appendix A. Code structure and design philosophy	17
4.1. Spin dynamics	6	Appendix B. Atomistic system generation in VAMPIRE	17
4.2. Langevin dynamics	7	Appendix C. Parallelization strategies	18
4.3. Time Integration of the LLG equation	7	C.1. Statistical parallelism	18
4.4. Monte Carlo methods	9	C.2. Geometric decomposition	18
		C.3. Replicated data	20
		C.4. Additional scaling tests	20
		References	21

<sup>1</sup> [www-users.york.ac.uk/~rfl500/](http://www-users.york.ac.uk/~rfl500/)

<sup>2</sup> [www.tomostler.co.uk](http://www.tomostler.co.uk)



## 1. Introduction

Atomistic models of magnetic materials, where the atoms are treated as possessing a local magnetic moment, originated with Ising in 1925 as the first model of the phase transition in a ferromagnet [1]. The Ising model has spin-up and spin-down only states, and is amenable to analytical treatment, at least in two dimensions. Although it is still extensively used in the study of phase transitions, it is limited in applicability to magnetic materials and cannot be used for dynamic simulations. A natural extension of the Ising model is to allow the atomic spin to vary freely in 3D space [2, 3] which yields the classical Heisenberg model, where quantum mechanical effects on the atomic spins are neglected [2]. Monte Carlo simulations of the classical Heisenberg model allowed the study of phase transitions, surface and finite size effects in simple magnetic systems. The study of dynamic phenomena however was intrinsically limited due to the use of Monte Carlo methods until the development of dynamic [4, 5] and stochastic Landau–Lifshitz–Gilbert atomistic spin models [6–8].

Today atomistic simulation of magnetic materials has become an essential tool in understanding the processes governing the complex behaviour of magnetic nanomaterials, including ultrafast laser-induced magnetization dynamics [9–11], exchange bias in core-shell nanoparticles [12–14] and multilayers [5, 15], surface anisotropy in magnetic nanoparticles [16, 17], microstructural effects [18–20], spin valves [21] and spin torque [22], temperature effects and properties [23–26] and magnetic recording media [27, 28]. A significant capability of the atomistic spin model is to bridge the gap between *ab initio* electronic structure calculations and micromagnetics by using a multiscale model [29–32]. Such a model is able to calculate effective parameters for larger scale micromagnetic simulations [33], such as anisotropies, and exchange constants [34]. The atomistic model is also able to interface directly with micromagnetic simulations to tackle extended systems by calculating interface properties atomistically while treating the bulk of the material with a micromagnetic discretization [35, 36]. Despite the broad applicability and importance of atomistic models, no easy-to-use and open-source software packages are presently available to researchers, unlike the mesoscopic micromagnetic approaches where several packages are currently available [37–39].

Today most magnetic modelling is performed using numerical micromagnetics in finite difference [37] and finite element [38, 39] flavours. The theoretical basis of micromagnetics is that the atomic dipoles which make up the magnetic material can be approximated as a continuous vector field where, due to the exchange interaction, the atomic dipoles in a small finite volume are perfectly aligned. Micromagnetics has proven to be an essential tool in understanding a range of complex magnetic effects [40–42] but due to the rapid pace of technological development in magnetic materials the continuum approximation at its heart precludes its application to many problems of interest at the beginning of the 21st century, such as heat assisted magnetic recording [43], ultrafast laser-induced demagnetization [44, 45], exchange bias in spin valves [46], surface and interface anisotropy [47, 48] and high

anisotropy materials for ultrahigh density recording media such as FePt [49]. The common theme to these problems is a sub-nanometre spatial variation in the magnetization caused by high temperatures, atomic level ordering (anti- and ferrimagnets), or atomic surface and interface effects. To tackle these problems requires a formalism to take account of the detailed atomic structure to express its impact on the macroscopic behaviour of a nano particle, grain or complete device.

Some, but not all, of these problems can adequately be tackled by next-generation micromagnetic approaches utilizing the Landau–Lifshitz–Bloch equation [50–52], which is based on a physically robust treatment of the coupling of a macrospin to a heat bath, allowing the study of high temperature processes [53], ultrafast demagnetization [54, 55] and switching [56]. However, true atomic scale variations of the magnetization, as apparent in antiferromagnets, surfaces and interfaces, still require an atomistic approach. Additionally with the decreasing size of magnetic elements, finite size effects begin to play an increasing role in the physical properties of magnetic systems [57].

In this article we present an overview of the common computational methods utilized in atomistic spin simulations and details of their implementation in the open-source VAMPIRE software package<sup>3</sup>. Testing of the code is an essential part of ensuring the accuracy of the model and so we also detail important tests of the various parts of the model and compare them to analytic results while exploring some interesting physics of small magnetic systems.

VAMPIRE is designed specifically with these problems in mind, and can easily simulate nanoparticles, multilayer films, interfacial mixing, surface anisotropy and roughness, core-shell systems, granular media and lithographically defined patterns, all with fully atomistic resolution. In addition, truly realistic systems predicted by Molecular Dynamics simulations [19, 20, 59] can also be used giving unprecedented detail about the relationships between shape and structure and the magnetic properties of nanoparticles. In addition to these physical features VAMPIRE also utilizes the computing power of multiprocessor machines through parallelization, allowing systems of practical interest to be simulated routinely, and large-scale problems on the 100+ nm length scale to be simulated on computing clusters. Further details of the VAMPIRE code and its architecture are presented in appendix A.

## 2. The atomistic spin model

The physical basis of the atomistic spin model is the localization of unpaired electrons to atomic sites, leading to an effective local atomistic magnetic moment. The degree of localization of electrons has historically been a contentious issue in 3d metals [60], due to the magnetism originating in the outer electrons which are notionally loosely bound to the atoms. *Ab initio* calculations of the electron density [61] show that in reality, even in ‘itinerant’ ferromagnets, the spin polarization is well-localized to the atomic sites. Essentially

<sup>3</sup> Details available from [vampire.york.ac.uk](http://vampire.york.ac.uk).

this suggests that the bonding electrons are unpolarized, and after taking into account the bonding charge the remaining d-electrons form a well-defined effective localized moment on the atomic sites.

Magnetic systems are fundamentally quantum mechanical in nature since the electron energy levels are quantized, the exchange interaction is a purely quantum mechanical effect, and other important effects such as magnetocrystalline anisotropy arise from relativistic interactions of electronic orbitals with the lattice, which are the province of *ab initio* models. In addition to these properties at the electronic level, the properties of magnetic materials are heavily influenced by thermal effects which are typically difficult to incorporate into standard density functional theory approaches. Therefore models of magnetic materials should combine the quantum mechanical properties with a robust thermodynamic formalism. The simplest model of magnetism using this approach is the Ising model [1], which allows the atomic moments one of two allowed states along a fixed quantization axis. Although useful as a descriptive system, the forced quantization is equivalent to infinite anisotropy, limiting the applicability of the Ising model in relation to real materials. In the classical description the direction of the atomic moment is a continuous variable in 3D space allowing for finite anisotropies and dynamic calculations. In some sense the classical spin model is analogous to Molecular Dynamics, where the energetics of the system are determined primarily from quantum mechanics, but the time evolution and thermodynamic properties are treated classically.

### 2.1. The classical spin Hamiltonian

The extended Heisenberg spin model encapsulates the essential physics of a magnetic material at the atomic level, where the energetics of a system of interacting atomic moments is given by a *spin* Hamiltonian (which neglects non-magnetic effects such as the Coulomb term). The spin Hamiltonian  $\mathcal{H}$  typically has the form:

$$\mathcal{H} = \mathcal{H}_{\text{exc}} + \mathcal{H}_{\text{ani}} + \mathcal{H}_{\text{app}} \quad (1)$$

denoting terms for the exchange interaction, magnetic anisotropy, and externally applied magnetic fields respectively.

The dominant term in the spin Hamiltonian is the Heisenberg exchange energy, which arises due to the symmetry of the electron wavefunction and the Pauli exclusion principle [60] which governs the orientation of electronic spins in overlapping electron orbitals. Due to its electrostatic origin, the associated energies of the exchange interaction are around 1–2 eV, which is typically up to 1000 times larger than the next largest contribution and gives rise to magnetic ordering temperatures in the range 300–1300 K. The exchange energy for a system of interacting atomic moments is given by the expression

$$\mathcal{H}_{\text{exc}} = - \sum_{i \neq j} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j \quad (2)$$

where  $J_{ij}$  is the exchange interaction between atomic sites  $i$  and  $j$ ,  $\mathbf{S}_i$  is a unit vector denoting the local spin moment direction and  $\mathbf{S}_j$  is the spin moment direction of neighbouring

atoms. The unit vectors are taken from the actual atomic moment  $\boldsymbol{\mu}_s$  and given by  $\mathbf{S}_i = \boldsymbol{\mu}_s / |\boldsymbol{\mu}_s|$ . It is important to note here the significance of the sign of  $J_{ij}$ . For ferromagnetic materials where neighbouring spins align in parallel,  $J_{ij} > 0$ , and for antiferromagnetic materials where the spins prefer to align anti-parallel  $J_{ij} < 0$ . Due to the strong distance dependence of the exchange interaction, the sum in equation (2) is often truncated to include nearest neighbours only. This significantly reduces the computational effort while being a good approximation for many materials of interest. In reality the exchange interaction can extend to several atomic spacings [29, 30], representing hundreds of pairwise interactions.

In the simplest case the exchange interaction  $J_{ij}$  is isotropic, meaning that the exchange energy of two spins depends only on their relative orientation, not their direction. In more complex materials, the exchange interaction forms a tensor with components:

$$\mathbf{J}_{ij}^{\text{M}} = \begin{bmatrix} J_{xx} & J_{xy} & J_{xz} \\ J_{yx} & J_{yy} & J_{yz} \\ J_{zx} & J_{zy} & J_{zz} \end{bmatrix}, \quad (3)$$

which is capable of describing *anisotropic* exchange interactions, such as two-ion anisotropy [29] and the Dzyaloshinskii–Moriya interaction (off-diagonal components of the exchange tensor). In the case of tensorial exchange  $\mathcal{H}_{\text{exc}}^{\text{M}}$ , the exchange energy is given by the product:

$$\mathcal{H}_{\text{exc}}^{\text{M}} = - \sum_{i \neq j} [S_x^i S_x^j S_z^i] \begin{bmatrix} J_{xx} & J_{xy} & J_{xz} \\ J_{yx} & J_{yy} & J_{yz} \\ J_{zx} & J_{zy} & J_{zz} \end{bmatrix} \begin{bmatrix} S_x^j \\ S_y^j \\ S_z^j \end{bmatrix}. \quad (4)$$

Obtaining the components of the exchange tensor may be done phenomenologically, or via *ab initio* methods such as the relativistic torque method [62–65] or the spin-cluster expansion technique [30, 66–68]. The above expressions for the exchange energy also exclude higher-order exchange interactions such as three-spin and four-spin terms. In most materials the higher-order exchange terms are significantly smaller than the leading term and can safely be neglected.

While the exchange energy gives rise to magnetic ordering at the atomic level, the thermal stability of a magnetic material is dominated by the magnetic anisotropy, or preference for the atomic moments to align along a preferred spatial direction. There are several physical effects which give rise to anisotropy, but the most important is the magnetocrystalline anisotropy (namely the preference for spin moments to align with particular crystallographic axes) arising from the interaction of atomic electron orbitals with the local crystal environment [69, 70].

The simplest form of anisotropy is of the single-ion uniaxial type, where the magnetic moments prefer to align along a single axis,  $\mathbf{e}$ , often called the easy axis and is an interaction confined to the local moment. Uniaxial anisotropy is most commonly found in particles with elongated shape (shape anisotropy), or where the crystal lattice is distorted along a single axis as in materials such as hexagonal Cobalt and L1<sub>0</sub> ordered FePt. The uniaxial single-ion anisotropy energy is given by the expression:

$$\mathcal{H}_{\text{ani}}^{\text{uni}} = -k_{\text{u}} \sum_i (\mathbf{S}_i \cdot \mathbf{e})^2 \quad (5)$$

where  $k_u$  is the anisotropy energy per atom. Materials with a cubic crystal structure, such as iron and nickel, have a different form of anisotropy known as cubic anisotropy. Cubic anisotropy is generally much weaker than uniaxial anisotropy, and has three principal directions which energetically are easy, hard and very hard magnetization directions respectively. Cubic anisotropy is described by the expression:

$$\mathcal{H}_{\text{ani}}^{\text{cub}} = \frac{k_c}{2} \sum_i (S_x^4 + S_y^4 + S_z^4) \quad (6)$$

where  $k_c$  is the cubic anisotropy energy per atom, and  $S_x$ ,  $S_y$ , and  $S_z$  are the  $x$ ,  $y$ , and  $z$  components of the spin moment  $\mathbf{S}$  respectively.

Most magnetic problems also involve interactions between the system and external applied fields, denoted as  $\mathbf{H}_{\text{app}}$ . External fields can arise in many ways, for example a nearby magnetic material, or as an effective field from an electric current. In all cases the applied field energy is simply given by:

$$\mathcal{H}_{\text{app}} = - \sum_i \mu_s \mathbf{S}_i \cdot \mathbf{H}_{\text{app}}. \quad (7)$$

## 2.2. A note on magnetic units

The subject of magnetic units is controversial due to the existence of multiple competing standards and historical origins [60]. Starting from the atomic level however, the dimensionality of units is relatively transparent. Atomic moments are usually accounted for in multiples of the Bohr magneton ( $\mu_B$ ), the magnetic moment of an isolated electron, with units of  $\text{J T}^{-1}$ . Given a number of atoms of moment  $\mu_s$  in a volume, the moment per unit volume is naturally in units of  $\text{J T m}^{-3}$ , which is identical to the SI unit of  $\text{A m}^{-1}$ . However, the dimensionality (moment per unit volume) of the unit  $\text{A m}^{-1}$  is not as obvious as  $\text{J T}^{-1} \text{m}^{-3}$ , and so the latter form is used herein.

Applied magnetic fields are hence defined in Tesla, which comes naturally from the derivative of the spin Hamiltonian with respect to the local moment. The unit of Tesla for applied field is also beneficial for hysteresis loops, since the area enclosed a typical  $M-H$  loop is then given as an energy density ( $\text{J m}^{-3}$ ). A list of key magnetic parameters and variables and their units are shown in table 1.

## 3. System parameterization and generation

Unlike micromagnetic simulations where the magnetic system can be partitioned using either a finite difference or finite element discretization, atomistic simulations generally require some *a priori* knowledge of atomic positions. Most simple magnetic materials such as Fe, Co or Ni form regular crystals, while more complex systems such as oxides, antiferromagnets and Heusler alloys possess correspondingly complex atomic structures. For ferromagnetic metals, the details of atomic positions are generally less important due to the strong parallel orientation of moments, and so they can often (but not always) be represented using a simple cubic discretization. In contrast, the properties of ferrimagnetic and antiferromagnetic materials

**Table 1.** Table of key variables and their units.

Variable	Symbol	Unit	
Atomic magnetic moment	$\mu_s$	Joules/Tesla	( $\text{J T}^{-1}$ )
Unit cell size	$a$	Angstroms	( $\text{\AA}$ )
Exchange energy	$J_{ij}$	Joules/link	(J)
Anisotropy energy	$k_u$	Joules/atom	(J)
Applied field	$\mathbf{H}$	Tesla	(T)
Temperature	$T$	Kelvin	(K)
Time	$t$	Seconds	(s)

Parameter	Symbol	Value	
Bohr magneton	$\mu_B$	$9.2740 \times 10^{-24}$	$\text{J T}^{-1}$
Gyromagnetic ratio	$\gamma$	$1.76 \times 10^{11}$	$\text{T}^{-1} \text{s}^{-1}$
Permeability of free space	$\mu_0$	$4\pi \times 10^{-7}$	$\text{T}^2 \text{J}^{-1} \text{m}^3$
Boltzmann constant	$k_B$	$1.3807 \times 10^{-23}$	$\text{J K}^{-1}$

are inherently tied to the atomic positions due to frustration and exchange interactions, and so simulation of these materials must incorporate details of the atomic structure.

In addition to the atomic structure of the material, it is also necessary to parameterize the terms of the spin Hamiltonian given by equation (1), principally including exchange and anisotropy values but also with other terms. There are generally two ways in which this may be done: using experimentally determined properties or with a multiscale approach using *ab initio* density functional theory calculations as input to the spin model.

### 3.1. Atomistic parameters from *ab initio* calculations

*Ab initio* density functional theory (DFT) calculations utilize the Hohenberg–Kohn–Sham theory [71, 72] which states that the total energy  $E$  of a system can be written solely in terms the electron density,  $\rho$ . Thus, if the electron density is known then the physical properties of the system can be found. In practice, the both electron density and the spin density are used as fundamental quantities in the total energy expression for spin-polarized systems [73]. In many implementations DFT-based methods only consider the outer electrons of a system, since the inner electrons play a minimal role in the bonding and also partially screen the effect of the nuclear core. These effects are approximated by a pseudopotential which determines the potential felt by the valence electrons. In all-electron methods, however, the core electron density is also relaxed. By energy minimization, DFT enables the calculation of a wide range of properties, including lattice constants, and in the case of magnetic materials localized spin moments, magnetic ground state and the effective magnetocrystalline anisotropy. Standard software packages such as VASP [74], CASTEP [75, 76] and SIESTA [77] make such calculations readily accessible. At present determining site resolved properties such as anisotropy constants and pairwise exchange interactions is more involved and requires *ab initio* Green's functions techniques such as the fully relativistic Korringa–Kohn–Rostoker method [78, 79] or the LMTO method [80, 81] in conjunction with the magnetic force theorem [62]. An alternative approach for the calculation of exchange parameters is the utilization



of the generalized Bloch theorem for spin-spiral states in magnetic systems [82] together with a Fourier transformation of  $k$ -dependent spin-spiral energies [83, 84].

A number of studies have determined atomic magnetic properties from first principles calculations by direct mapping onto a spin model, including the principal magnetic elements Co, Ni and Fe [81], metallic alloys including FePt [29], IrMn [31], oxides [85] and spin glasses [86], and also bilayer systems such as Fe/FePt [87]. Such calculations give detailed insight into microscopic magnetic properties, including atomic moments, long-ranged exchange interactions, magnetocrystalline anisotropies (including surface and two-ion interactions) and other details not readily available from phenomenological theories. Combined with atomistic models it is possible to determine macroscopic properties such as the Curie temperature, temperature-dependent anisotropies, and magnetic ground states, often in excellent agreement with experiment. However, the computational complexity of DFT calculations also means that the systems which can be simulated with this multi scale approach are often limited to small clusters, perfect bulk systems and 2D periodic systems, while real materials of course often contain a plethora of defects disrupting the long range order. Some studies have also investigated the effects of disorder in magnetic systems combined with a spin model mapping, such as dilute magnetic semiconductors [88] and metallic alloys [89].

Magnetic properties calculated at the electronic level have a synergy with atomistic spin models, since the electronic properties can often be mapped onto a Heisenberg spin model with effective local moments. This multiscale electronic and atomistic approach avoids the continuum approximations of micromagnetics and treats magnetic materials at the natural atomic scale.

### 3.2. Atomistic parameters from macroscopic properties

The alternative approach to multiscale atomistic/density-functional-theory simulations is to derive the parameters from experimentally determined values. This has the advantage of speed and lower complexity, whilst foregoing microscopic details of the exchange interactions or anisotropies. Another key advantage of generic parameters is the possibility of parametric studies, where parameters are varied explicitly to determine their importance for the macroscopic properties of the system, such as has been done for studies of surface anisotropy [17] and exchange bias [13].

Unlike micromagnetic simulations, the fundamental thermodynamic approach of the atomistic model means that all parameters must be determined for *zero* temperature. The spin fluctuations then determine the intrinsic temperature dependence of the *effective* parameters which are usually put into micromagnetic simulations as parameters. Fortunately determination of the atomic moments, exchange constants and anisotropies from experimental values is relatively straightforward for most systems.

**3.2.1. Atomic spin moment.** The atomic spin moment  $\mu_s$  is related to the saturation magnetization simply by:

$$\mu_s = \frac{M_s a^3}{n_{\text{at}}} \quad (8)$$

where  $M_s$  is the saturation magnetization at 0 K in  $\text{J T}^{-1} \text{m}^{-3}$ ,  $a$  is the unit cell size (m), and  $n_{\text{at}}$  is the number of atoms per unit cell. We also note the usual convention of expressing atomic moments in multiples or fractions of the Bohr magneton,  $\mu_B$  owing to their electronic origin. Taking BCC iron as an example, the zero temperature saturation magnetization is  $1.75 \text{ MA m}^{-1}$  [90], unit cell size of  $a = 2.866 \text{ \AA}$ , this gives an atomic moment of  $2.22 \mu_B/\text{atom}$ .

**3.2.2. Exchange energy.** For a generic atomistic model with  $z$  nearest neighbour interactions, the exchange constant is given by the mean-field expression:

$$J_{ij} = \frac{3k_B T_c}{\epsilon z} \quad (9)$$

where  $k_B$  is the Boltzmann constant and  $T_c$  is the Curie temperature  $z$  is the number of nearest neighbours.  $\epsilon$  is a correction factor from the usual mean-field expression which arises due to spin waves in the 3D Heisenberg model [91]. Because of this  $\epsilon$  is also dependent on the crystal structure and coordination number, and so the calculated  $T_c$  will vary slightly according to the specifics of the system. For Cobalt with a  $T_c$  of 1388 K and assuming a hexagonal crystal structure with  $z = 12$ , this gives a nearest neighbour exchange interaction  $J_{ij} = 6.064 \times 10^{-21} \text{ J/link}$ .

**3.2.3. Anisotropy energy.** The atomistic magnetocrystalline anisotropy  $k_u$  is derived from the macroscopic anisotropy constant  $K_u$  by the expression:

$$k_u = \frac{K_u a^3}{n_{\text{at}}} \quad (10)$$

where  $K_u$  is given in  $\text{J m}^{-3}$ . In addition to the atomistic parameters, it is also worth noting the analogous expressions for the anisotropy field  $H_a$  for a single domain particle:

$$H_a = \frac{2K_u}{M_s} = \frac{2k_u}{\mu_s} \quad (11)$$

where symbols have their usual meaning. At this point it is worth mentioning that the measured anisotropy is a free energy difference. While the intrinsic  $k_u$  remains (to a first approximation) temperature independent, at a non-zero temperature the free energy in the easy/hard directions is increased/decreased due to the magnetization fluctuations. Thus the macroscopic anisotropy value decreases with increasing temperature, vanishing at  $T_c$ . The thermodynamic basis of atomistic models makes them highly suitable for the investigation of such phenomena, as we show later. Applying the above, parameters for the key ferromagnetic elements are given in table 2.

**Table 2.** Table of derived constants for the ferromagnetic elements Fe, Co, Ni and Gd.

	Fe	Co	Ni	Gd	Unit
Crystal structure	bcc	hcp	fcc	hcp	—
Unit cell size $a$	2.866	2.507	3.524	3.636	Å
Interatomic spacing $r_{ij}$	2.480	2.507	2.492	3.636	Å
Coordination number $z$	8	12	12	12	—
Curie temperature $T_c$	1043	1388	631	293	K
Spin-wave MF correction [91, 92] $\epsilon$	0.766	0.790	0.790	0.790	—
Atomic spin moment $\mu_s$	2.22	1.72	0.606	7.63	$\mu_B$
Exchange energy $J_{ij}$	$7.050 \times 10^{-21}$	$6.064 \times 10^{-21}$	$2.757 \times 10^{-21}$	$1.280 \times 10^{-21}$	J/link
Anisotropy energy [93] $k$	$5.65 \times 10^{-25}$	$6.69 \times 10^{-24}$	$5.47 \times 10^{-26}$	$5.93 \times 10^{-24}$	J/atom

**3.2.4. Ferrimagnets and antiferromagnets.** In the case of ferrimagnets and antiferromagnets the above methods for anisotropy and moment determination do not work due to the lack of macroscopic measurements, although the estimated exchange energies apply equally well to the Néel temperature provided no magnetic frustration (due to lattice symmetry) is present. In general, other theoretical calculations or formalisms are required to determine parameters, such as mean-field approaches [9] or density functional theory calculations [31].

### 3.3. Atomistic system generation

In addition to determining the parameters of the spin Hamiltonian, an essential part of the atomistic model is the determination of the nuclear, or atomic, positions in the system. In the multiscale approach utilizing *ab initio* parameterization of the system, the spin Hamiltonian is intrinsically tied to the atomic positions. The additional detail offered by first principles calculations is highly desirable even for perfect crystals and atomically sharp interfaces, however the computational complexity of the calculations limits the ability to parameterize a spin Hamiltonian for systems of extended defects over 10 nm+ length scales, including physical effects such as vacancies, impurities, dislocations and even amorphous materials.

For systems modelled using the nearest neighbour approximation, the atomic structures are much less restricted, allowing for simulations of material defects such as interface roughness [13] and intermixing [21], magnetic multilayers [94], disordered magnetic alloys [9], surface [17] and finite size effects [57]. VAMPIRE includes extensive functionality to generate such systems, the details of which are included in appendix B. In addition to crystallographic and molecular systems [95, 96] it is also possible to investigate magnetic effects in disordered materials and nanoparticles by incorporating the results of Molecular Dynamics simulations [19, 20, 97].

## 4. Integration methods

Although the spin Hamiltonian describes the energetics of the magnetic system, it provides no information regarding its time evolution, thermal fluctuations, or the ability to determine the ground state for the system. In the following the commonly utilized integration methods for atomistic spin models are introduced.

### 4.1. Spin dynamics

The first understanding of spin dynamics came from ferromagnetic resonance experiments, where the time-dependent behaviour of a magnetic materials is described by the equation derived by Landau and Lifshitz [98], and in the modern form given by:

$$\frac{\partial \mathbf{m}}{\partial t} = -\gamma[\mathbf{m} \times \mathbf{H} + \alpha \mathbf{m} \times (\mathbf{m} \times \mathbf{H})] \quad (12)$$

where  $\mathbf{m}$  is a unit vector describing the direction of the sample magnetization,  $\mathbf{H}$  is the effective applied field acting on the sample,  $\gamma$  is the gyromagnetic ratio and  $\alpha$  is a phenomenological damping constant which is a property of the material. The physical origin of the Landau–Lifshitz (LL) equation arises due to two distinct physical effects. The precession of the magnetization (first term in equation (12)) arises due to the quantum mechanical interaction of an atomic spin with an applied field. The relaxation of the magnetization (second term in equation (12)) is an elementary formulation of energy transfer representing the coupling of the magnetization to a heat bath which aligns the magnetization along the field direction with a characteristic coupling strength determined by  $\alpha$ . In the LL equation the relaxation rate of the magnetization to the field direction is a linear function of the damping parameter, which was shown by Gilbert to yield incorrect dynamics for materials with high damping [99]. Subsequently Gilbert introduced critical damping, with a maximum effective damping for  $\alpha = 1$ , to arrive at the Landau–Lifshitz–Gilbert (LLG) equation. Although initially derived to describe the dynamics of the macroscopic magnetization of a sample, the LLG is the standard equation of motion used in numerical micromagnetics, describing the dynamics of small magnetic elements.

The same equation of motion can also be applied at the atomistic level. The precession term arises quantum mechanically for atomic spins and the relaxation term now describes direct angular momentum transfer between the spins and the heat bath, which includes contributions from the lattice [100] and the electrons [101]. A distinction between the macroscopic LLG and the atomistic LLG now appears in terms of the effects included within the damping parameter. In the macroscopic LLG,  $\alpha$  includes all contributions, both intrinsic (such as spin–lattice and spin–electron interactions) and extrinsic (spin–spin interactions arising from demagnetization fields, surface

defects [102], doping [103] and temperature [50]), while the atomistic LLG only includes the local intrinsic contributions. To distinguish the different definitions of damping we therefore introduce a microscopic damping parameter  $\lambda$ . Although the form of the LLG is identical for atomistic and macroscopic length scales, the microstructural detail in the atomistic model allows for calculations of the effective damping including extrinsic effects, such as rare-earth doping [103]. Including a microscopic damping  $\lambda$  the atomistic Landau–Lifshitz–Gilbert equation is given by

$$\frac{\partial \mathbf{S}_i}{\partial t} = -\frac{\gamma}{(1+\lambda^2)} [\mathbf{S}_i \times \mathbf{H}_{\text{eff}}^i + \lambda \mathbf{S}_i \times (\mathbf{S}_i \times \mathbf{H}_{\text{eff}}^i)] \quad (13)$$

where  $\mathbf{S}_i$  is a unit vector representing the direction of the magnetic spin moment of site  $i$ ,  $\gamma$  is the gyromagnetic ratio and  $\mathbf{H}_{\text{eff}}^i$  is the net magnetic field on each spin. The atomistic LLG equation describes the interaction of an atomic spin moment  $i$  with an effective magnetic field, which is obtained from the negative first derivative of the complete spin Hamiltonian, such that:

$$\mathbf{H}_{\text{eff}}^i = -\frac{1}{\mu_s} \frac{\partial \mathcal{H}}{\partial \mathbf{S}_i} \quad (14)$$

where  $\mu_s$  is the local spin moment. The inclusion of the spin moment within the effective field is significant, in that the field is then expressed in units of Tesla, given a Hamiltonian in Joules. Given typical energies in the Hamiltonian of 10  $\mu\text{eV}$ –100 meV range. This gives fields typically in the range 0.1–1000 T, given a spin moment of the same order as the Bohr magneton ( $\mu_B$ ).

#### 4.2. Langevin dynamics

In its standard form the LLG equation is strictly only applicable to simulations at zero temperature. Thermal effects cause thermodynamic fluctuations of the spin moments which at sufficiently high temperatures are stronger than the exchange interaction, giving rise to the ferromagnetic-paramagnetic transition. The effects of temperature can be taken into account by using Langevin Dynamics, an approach developed by Brown [104]. The basic idea behind Langevin Dynamics is to assume that the thermal fluctuations on each atomic site can be represented by a Gaussian white noise term. As the temperature is increased, the width of the Gaussian distribution increases, thus representing stronger thermal fluctuations. The established Langevin Dynamics method is widely used for spin dynamics simulations and incorporates an effective thermal field into the LLG equation to simulate thermal effects [105–107]. The thermal fluctuations are represented by a Gaussian distribution  $\mathbf{\Gamma}(t)$  in three dimensions with a mean of zero. At each time step the instantaneous thermal field on each spin  $i$  is given by:

$$\mathbf{H}_{\text{th}}^i = \mathbf{\Gamma}(t) \sqrt{\frac{2\lambda k_B T}{\gamma \mu_s \Delta t}} \quad (15)$$

where  $k_B$  is the Boltzmann constant,  $T$  is the system temperature,  $\lambda$  is the Gilbert damping parameter,  $\gamma$  is the absolute value of the gyromagnetic ratio,  $\mu_s$  is the magnitude of the

atomic magnetic moment, and  $\Delta t$  is the integration time step. The effective field for application in the LLG equation with Langevin Dynamics then reads:

$$\mathbf{H}_{\text{eff}}^i = -\frac{1}{\mu_s} \frac{\partial \mathcal{H}}{\partial \mathbf{S}_i} + \mathbf{H}_{\text{th}}^i. \quad (16)$$

Given that for each time step three Gaussian distributed random numbers are required for every spin, efficient generation of such numbers is essential. VAMPIRE makes use the Mersenne Twister [108] uniform random number generator and the Ziggurat method [109] for generating the Gaussian distribution.

It is useful at this point to address the applicability of the atomistic LLG equation to slow and fast problems respectively. In reality the thermal and magnetic fluctuations are correlated at the atomic level, arising from the dynamic interactions between the atoms and lattice/electron system. These correlations may be important in terms of the thermal fluctuations experienced by the atomistic spins. In the conventional Langevin dynamics approach described above, the noise term is completely uncorrelated in time and space. For short timescales however, the thermal fluctuations are correlated in time, and so the noise is *coloured* [110]. The effect of the coloured noise is to lessen the effect of sudden temperature changes on the magnetization dynamics. However, the existence of ultrafast magnetization dynamics [11, 44], and that it is driven by a thermal rather than quantum mechanical process [111], requires that the effective correlation time is short, with an upper bound of around 1 fs. Given that the correlation time is close to the integration timestep, the applicability of the LLG to problems with timescales  $\geq 1$  fs is sound. There will be a point however where the LLG is no longer valid, where direct simulation of the microscopic damping mechanisms will be necessary. Progress has been made in linking molecular dynamics and spin models [100, 112, 113] which enables the simulation of spin–lattice interactions, which is particularly relevant for slow problems, such as conventional magnetic recording where switching occurs over nanosecond timescales. However, it is also essential to consider spin–electron effects [101, 114] necessary for ultrafast demagnetization processes, although the physical origins are still currently debated [115].

#### 4.3. Time Integration of the LLG equation

In order to determine the time evolution of a system of spins, it is necessary to solve the stochastic LLG equation, as given by equations (13) and (16), numerically. The choice of solver is limited due to the stochastic nature of the equations. Specifically, it is necessary to ensure convergence to the Stratonovich solution. This has been considered in detail by Garcia-Palacios [105], but the essential requirement [116] is that the solver enforces the conservation of the magnitude of the spin, either implicitly or by renormalization. The most primitive integration scheme uses Euler’s method, which assumes a linear change in the spin direction in a single discretized time step,  $\Delta t$ . An improved integration scheme, known as the Heun method [105] is commonly used, which



allows the use of larger time steps because of its use of a predictor–corrector algorithm. Other more advanced integration schemes have been suggested, such as the midpoint method [117] and modified predictor–corrector midpoint schemes [103, 118]. The principal advantage of the midpoint scheme is that the length of the spin vector is preserved during the integration which allows larger time steps to be used. However for the midpoint scheme the significant increase in computational complexity outweighs the benefits of larger time steps [118]. Modified predictor–corrector schemes [103, 118] reduce the computational complexity of the midpoint scheme, but with a loss of accuracy, particularly in the time-dependent dynamics [103]. For valid integration of the stochastic LLG equation it is also necessary for the numerical scheme to converge to the Stratonovich solution [105, 119]. Although proven for the midpoint and Heun numerical schemes, the validity of the predictor–corrector midpoint schemes for the stochastic LLG have yet to be confirmed. On balance the Heun scheme, despite its relative simplicity, is sufficiently computationally efficient that it is still the most widely used integration scheme for stochastic magnetization dynamics, and so we proceed to describe its implementation in detail.

In the Heun method, the first (predictor) step calculates the new spin direction,  $\mathbf{S}'_i$ , for a given effective field  $\mathbf{H}_{\text{eff}}^i$  by performing a standard Euler integration step, given by:

$$\mathbf{S}'_i = \mathbf{S}_i + \Delta\mathbf{S}\Delta t \quad (17)$$

where

$$\Delta\mathbf{S} = -\frac{\gamma}{(1+\lambda^2)}[\mathbf{S}_i \times \mathbf{H}_{\text{eff}}^i + \lambda\mathbf{S}_i \times (\mathbf{S}_i \times \mathbf{H}_{\text{eff}}^i)]. \quad (18)$$

The Heun scheme does not preserve the spin length and so it is essential to renormalize the spin unit vector length  $\mathbf{S}_i$  after both the predictor and corrector steps to ensure numerical stability and convergence to the Stratonovich solution. After the first step the effective field must be re-evaluated as the intermediate spin positions have now changed. It should be noted that the random thermal field does not change between steps. The second (corrector) step then uses the predicted spin position and revised effective field  $\mathbf{H}_{\text{eff}}^{i'}$  to calculate the final spin position, resulting in a complete integration step given by:

$$\mathbf{S}_i^{+\Delta t} = \mathbf{S}_i + \frac{1}{2} [\Delta\mathbf{S} + \Delta\mathbf{S}'] \Delta t \quad (19)$$

where

$$\Delta\mathbf{S}' = -\frac{\gamma}{(1+\lambda^2)}[\mathbf{S}'_i \times \mathbf{H}_{\text{eff}}^{i'} + \lambda\mathbf{S}'_i \times (\mathbf{S}'_i \times \mathbf{H}_{\text{eff}}^{i'})]. \quad (20)$$

The predictor step of the integration is performed on every spin in the system before proceeding to evaluate the corrector step for every spin. This is then repeated many times so that the time evolution of the system can be simulated. Although the Heun scheme was derived specifically for a stochastic equation with multiplicative noise, in the absence of the noise term the Heun method reduces to a standard second order Runge–Kutta method [120]. In order to test the implementation of the Heun integration scheme, it is possible to compare the calculated result with the analytical solution for the LLG. For the simple

case of a single spin initially along the  $x$ -axis in an applied field along the  $z$ -axis, the time evolution [121] is given by:


$$\begin{aligned} S_x(t) &= \text{sech}\left(\frac{\lambda\gamma H}{1+\lambda^2}t\right) \cos\left(\frac{\gamma H}{1+\lambda^2}t\right) \\ S_y(t) &= \text{sech}\left(\frac{\lambda\gamma H}{1+\lambda^2}t\right) \sin\left(\frac{\gamma H}{1+\lambda^2}t\right) \\ S_z(t) &= \tanh\left(\frac{\lambda\gamma H}{1+\lambda^2}t\right). \end{aligned} \quad (21)$$

The expected and simulated time evolution for a single spin with  $H = 10$  T,  $\Delta t = 1 \times 10^{-15}$  s and  $\lambda = 0.1, 0.05$  is plotted in figure 1. Superficially the simulated and expected time evolution agree very well, with errors around  $10^{-6}$ . The error gives a characteristic trace the size and shape of which is indicative of a correct implementation of the Heun integration scheme.

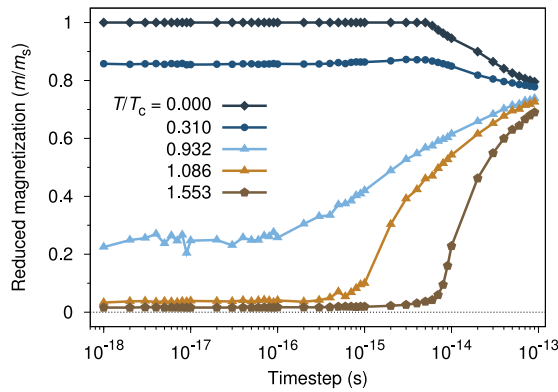
Ideally one would like to use the largest time step possible so as to simulate systems for the longest time. For micromagnetic simulations at zero temperature, the minimum time step is a well defined quantity since the largest field (usually the exchange term) essentially defines the precession frequency. However, for atomistic simulations using the stochastic LLG equation with Langevin dynamics, the effective field becomes temperature dependent. The consequence of this is that for atomistic models the most difficult region to integrate is in the immediate vicinity of the Curie point. Errors in the integration of the system will be apparent from a non-converged value for the average magnetization. This gives a relatively simple case which can then be used to test the stability of integration schemes for the stochastic LLG model. A plot of the mean magnetization as a function of temperature is shown in figure 2 for a representative system consisting of  $22 \times 22 \times 22$  unit cells with generic material parameters of FePt with an fcc crystal structure, nearest neighbour exchange interaction of  $J_{ij} = 3.0 \times 10^{-21}$  J/link and uniaxial anisotropy of  $1.0 \times 10^{-23}$  J/atom. The system is first equilibrated for 10 ps at each temperature and then the mean magnetization is calculated over a further 10 ps.

First, comparing the effect of temperature on the minimum allowable time step, the data show that for low temperatures reasonably large time steps of  $1 \times 10^{-15}$  s give the correct solution of the LLG equations, while near the Curie point (690 K) the deviations from the correct equilibrium value are significant. Consequently for simulations studying high temperature reversal processes time steps of  $1 \times 10^{-16}$  s are necessary. It should be noted that the time steps which can be used are material-dependent—specifically if a material with higher Curie temperature is used then the usable time steps will be correspondingly lower due to the increased exchange field.

From a practical perspective a significant advantage of the spin dynamics method is the ability to parallelize the integration system by domain decomposition, details of which are given in appendix C. This allows the efficient simulation of relatively large systems consisting of tens or hundreds of grains or nano structures with dimensions greater than 100 nm for nanosecond timescales, with typical numbers of spins in the range  $10^6$ – $10^8$ .



**Figure 1.** Time evolution of a single isolated spin in an applied field of 10 T and time step of 1 fs. Magnetization traces (a) and (c) show relaxation of the magnetization to the  $z$ -direction and precession of the  $x$  component (the  $y$ -component is omitted for clarity) for damping constants  $\lambda = 0.1$  and  $\lambda = 0.05$  respectively. The points are the result of direction integration of the LLG and the lines are the analytical solution plotted according to equation (21). Panels (b) and (d) show the corresponding error traces (difference between the expected and calculated spin components) for the two damping constants for (a) and (c) respectively. For  $\lambda = 0.1$  the error is below  $10^{-6}$ , while for lower damping the numerical error increases significantly due to the increased number of precessions, highlighting the damping dependence of the integration time step.



**Figure 2.** Time step dependence of the mean magnetization for different reduced temperatures for the Heun integration scheme. Low ( $T \ll T_c$ ) and high ( $T \gg T_c$ ) temperatures integrate accurately with a 1fs timestep, but in the vicinity of  $T_c$  a timestep of around  $10^{-16}$  is required for this system.

#### 4.4. Monte Carlo methods

While spin dynamics are particularly useful for obtaining dynamic information about the magnetic properties or reversal processes for a system, they are often not the optimal method for determining the equilibrium properties, for


example the temperature-dependent magnetization. The Monte Carlo Metropolis algorithm [122] provides a natural way to simulate temperature effects where dynamics are not required due to the rapid convergence to equilibrium and relative ease of implementation.

The Monte Carlo metropolis algorithm for a classical spin system proceeds as follows. First a random spin  $i$  is picked and its initial spin direction  $\mathbf{S}_i$  is changed randomly to a new trial position  $\mathbf{S}'_i$ , a so-called trial move. The change in energy  $\Delta E = E(\mathbf{S}'_i) - E(\mathbf{S}_i)$  between the old and new positions is then evaluated, and the trial move is then accepted with probability

$$P = \exp\left(-\frac{\Delta E}{k_B T}\right) \quad (22)$$

by comparison with a uniform random number in the range 0–1. Probabilities greater than 1, corresponding with a reduction in energy, are accepted unconditionally. This procedure is then repeated until  $N$  trial moves have been attempted, where  $N$  is the number of spins in the complete system. Each set of  $N$  trial moves comprises a single Monte Carlo step.

The nature of the trial move is important due to two requirements of any Monte Carlo algorithm: ergodicity and reversibility. Ergodicity expresses the requirement that all possible states of the system are accessible, while reversibility



**Figure 3.** Schematic showing the three principal Monte Carlo moves: (a) spin flip; (b) Gaussian; and (c) random.

requires that the transition probability between two states is invariant, explicitly  $P(\mathbf{S}_i \rightarrow \mathbf{S}'_i) = P(\mathbf{S}'_i \rightarrow \mathbf{S}_i)$ . From equation (22) reversibility is obvious since the probability of a spin change depends only on the initial and final energy. Ergodicity is easy to satisfy by moving the selected spin to a random position on the unit sphere, however this has an undesirable consequence at low temperatures since large deviations of spins from the collinear direction are highly improbable due to the strength of the exchange interaction. Thus at low temperatures a series of trial moves on the unit sphere will lead to most moves being rejected. Ideally a move acceptance rate of around 50% is desired, since very high and very low rates require significantly more Monte Carlo steps to reach a state representative of true thermal equilibrium.


One of the most efficient Monte Carlo algorithms for classical spin models was developed by Hinzke and Nowak [123], involving a combinational approach using a mixture of different trial moves. The principal advantage of this method is the efficient sampling of all available phase space while maintaining a reasonable trial move acceptance rate. The Hinzke–Nowak method utilizes three distinct types of move: spin flip, Gaussian and random, as illustrated schematically in figure 3.

The spin flip move simply reverses the direction of the spin such that  $\mathbf{S}'_i = -\mathbf{S}_i$  to explicitly allow the nucleation of a switching event. The spin flip move is identical to a move in Ising spin models. It should be noted that spin flip moves do not by themselves satisfy ergodicity in the classical spin model, since states perpendicular to the initial spin direction are inaccessible. However, when used in combination with other ergodic trial moves this is quite permissible. The Gaussian trial move takes the initial spin direction and moves the spin to a point on the unit sphere in the vicinity of the initial position according to the expression

$$\mathbf{S}'_i = \frac{\mathbf{S}_i + \sigma_g \mathbf{\Gamma}}{|\mathbf{S}_i + \sigma_g \mathbf{\Gamma}|} \quad (23)$$

where  $\mathbf{\Gamma}$  is a Gaussian distributed random number and  $\sigma_g$  is the width of a cone around the initial spin  $\mathbf{S}_i$ . After generating the trial position  $\mathbf{S}'_i$  the position is normalized to yield a spin of unit length. The choice of a Gaussian distribution is deliberate since after normalization the trial moves have a uniform sampling over the cone. The width of the cone is generally chosen to be temperature dependent and of the form

$$\sigma_g = \frac{2}{25} \left( \frac{k_B T}{\mu_B} \right)^{1/5}. \quad (24)$$



**Figure 4.** Visualization of Monte Carlo sampling on the unit sphere for (a) random and (b) Gaussian sampling algorithms at  $T = 10$  K. The dots indicate the trial moves. The random algorithm shows a uniform distribution on the unit sphere, and no preferential biasing along the axes. The Gaussian trial moves are clustered around the initial spin position, along the  $z$ -axis.

The Gaussian trial move thus favours small angular changes in the spin direction at low temperatures, giving a good acceptance probability for most temperatures.

The final random trial move picks a random point on the unit sphere according to

$$\mathbf{S}'_i = \frac{\mathbf{\Gamma}}{|\mathbf{\Gamma}|} \quad (25)$$

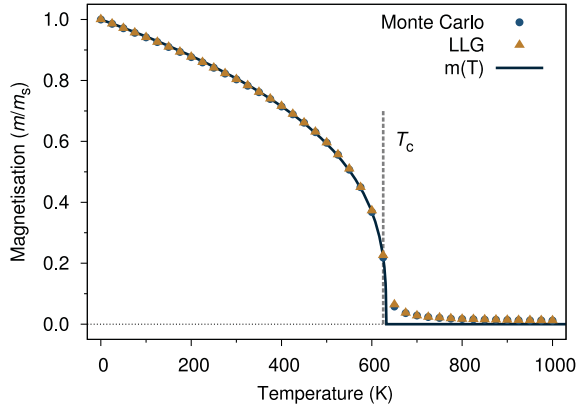
which ensures ergodicity for the complete algorithm and ensures efficient sampling of the phase space at high temperatures. For each trial step one of these three trial moves is picked randomly, which in general leads to good algorithmic properties.

To verify that the random sampling and Gaussian trial moves give the expected behaviour, a plot of the calculated trial moves on the unit sphere for the different algorithms is shown in figure 4. The important points are that the random trial move is uniform on the unit sphere, and that the Gaussian trial move is close to the initial spin direction, along the  $z$ -axis in this case.

At this point it is worthwhile considering the relative efficiencies of Monte Carlo and spin dynamics for calculating equilibrium properties. Figure 5 shows the simulated temperature-dependent magnetization for a test system using both LLG spin dynamics and Monte Carlo methods. Agreement between the two methods is good, but the spin dynamics simulation takes around twenty times longer to compute due to the requirements of a low time step and slower convergence to equilibrium. However, Monte Carlo algorithms are notoriously difficult to parallelize, and so for larger systems LLG spin dynamic simulations are generally more efficient than Monte Carlo methods.

## 5. Test simulations

Having outlined the important theoretical and computational methods for the atomistic simulation of magnetic materials, we now proceed to detail the tests we have refined to ensure the correct implementation of the main components of the model. Such tests are particularly helpful to those wishing to implement these methods. Similar tests developed for micro-magnetic packages [124] have proven an essential benchmark for the implementation of improved algorithms and codes with different capabilities but the same core functionality.



**Figure 5.** Comparative simulation of temperature-dependent magnetization for Monte Carlo and LLG simulations. Simulation parameters assume a nearest neighbour exchange of  $6.0 \times 10^{-21}$  J/link with a simple cubic crystal structure, periodic boundary conditions and 21952 atoms. The Monte Carlo simulations use 50 000 equilibration and averaging steps, while the LLG simulations use 5 000 000 equilibration and averaging steps with critical damping ( $\lambda = 1$ ) and a time step of 0.01 fs. The value of  $T_c \sim 625$  K calculated from equation (9) is shown by the dashed vertical line. The temperature-dependent magnetization is fitted to the expression  $m(T) = (1 - T/T_c)^\beta$  (shown by the solid line) which yields a fitted  $T_c = 631.82$  K and exponent  $\beta = 0.334\,297$ .

### 5.1. Angular variation of the coercivity

Assuming a correct implementation of an integration scheme as described in the previous section, the first test case of interest is the correct implementation of uniaxial magnetic anisotropy. For a single spin in an applied field and at zero temperature, the behaviour of the magnetization is essentially that of a Stoner–Wohlfarth particle, where the angular variation of the coercivity, or reversing field, is well known [125]. This test serves to verify the static solution for the LLG equation by ensuring an easy axis loop gives a coercivity of  $H_k = 2k_u/\mu_s$  as expected analytically. For this problem the Hamiltonian reads

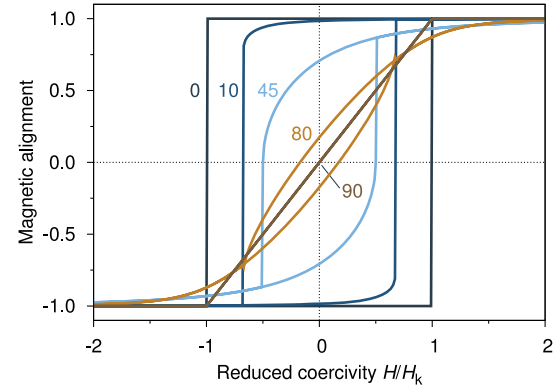
$$\mathcal{H} = -k_u S_z^2 - \mu_s \mathbf{S} \cdot \mathbf{H}_{\text{app}} \quad (26)$$

where  $k_u$  is the on-site uniaxial anisotropy constant and  $\mathbf{H}_{\text{app}}$  is the external applied field. The spin is initialized pointing along the applied field direction, and then the LLG equation is solved for the system, until the net torque on the system  $\mathbf{S} \times \mathbf{H}_{\text{eff}} \leq |10^{-6}|$  T, essentially a condition of *local* minimum energy.


The field strength is decreased from saturation in steps of  $0.01 H/H_k$  and solved again until the same condition is reached. A plot of the calculated alignment of the magnetization to the applied field ( $\mathbf{S} \cdot \mathbf{H}_{\text{app}}$ ) for different angles from the easy axis is shown in figure 6. The calculated hysteresis curve conforms exactly to the Stoner–Wohlfarth solution.

### 5.2. Boltzmann distribution for a single spin

To quantitatively test the thermal effects in the model and the correct implementation of the stochastic LLG or Monte Carlo integrators, the simplest case is that of the Boltzmann



**Figure 6.** Plot of alignment of magnetization with the applied field for different angles of from the easy axis. The  $0^\circ$  and  $90^\circ$  loops were calculated for very small angles from the easy and hard axes respectively, since in the perfectly aligned case the net torque is zero and no change of the spin direction occurs.



**Figure 7.** Calculated angular probability distribution for a single spin with anisotropy for different effective temperatures  $k_u/k_B T$ . The lines show the analytic solution given by equation (27).

distribution for a single spin with anisotropy (or applied field), where the probability distribution is characteristic of the temperature and the anisotropy energy. The Boltzmann distribution is given by:

$$P(\theta) \propto \sin \theta \exp\left(-\frac{k_u \sin^2 \theta}{k_B T}\right) \quad (27)$$

where  $\theta$  is the angle from the easy axis. The spin is initialized along the easy axis direction and the system is allowed to evolve for  $10^8$  time steps after equilibration, recording the angle of the spin to the easy axis at each time. Since the anisotropy energy is symmetric along the easy axis, the probability distribution is reflected and summed about  $\pi/2$ , since at low temperatures the spin is confined to the upper well ( $\theta < \pi/2$ ). Figure 7 shows the normalized probability distribution for three reduced temperatures.

The agreement between the calculated distributions is excellent, indicating a correct implementation of the stochastic LLG equation.



### 5.3. Curie temperature

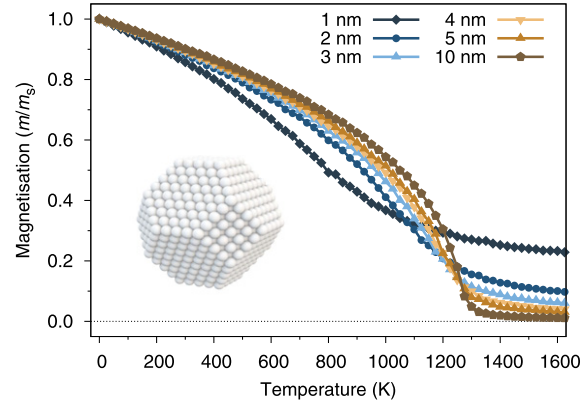
Tests such as the Stoner–Wohlfarth hysteresis or Boltzmann distribution are helpful in verifying the mechanical implementation of an algorithm for a single spin, but interacting systems of spins present a significant challenge in that no analytical solutions exist. Hence it is necessary to calculate some well-defined macroscopic property which ensures the correct implementation of interactions in a system. The Curie temperature  $T_c$  of a nanoparticle is primarily determined by the strength of the exchange interaction between spins and so makes an ideal test of the exchange interaction. As discussed previously the bulk Curie temperature is related to the exchange coupling by the mean-field expression given in equation (9). However, for nanoparticles with a reduction in coordination number at the surface and a finite number of spins, the Curie temperature and *criticality* of the temperature-dependent magnetization will vary significantly with varying size [57].

To investigate the effects of finite size and reduction in surface coordination on the Curie temperature, the equilibrium magnetization for different sizes of truncated octahedron nanoparticles was calculated as a function of temperature. The Hamiltonian for the simulated system is

$$\mathcal{H} = - \sum_{i \neq j} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j \quad (28)$$

where  $J_{ij} = 5.6 \times 10^{-21}$  J/link, and the crystal structure is face-centred-cubic, which is believed to be representative of Cobalt nanoparticles. Given the relative strength of the exchange interaction, anisotropy generally has a negligible impact on the Curie temperature of a material, and so the omission of anisotropy from the Hamiltonian is purely for simplicity. The system is simulated using the Monte Carlo method with 10 000 equilibration and 20 000 averaging steps. The system is heated sequentially in 10 K steps, with the final state of the previous temperature taken as the starting point of the next temperature to minimize the number of steps required to reach thermal equilibrium. The mean temperature-dependent magnetization for different particle sizes is plotted in figure 8.

From equation (9) the expected Curie temperature is 1282 K, which is in agreement with the results for the 10 nm diameter nanoparticle. For smaller particle sizes the magnetic behaviour close to the Curie temperature loses its criticality, making  $T_c$  difficult to determine. Traditionally the Curie point is taken as the maximum of the gradient  $dm/dT$  [57], however this significantly underestimates the actual temperature at which magnetic order is lost (which is, by definition, the Curie temperature). Other estimates of the Curie point such as the divergence in the susceptibility are probably a better estimate for finite systems, but this is beyond the scope of the present article. Another effect visible for very small particle sizes is the appearance of a magnetization above the Curie point, an effect first reported by Binder [126]. This arises from local moment correlations which exist above  $T_c$ . It is an effect only observable in nanoparticles where the system size is close to the magnetic correlation length.



**Figure 8.** Calculated temperature-dependent magnetization and Curie temperature for truncated octahedron nanoparticles with different size. A visualization of a 3 nm diameter particle is inset.

### 5.4. Demagnetizing fields

For systems larger than the single domain limit [33] and systems which have one dimension significantly different from another, the demagnetizing field can have a dominant effect on the macroscopic magnetic properties. In micromagnetic formalisms implemented in software packages such as OOMMF [37], MAGPAR [38] and NMAG [39], the calculation of the demagnetization fields is calculated accurately due to the routine simulation of large systems where such fields dominate. Due to the long-ranged interaction the calculation of the demagnetization field generally dominates the compute time and so computational methods such as the fast-Fourier-transform [127, 128] and multipole expansion [129] have been developed to accelerate their calculation.


In large-scale atomistic calculations, it is generally sufficient to adopt a micromagnetic discretization for the demagnetization fields, since they only have a significant effect on nanometre length scales [7]. Additionally due to the generally slow variation of magnetization, the timescales associated with the changes in the demagnetization field are typically much longer than the time step for atomistic spins. Here we present a modified finite difference scheme for calculating the demagnetization fields, described as follows.

The complete system is first discretized into macrocells with a fixed cell size, each consisting of a number of atoms, as shown in figure 9(a). The cell size is freely adjustable from atomistic resolution to multiple unit cells depending on the accuracy required. The position of each macrocell  $p_{mc}$  is determined from the magnetic ‘centre of mass’ given by the expression

$$p_{mc}^{\alpha} = \frac{\sum_i^n \mu_i p_i^{\alpha}}{\sum_i^n \mu_i} \quad (29)$$

where  $n$  is the number of atoms in the macrocell,  $\mu_i$  is the local (site-dependent) atomic spin moment and  $\alpha$  represents the spatial dimension  $x, y, z$ . For a magnetic material with the same magnetic moment at each site, equation (29) corrects for partial occupation of a macrocell by using the mean atomic position as the origin of the macrocell dipole, as shown in figure 9(b). For a sample consisting of two materials with different atomic moments, the ‘magnetic centre of mass’ is





**Figure 9.** (a) Visualization of the macrocell approach used to calculate the demagnetization field, with the system discretized into cubic macrocells. Each macrocell consists of several atoms, shown schematically as cones. (b) Schematic of the macrocell discretization at the curved surface of a material, indicated by the dashed line. The mean position of the atoms within the macrocell defines the centre of mass where the effective macrocell dipole is located. (c) Schematic of a macrocell consisting of two materials with different atomic moments. Since the magnetization is dominated by one material, the magnetic centre of mass moves closer to the material with the higher atomic moments.

closer to the atoms with the higher atomic moments, as shown in figure 9(c). This modified micromagnetic scheme gives a good approximation of the demagnetization field without having to use computationally costly atomistic resolution calculation of the demagnetization field.

The total *moment* in each macrocell  $m_{\text{mc}}$  is calculated from the vector sum of the atomic moments within each cell, given by

$$m_{\text{mc}}^{\alpha} = \sum_i^n \mu_i S_i^{\alpha}. \quad (30)$$

Depending on the particulars of the system, the macrocell moments can vary significantly depending on position, composition and temperature. At elevated temperatures close to the Curie point, the macrocell magnetization becomes small, and so the effects of the demagnetizing field are much less important. Similarly in compensated ferrimagnets consisting of two competing sublattices the overall macrocell magnetization can also be small again leading to a reduced influence of the demagnetizing field.

The demagnetization field within each macrocell  $p$  is given by

$$\mathbf{H}_{\text{demag}}^{\text{mc},p} = \frac{\mu_0}{4\pi} \left( \sum_{p \neq q} \frac{3(\mathbf{m}_{\text{mc}}^q \cdot \hat{\mathbf{r}})\hat{\mathbf{r}} - \mathbf{m}_{\text{mc}}^q}{r^3} \right) - \frac{\mu_0}{3} \frac{\mathbf{m}_{\text{mc}}^p}{V_{\text{mc}}^p} \quad (31)$$

where  $r$  is the separation between dipoles  $p$  and  $q$ ,  $\hat{\mathbf{r}}$  is a unit vector in the direction  $p \rightarrow q$ , and  $V_{\text{mc}}^p$  is the volume of the macrocell  $p$ . The first term in equation (31) is the usual dipole term arising from all *other* macrocells in the system, while the second term is the self-demagnetization field of the macrocell, taken here as having a demagnetization factor 1/3. Strictly this is applicable only for the field at the centre of a cube. However, the non-uniformity of the field inside a uniformly magnetized cube is not large and the assumption of a uniform demagnetization field is a reasonable approximation. The self-demagnetization term is often neglected in the literature, but in fact is essential when calculating the field inside a magnetic material. Once the demagnetization field is calculated for each macrocell, this is applied uniformly to all atoms as an effective field within the macrocell. It should be noted however that the macrocell size cannot be larger than the smallest sample dimension, otherwise significant errors in the calculation of the demagnetizing field will be incurred.

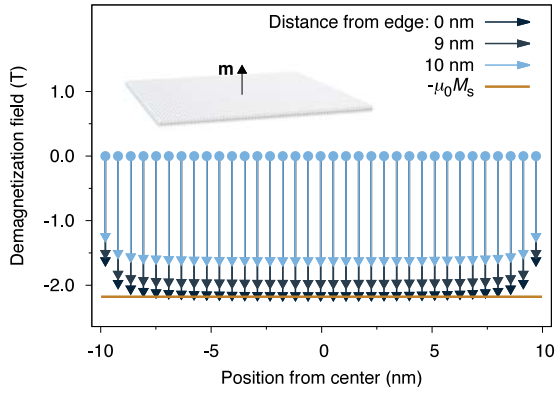
The volume of the macrocell  $V_{\text{mc}}$  is an effective volume determined from the number of atoms in each cell and given by

$$V_{\text{mc}} = n_{\text{mc}}^{\text{a}} V_{\text{atom}} = n_{\text{mc}}^{\text{a}} \frac{V_{\text{uc}}}{n_{\text{uc}}^{\text{a}}} \quad (32)$$

where  $n_{\text{mc}}^{\text{a}}$  is the number of atoms in the macrocell,  $n_{\text{uc}}^{\text{a}}$  is the number of atoms in the unit cell and  $V_{\text{uc}}$  is the volume of the unit cell. The macrocell volume is necessary to determine the magnetization (moment per volume) in the macrocell. For unit cells much smaller than the system size, equation (32) is a good approximation, however for a large unit cell with significant free space, for example a nanoparticle in vacuum, the free space contributes to the effective volume which reduces the effective macrocell volume.

**5.4.1. Demagnetizing field of a platelet.** To verify the implementation of the demagnetization field calculation it is necessary to compare the calculated fields with some analytic solution. Due to the complexity of demagnetization fields analytical solutions are only available for simple geometric shapes such as cubes and cylinders [130], however for an infinite perpendicularly magnetized platelet the demagnetization field approaches the magnetic saturation  $-\mu_0 M_s$ . To test this limit we have calculated the demagnetizing field of a  $20 \text{ nm} \times 20 \text{ nm} \times 1 \text{ nm}$  platelet as shown in figure 10. In the centre of the film agreement with the analytic value is good, while at the edges the demagnetization field is reduced as expected.

**5.4.2. Performance characteristics.** In micromagnetic simulations, calculation of the demagnetization field usually dominates the runtime of the code and generally it is preferable to have as large a cell size as possible. For atomistic calculations however, additional flexibility in the frequency of updates of



**Figure 10.** Calculated cross-section of the demagnetization fields in a  $20 \text{ nm} \times 20 \text{ nm} \times 1 \text{ nm}$  platelet (visualization inset) with magnetization perpendicular to the film plane. A macrocell size of 2 unit cells is used. In the centre of the film the calculated demagnetization field is  $-2.236 \text{ T}$  which compares well to the analytic solution of  $\mathbf{H}_{\text{demag}} = -\mu_0 \mathbf{M} = -2.18 \text{ T}$ . Note that the 2D grid used slightly overestimates the demagnetization field.

the demagnetization field is permitted due to the short time steps used and the fact that the magnetization is generally a slowly varying property.

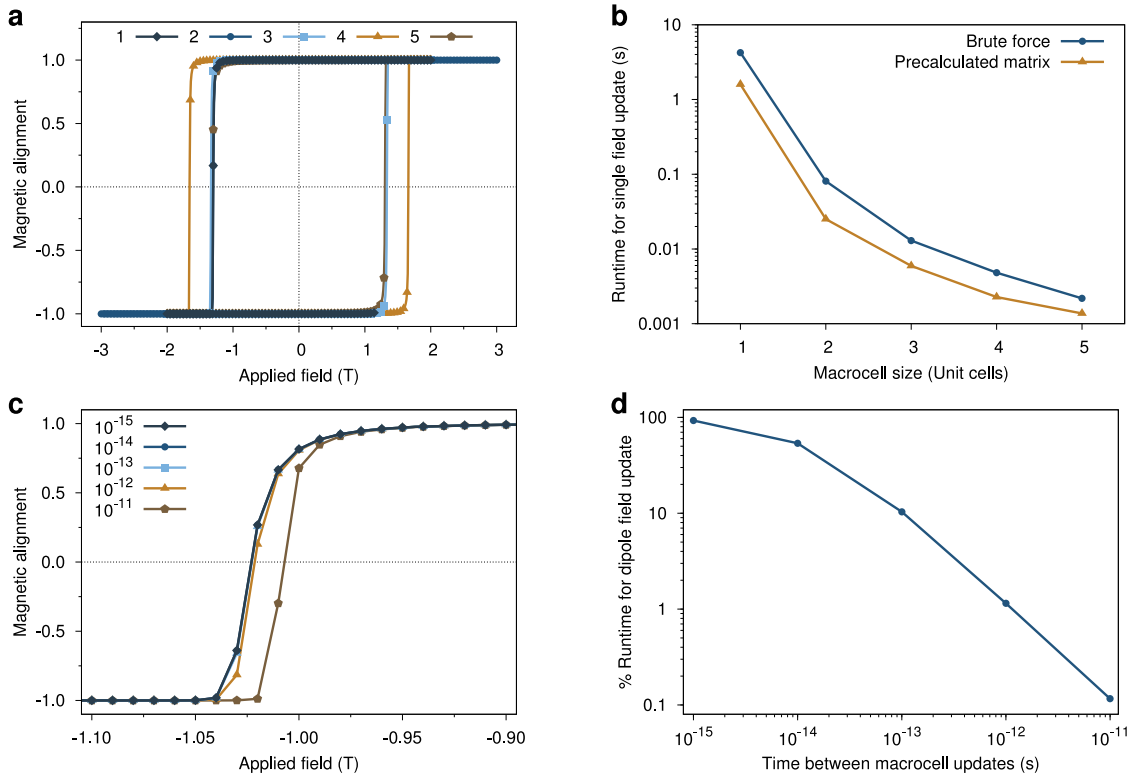
To investigate the effects of different macrocell sizes and the time taken between updates of the demagnetization field we have simulated hysteresis loops of a nanodot of diameter  $40 \text{ nm}$  and height of  $1.4 \text{ nm}$ . Figure 11(a) shows hysteresis loops calculated for different macrocell sizes for the nanodot. For

most cell sizes the results of the calculation agree quite well, however, for a cell size of 4 unit cells, the calculated coercivity is significantly larger, owing to the creation of a flat macrocell (with dimensions  $4 \times 4 \times 1$  unit cells). This illustrates that for systems with small dimensions, extra care must be taken when specifying the macro cell size in order to avoid non-cubic cells. In general, the problem with asymmetric macrocells is not trivial to solve within the finite difference formalism, since the problem arises due to a modification of both the intracell and intercell contributions to the demagnetizing field.

Figure 11(b) shows the runtime for a single update of the demagnetizing field on a single CPU for different macrocell size discretizations. Noting the logarithmic scale for the simulation time, single unit cell discretizations are computationally costly while not giving significantly better results than larger macrocell discretizations. Although the demagnetization field calculation is an  $n_{\text{mc}}^2$  problem, it is possible to pre-calculate the distances between the macrocells at the cost of increased memory usage. Due to the computational cost of calculating the position vectors, this method is often much faster than the brute force calculation. However, due to the fact that memory usage increases proportionally to  $n_{\text{mc}}^2$ , fine discretizations for large systems can require many GBs of memory.

By collating terms in equation (31) it is possible to construct the following matrix  $\mathbf{M}_{pq}$  for each pairwise interaction:

$$\mathbf{M}_{pq} = \begin{bmatrix} (3r_x r_x - 1)/r_{pq}^3 - 1/3 & 3r_x r_y & 3r_x r_z \\ 3r_y r_y & (3r_y r_y - 1)/r_{pq}^3 - 1/3 & 3r_y r_z \\ 3r_z r_z & 3r_z r_y & (3r_z r_z - 1)/r_{pq}^3 - 1/3 \end{bmatrix} \quad (33)$$



**Figure 11.** Simulated hysteresis loops and computational efficiency for a  $40 \text{ nm} \times 40 \text{ nm} \times 1 \text{ nm}$  nanodot for different cell sizes (multiples of unit cell size) ((a), (b)) and update rates (seconds between update calculations) ((c), (d)).

where  $r_x, r_y, r_z$  are the components of the unit vector in the direction  $p \rightarrow q$ , and  $r_{pq}$  is the separation of macrocells. Since the matrix is symmetric along the diagonal only six numbers need to be stored in memory. The total demagnetization field for each macrocell  $p$  is then given by:

$$\mathbf{H}_{\text{demag}}^{\text{mc},p} = \frac{\mu_0}{4\pi} \left( \sum_{p \neq q} \mathbf{M}_{pq} \cdot \mathbf{m}_{\text{mc}}^q \right) - \frac{\mu_0}{3} \frac{\mathbf{m}_{\text{mc}}^p}{V_{\text{mc}}^p}. \quad (34)$$

The relative performance of the matrix optimization is plotted for comparison in figure 11(b), showing a significant reduction in runtime. Where the computer memory is sufficiently large, the recalculated matrix should always be employed for optimal performance.

In addition to variable macrocell sizes, due to the small time steps employed in atomistic models and that the magnetization is generally a slowly varying property, it is not always necessary to update the demagnetization fields every single time step. Hysteresis loops for different times between updates of the demagnetization field are plotted in figure 11(c). In general hysteresis calculations are sufficiently accurate with a picosecond update of the demagnetizing field, which significantly reduces the computational cost.

In general good accuracy for the demagnetizing field calculation can be achieved with coarse discretization and infrequent updates, but fast dynamics such as those induced by laser excitation require much faster updates, or simulation of domain wall processes in high anisotropy materials requires finer discretizations to achieve correct results.


**5.4.3. Demagnetizing field in a prolate ellipsoid.** Since the macrocell approach works well in platelets and nanodots, it is also interesting to apply the same method to a slightly more complex system: a prolate ellipsoid. An ellipsoid adds an effective shape anisotropy due to the demagnetization field, and so for a particle with uniaxial magnetocrystalline anisotropy along the elongated direction ( $z$ ), the calculated coercivity should increase according to the difference in the demagnetization field along  $x$  and  $z$ , given by:

$$H_{\text{dm}}^{\text{shape}} = +\Delta N \mu_0 M_s \quad (35)$$

where  $\Delta N = N_z - N_x$ . The demagnetizing factors  $N_x, N_y$ , and  $N_z$  are known analytically for various ellipticities [131], and here we assume  $a/c = b/c = 0.5$ , where  $a, b$ , and  $c$  are the extent of the ellipsoid along  $x, y$  and  $z$  respectively.

To verify the macrocell approach gives the same expected increase of the coercivity we have simulated a generic ferromagnet with atomic moment  $1.5 \mu_B$ , an FCC crystal structure with lattice spacing  $3.54 \text{ \AA}$  and anisotropy field of  $H_a = 1 \text{ T}$ . The particle is cut from the lattice in the shape of an ellipsoid, of diameter  $10 \text{ nm}$  and height of  $20 \text{ nm}$ , as shown inset in figure 12. A macrocell size of 2 unit cells is used, which is updated every 100 time steps ( $0.1 \text{ ps}$ ).

As expected the coercivity increases due to the shape anisotropy. From [131] the expected increase in the coercivity is  $H_{\text{dm}}^{\text{shape}} = 0.37 \text{ T}$  which compares well to the simulated increase of  $0.33 \text{ T}$ .



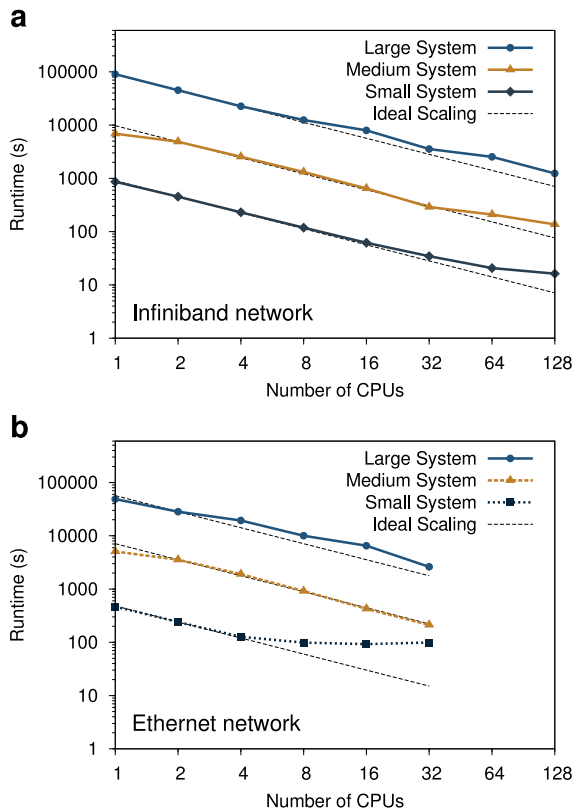
**Figure 12.** Simulated hysteresis loops for an ellipsoidal nanoparticle with an axial ratio of 2 showing the effect of the demagnetizing field calculated with the macrocell approach. A visualization of the simulated particle is inset.

## 6. Parallel implementation and scaling

Although the algorithms and methods discussed in the preceding sections describe the mechanics of atomistic spin models, it is important to note finally the importance of parallel processing in simulating realistic systems which include many-particle interactions, or nano patterned elements with large lateral sizes. Details of the parallelization strategies which have been adopted to enable the optimum performance of VAMPIRE for different problems are presented in appendix C. In general terms the parallelization works by subdividing the simulated system into sections, with each processor simulating part of the complete system. Spin orientations at the processor boundaries have to be exchanged with neighbouring processors to calculate the exchange interactions, which for small problems and large numbers of processors can significantly reduce the parallel efficiency. The use of latency hiding, where the local spins are calculated in parallel with the inter-processor communications, is essential to ensure good scaling for these problems.

To demonstrate the performance and scalability of VAMPIRE, we have performed tests for three different system sizes: small ( $10628$  spins), medium ( $8 \times 10^5$  spins), and large ( $8 \times 10^6$  spins). We have access to two Beowulf-class clusters; one with 8 cores/node with an Infiniband 10 Gbps low-latency interconnect, and another with 4 cores/node with a Gigabit Ethernet interconnect. For parallel simulations the interconnect between the nodes can be a limiting factor for increasing performance with increasing numbers of processors, since as more processors are added, each has to do less work per time step. Eventually network communication will dominate the calculation since processors with small amounts of work require the data from other processors in shorter times, leading to a drop in performance. The scaling performance of the code for 100 000 time steps on both machines is presented in figure 13.

The most challenging case for parallelization is the small system size, since a significant fraction of the system must be communicated to other processors during each timestep.



**Figure 13.** Runtime scaling of VAMPIRE for three different problem sizes on the Infiniband network (a) and Ethernet network (b), normalized to the runtime for 2 cores for each problem size.

On the Ethernet network system for the smallest system size reasonable scaling is seen only for 4 CPUs due to the high latency of the network. However larger problems are much less sensitive to network latency due to latency hiding, and show excellent scalability up to 32 CPUs. Essentially this means that larger problems scale much better than small ones, allowing more processors to be utilized. This is of course well known for parallel scaling problems, but even relatively modest systems consisting of  $10^5$  spins show significant improvements with more processors.

For the system with the low-latency Infiniband network, excellent scalability is seen for all problems up to 64 CPUs. Beyond 64 CPUs the reduced scalability for all problem sizes is likely due to a lack of network bandwidth. The bandwidth requirements are similar for all problem sizes, since smaller problems complete more time steps in a given period of time and so have to send several sets of data to other processors. Nevertheless improved performance is seen with increasing numbers of CPUs allowing for continued reductions in compute time. Although not shown, initial tests on an IBM Blue Gene class system have demonstrated excellent scaling of VAMPIRE up to 16 000 CPUs, allowing the real possibility for atomistic simulations with lateral dimensions of micrometres. Additional scaling tests for systems including calculation of the demagnetizing field and a long-ranged exchange interaction are presented in appendix C.

## 7. Conclusions and perspectives

We have described the physical basis of the rapidly developing field of atomistic spin models, and given examples via its implementation in the form of the VAMPIRE code. Although the basic formalism underpinning atomistic spin models is well established, ongoing developments in magnetic materials and devices means that new approaches will need to be developed to simulate a wider range of physical effects at the atomistic scale. One of the most important phenomena is spin transport and magnetoresistance which is behind an emergent field of spin–electronics, or spintronics. Simulation of spin transport and spin torque switching is already in development, and must be included in atomistic level models in order to simulate a wide range of spintronic materials and devices, including read sensors and MRAM (magnetic random access memory). Other areas of interest include ferroelectrics, the spin Seebeck effect [132], and Coloured noise [110] where simulation capabilities are desirable, and incorporation of these effects are planned in future. In addition to modelling known physical effects, it is hoped that improved models of damping incorporating phononic and electronic mechanisms will be developed which enable the study of magnetic properties of materials at sub-femtosecond timescales.

The ability of atomistic models to incorporate magnetic parameters from density functional theory calculations is a powerful combination which allows complex systems such as alloys, surfaces and defects to be accurately modelled. This multiscale approach is essential to relate microscopic quantum mechanical effects to a macroscopic length scale accessible to experiment. Such a multiscale approach leads to the possibility of simulation driven technological development, where the magnetic properties of a complete device can be predicted and optimized through a detailed understanding of the underlying physics. Due to the potential of multiscale simulations, it is planned in future to develop links to existing DFT codes such as CASTEP [75, 76] to allow easier integration of DFT parameters and atomistic spin models.

The computational methods presented here provide a sound basis for atomistic simulation of magnetic materials, but further improvements in both algorithmic and computational efficiency are of course likely. One area of potential computational improvement is GPGPU (general purpose graphics processing unit) computation, which utilizes the massively parallel nature of GPUs to accelerate simulations, with speed ups over a single CPU of 75 times routinely reported. With several supercomputers moving to heterogenous computing architectures utilizing both CPUs and GPUs, supporting GPGPU computation is likely to be important in future, and an implementation in our VAMPIRE code is currently planned. In terms of algorithmic improvements it should be noted that the Heun numerical scheme although simple is relatively primitive by modern standards, and moving to a midpoint scheme may allow for larger time steps to be used than currently.

With the continuing improvements in computer power, atomistic simulations have become a viable option for the routine simulation of magnetic materials. With the increasing complexity of devices and material properties, atomistic



models are a significant and important development. While micromagnetic models remain a powerful tool for the simulation and design of devices, the limitations of the (continuum) micromagnetic formalism are increasingly exposed by its failure to deal with the complex physics of elevated temperatures, ultrafast magnetization processes and interfaces. The atomistic model also allows for the study of complex microstructural effects which is pertinent to Brown's paradox [133], where there is generally a large disparity between the intrinsic and measured magnetic properties arising from nucleation caused by defects in the sample structure. While micromagnetics will remain the computational model of choice for large-scale and coarse-grained applications, the ability to accurately model the effects of microscopic details, temperature effects and ultrafast dynamics make atomistic models an essential tool to understand the physics of magnetic materials at the forefront of the field.

## Acknowledgments

The authors wish to thank Jerome Jackson for helpful discussions particularly with regard to the parallel implementation of the code; Thomas Schrefl for insisting on the name of VAMPIRE; Matt Probert, Phil Hasnip, Joe Barker, Uli Nowak and Ondrej Hovorka for helpful discussions; and Laszlo Szunyogh, Phivos Mavropoulos and Stefan Blügel for assistance with the section on the *ab initio* parameterization of the atomistic model. We also acknowledge the financial support from the EU Seventh Framework Programme grant agreement No. 281043 FEMTOSPIN and Seagate Technology Inc. This work made use of the facilities of N8 HPC provided and funded by the N8 consortium and EPSRC (Grant No. EP/K000225/1) co-ordinated by the Universities of Leeds and Manchester and the EPSRC Small items of research equipment at the University of York ENERGY (Grant No. EP/K031589/1).

## Appendix A. Code structure and design philosophy

In addition to implementing the necessary computational methods for magnetic atomistic calculations, it is also important to provide a framework structure for the code, where new additions in the form of features or improvements can be made with minimal disturbance to other sections of the code. Equally important for intensive computational problems is ensuring high performance of the code so that simulations proceed as rapidly as possible.

In VAMPIRE this is achieved through hybrid coding using a mixture of object-oriented and functional programming styles. Object-oriented programming is widely used in modern software projects as a level of abstraction around the data, or objects, in the code. This abstraction makes it easy to store information about an element, for example an atom, as a single unified data type, known as a class. One significant caveat with object-oriented code is that it is generally hard to optimize for maximum performance. High performance codes generally utilize a different coding approach known as functional programming, where the focus is on functions which operate on numerous data sets. However the organization of data into large

blocks in functional programming generally makes it harder to organize the data. VAMPIRE therefore makes use of both methodologies to gain the benefits of object-oriented design during the initialization phase, while for the performance-critical parts of the code the data is re-arranged to use a functional style for maximum performance. Due to the requirements of high performance, object-oriented design and parallelization, the C++ programming language was chosen for all parts of the code. The popularity of the C++ language also allows for easy future integration of other libraries, such as NVIDIA's CUDA framework for utilizing graphics processing units. For portability the code also has a minimal dependence on external libraries and also conforms to the published standard allowing simple compilation on the widest possible variety of computer architectures.

In addition to the low-level structure described in terms of object-oriented and functional programming styles, the code is also designed in a modular fashion so that most of the mechanistic operations (such as the parallelization and data analysis) are separated from high level functions which control the various simulation types. This enables users to easily add new simulation types or physical effects to the code, without having to be concerned with the inner workings.

## Appendix B. Atomistic system generation in VAMPIRE

VAMPIRE has a number of dedicated functions for generating atomic systems within the nearest neighbour approximation. The principal advantage of the nearest neighbour approximation is its simplicity and ability to consider a wide range of physical effects such as finite size, surfaces, ferri and antiferromagnets, disordered systems etc. with relative ease. VAMPIRE also includes in-built particle structures to enable generation of systems with simple geometric shapes such as spheres, cylinders, truncated octahedra and cubes.

The first step is to generate a crystal lattice of the desired type and dimensions sufficiently large to incorporate the complete simulated system. For the nearest neighbour approximation the Hamiltonian is generally only well defined for a single unified crystal structure, and therefore such generic simulations require a single crystal from which the system is cut. More complex structures are readily simulated, however the user must define the complete Hamiltonian for the system, taking into account the realistic interfaces between different crystals. VAMPIRE uses the unit cell as the essential building block of the atomic structure, since the exchange interactions of atoms between neighbouring unit cells are known before the structure is generated. The global crystal is generated by replicating the basic unit cell on a grid in  $x$ ,  $y$  and  $z$ . This bare crystal structure is then cut into the desired geometry, for example a single nanoparticle, Voronoi granular structure, or a user defined 2D geometry by removing atoms from the complete generated crystal. Atoms within this geometry are then assigned to one or more materials as desired (each material having different magnetic properties such as atomic spin moments, anisotropy or exchange interactions), generating the complete atomic system. The assignment of



different parts of the system to different materials enables the easy simulation of multilayers and core-shell nanoparticles, as well as combinations of these for systems such as multilayer magnetic recording media. As an example, figure 14 shows a visualization of a multilayer magnetic recording media generated using VAMPIRE.

Once the structure is defined the exchange interactions for all atoms in the system are calculated from a list of nearest neighbour interactions for the defined unit cell. Since each cell on the grid contains a fixed number of atoms, and the exchange interactions of those atoms with other neighbouring cells are known relative to the local cell, the interaction list is trivial to generate. For computational efficiency the final interaction list is then stored as a one dimensional (1D) linked list.

## Appendix C. Parallelization strategies

A consistent trend among computers today is the drive towards parallel architectures, designed to improve overall performance in a consistent and scalable fashion. The downside of this approach is that software must be specially modified to take advantage of the hardware, which still presents a significant challenge. In order to make the best use of parallel computers, we have adopted a number of distinct parallelization strategies. This approach means that for any given problem, an optimal strategy can be utilized to achieve maximum performance.

### C.1. Statistical parallelism

The most trivial form of parallelism is batch or statistical, where the statistical properties of a system are determined by a series of independent calculations, each of which can be run in parallel. Statistical parallelism has the prime advantage that the division of work leads to an ideal scaling behaviour, since each of the runs are entirely independent and require no intercommunication.

In magnetic simulations, the most common applications of statistical parallelism are sweeps of the parameter space for a particular system, or in determining thermodynamic averages. For the former, a given system is calculated for different values of key parameters, for example, anisotropy or exchange constants; for the latter, the same system is simulated, but each run is given a different seed for the random number generator. This leads to a different thermodynamic evolution, which can provide information about the statistical behaviour of the system. It should be noted that the correct seeding of the random number generator, where a number of uncorrelated random number sequences are generated, is quite complex [108]. For magnetic simulations, the chaotic nature of the system, whereby a small change in the time evolution rapidly leads to a significantly different result, means that crude sequential number seeding is quite satisfactory.

### C.2. Geometric decomposition

Although statistical parallelism is useful for some types of simulation, it has one significant limitation: it can only be

applied to relatively small systems, as the entire problem must be solved on a single processor. For larger systems it is necessary to divide the system into smaller parts for parallel execution. The most efficient method for such parallelization is generally geometric decomposition, where the space is divided into cells, and each processor is assigned a cell to simulate. If well implemented, geometric decomposition can be scaled to run on thousands of CPUs, and this is one of the key aims of our implementation.

The starting point for geometric decomposition is efficiently dividing the space to run on  $N_{\text{CPU}}$  CPUs. In order to achieve this, we have devised an algorithm which takes into account the physical system dimensions and which searches for a solution where

$$n_x \cdot n_y \cdot n_z = N_{\text{CPU}} \quad (\text{C.1})$$

while minimizing the surface to volume ratio. If the dimensions of the overall system are given by  $l_x$ ,  $l_y$ , and  $l_z$ , then the volume of each cell is:

$$V_{\text{cell}} = \frac{l_x}{n_x} \cdot \frac{l_y}{n_y} \cdot \frac{l_z}{n_z} \quad (\text{C.2})$$

and the surface area of each cell is:

$$A_{\text{cell}} = 2 \left[ \frac{l_x l_y}{n_x n_y} + \frac{l_y l_z}{n_y n_z} + \frac{l_x l_z}{n_x n_z} \right]. \quad (\text{C.3})$$


The surface to volume ratio is then given by:

$$\frac{A_{\text{cell}}}{V_{\text{cell}}} = 2 \left[ \frac{n_z}{l_z} + \frac{n_y}{l_y} + \frac{n_x}{l_x} \right]. \quad (\text{C.4})$$


It is clear that the minimum in the surface to volume ratio occurs for  $n_\alpha/l_\alpha = 1$  for all three dimensions, essentially showing that longer dimensions parallelize better with more CPUs.

Given that the dimensions of the system are fixed, the only free variables are the number of CPUs in each dimension,  $n_x$ ,  $n_y$ , and  $n_z$ . These are further constrained by equation (C.1). In order to find the optimal solution for a given number of CPUs, the starting point is  $n_\alpha = \frac{l_\alpha}{\sqrt{N_{\text{CPU}}}}$ . Exact solutions for  $n_x$ ,  $n_y$  and  $n_z$  are then searched for and the one with the lowest surface to volume ratio is selected. This approach is very flexible and allows for efficient decomposition for any number of CPUs. The only problematic solution is for prime numbers of CPUs, where only one exact solution exists, though this is a rare occurrence for large numbers of CPUs. A visualization of a cubic system decomposed into 48 blocks is shown in figure 15.


Having decomposed the system, each CPU is allocated a cell which defines its own spatial domain. In order to maintain maximum scalability, each CPU generates its own portion of the complete system, and all associated data. This has the advantage of minimizing the memory footprint and also parallelizing the system creation, which can become a significant bottleneck for very large numbers of processors. Once the local atoms have been generated it is necessary to know which atoms on remote processors (halo atoms) are potentially interacting with atoms on the local processor (boundary atoms), as well as which atoms are interacting



**Figure 14.** Visualization of a magnetic recording medium generated using VAMPIRE. The medium consists of magnetically hard and soft layers with interfacial mixing of atoms between the layers. The material is granular in nature, and so a Voronoi tessellation is overlaid on top of the layers to form the isolated magnetic grains. Finally, a dilute intermixing layer is applied between the grains representing the diffusion of magnetic atoms into the  $\text{SiO}_2$  between the grains, as seen in real media.



**Figure 15.** Visualization of the decomposition of a cubic system into 48 blocks of equal volume.



**Figure 16.** Visualization of the different categorizations of atoms on a processor, determined by their spatial location. The core and boundary regions exist on the local processor, denoted by the regions within the dashed line. The halo region contains atoms on remote processors which atoms on the local processor potentially interact with.

locally only (core atoms). This essentially defines three distinct regions, as shown schematically in figure 16.

The maximum interaction range of the atoms is known globally, and so provided all atoms in this range are included, generation of the neighbour list is trivially the same as the serial

case. In practice this is implemented by a global broadcast of each processor's domain, i.e., which regions of space are 'owned' by each processor. Each processor then looks at each atom in its boundary region, and then dispatches a copy of the atom to the appropriate neighbouring processors. This method has the advantage that it is quite general, and can be applied to any decomposition method, not necessarily cubes. At this point parallel periodic boundary conditions are easily implemented in the same manner, by copying atoms at the edge of the system to the desired processors. Once all boundary atoms have been sent, and all halo atoms have been received, the neighbour list is generated in the usual fashion with a linked-cell algorithm. After the actual neighbour list has been generated, it is likely that some of the copied halo atoms are in fact not needed, and so these atoms are deleted. Similarly some atoms in the in the boundary region may not interact with the halo, and these atoms are re-assigned to the core region. Following this book-keeping exercise, parallel simulation of the system can begin.

The method we have adopted for parallel simulation of the system makes use of latency hiding, where requests for data from other processors are made prior to a locally compute-intensive period, after which the requested data should have arrived. Such latency hiding is an important consideration when running the code on many processors. In practice atoms on each processor are ordered according to their interaction classification, i.e.: core; boundary; and finally halo atoms. The integration of the system proceeds as follows:

- A request is made for all halo data from other processors.
- The core region is then integrated.
- If the halo data has not arrived, then wait for it.
- Integrate the boundary region.
- Global synchronization.

The parallel integration is repeated the desired number of times during the simulation.

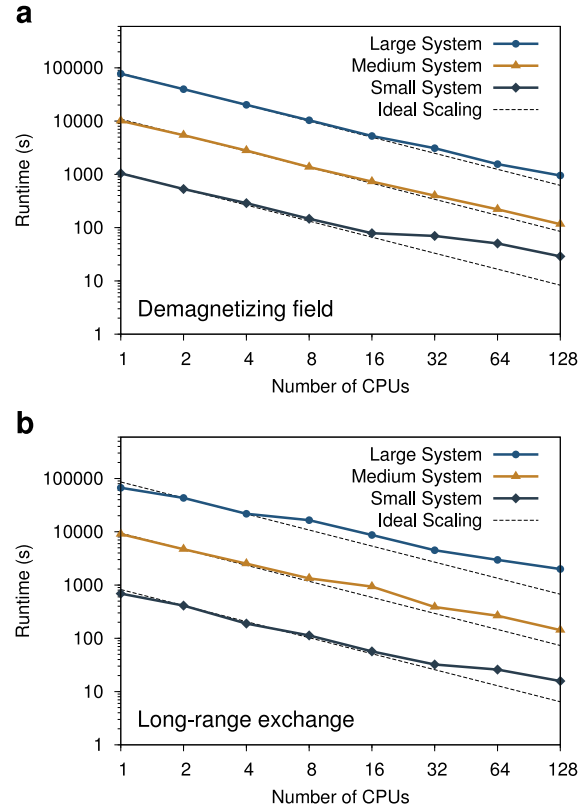
### C.3. Replicated data

For continuous systems, geometric decomposition provides an efficient way of parallelizing the calculation. However, for sparse systems geometric decomposition can be inefficient due to poor load balancing, where some processors have many more atoms than others. This means some processors spend a significant amount of time waiting for others to complete the integration step, leading to a reduction in scalability. In magnetism, such systems are typically granular, consisting of a small number of grains. One solution to sparse systems is to utilize a replicated data approach, where each processor has a complete copy of all data, similar to the statistical parallelism method. Each processor then simulates  $(1/n_{\text{cpu}})$ th of the total system, without any constraints on spatial locality. The atoms are classified in a similar way to the geometric decomposition approach, as core, boundary, and halo, and integrated in exactly the same way.

The principal disadvantages of the replicated data approach are the increased memory footprint (each processor must generate a complete copy of the system), and the tendency to have a high proportion of boundary spins. The latter can be mitigated by re-ordering spins in memory to ensure some degree of spatial locality. For granular systems this is fact trivial, since the assignment of each spin to a grain provides the necessary geometric information, and so the spins are ordered by a unique grain identification, which is also spatially correlated. In addition to its use in simulation of sparse systems, the replicated data approach is also the strategy adopted for the parallel calculation of the long-ranged demagnetization fields. The method for the parallel code is identical to that described earlier, where the spins are allocated to macrocells which then interact with each other. Since the calculation of the demagnetization field in each cell requires knowledge of all other cells, replicated data is the logical choice for the parallelization. The demagnetization field calculation proceeds as follows:

- The macrocell moments are initialized to zero on all processors.
- Each processor determines the contribution of its spins to each macrocell.
- The macrocell moments are summed globally, so that each processor has a complete copy of the macrocell magnetizations.
- Each processor calculates the demagnetization field only for macrocells which contain local spins.
- The local demagnetization field on each spin is determined from its macrocell demagnetization field.

This approach leads to excellent scaling, as for reasonable macrocell sizes the communication costs are minimal, and for  $\gg 1$  macrocell per processor the method scales linearly with  $n_{\text{cpu}}$ . Fine macrocell discretizations ( $\ll 27$  unit cells/macrocell) can lead to significant memory and computation costs, but in general this is unnecessary for most atomic scale calculations.



**Figure 17.** Runtime scaling of VAMPIRE for three different problem sizes including demagnetizing fields (a) and a spin Hamiltonian including a long-ranged exchange interaction (b), normalized to the runtime for 2 cores for each problem size. The cluster has 8 processors per node, and an Infiniband interconnect for inter-node communications.

### C.4. Additional scaling tests

VAMPIRE has already been shown to scale well for a generic system with nearest neighbour exchange interactions, but in order to verify the general usefulness of the parallelization we have also considered an extended system including the demagnetization field calculation, and a system using *ab initio* parameters for FePt [29] which includes a long-ranged exchange interaction extending over several unit cells, as shown in figure 17.

The simulations in figure 17(a) including the demagnetizing fields use a  $2 \times 2$  unit cell macrocell size, updated every 10 steps, for system sizes of approximately  $10^4$ ,  $10^5$  and  $10^6$  spins respectively. The high spatial and temporal resolution of the demagnetizing field calculation are in some sense a worst case scenario as these are probably not needed for most problems, and so for these simulations, calculation of the demagnetizing fields dominates the run time. Nevertheless, the scaling of the code remains very good, showing the effectiveness of the parallelization of the demagnetizing field calculation. For small system sizes the scaling breaks down as the number of macrocells approaches the number of processors. Here the scaling is limited by the time required to update all of the cell magnetizations and the time required to calculate the contributions of the reduced number of local macrocells. For all problem sizes scaling begins to reduce somewhat due to limitations in the network bandwidth.



For the long-ranged exchange interactions shown in figure 17(b) the scaling is only good within a single node (up to 8 processors). Above that, the larger the system the worse the scaling. This arises due to the large amount of data which has to be shared between processors. For the long-ranged exchange interaction, each atom to be simulated must know the spin configurations of over 1000 neighbouring atoms. In a parallel simulation these spin directions must be passed between processors twice per time step, which is a bandwidth intensive operation. Thus, the reduced scaling, particularly with larger system size, is due to saturation of the network link. Due to the long range nature of the exchange interaction, memory use also becomes an issue in terms of storing the neighbour list. For the largest simulation size of around  $1.6 \times 10^6$  atoms, 33GB of RAM is required to store all the interactions. However, due to the parallel system generation this divides nicely between all of the processors, so the memory required per processor is quite reasonable for larger numbers of processors.

## References

- [1] Ising E 1925 *Z. Phys.* **31** 253
- [2] Watson R E, Blume M and Vineyard G H 1969 *Phys. Rev.* **181** 811
- [3] Binder K and Rauch H 1969 *Z. Phys.* **219** 201
- [4] Kodama R and Berkowitz A 1999 *Phys. Rev. B* **59** 6321
- [5] Mitsumata C, Sakuma A and Fukamichi K 2003 *Phys. Rev. B* **68** 014437
- [6] Nowak U 2007 Classical spin models *Handbook of Magnetism and Advanced Magnetic Materials* (New York: Wiley)
- [7] Boerner E, Chubykalo-Fesenko O, Chantrell R, Heinonen O and Mryasov O 2005 *IEEE Trans. Magn.* **41** 936
- [8] Skubic B, Hellsvik J, Nordström L and Eriksson O 2008 *J. Phys.: Condens. Matter* **20** 315203
- [9] Ostler T *et al* 2011 *Phys. Rev. B* **84**
- [10] Ostler T A *et al* 2012 *Nature Commun.* **3** 666
- [11] Radu I *et al* 2011 *Nature* **472** 205
- [12] Iglesias Ò, Batlle X and Labarta A 2005 *Phys. Rev. B* **72** 212401
- [13] Evans R F L, Bate D, Chantrell R W, Yanes R and Chubykalo-Fesenko O 2011 *Phys. Rev. B* **84** 092404
- [14] Evans R F L, Yanes R, Mryasov O, Chantrell R W and Chubykalo-Fesenko O 2009 *Europhys. Lett.* **88** 57004
- [15] Ali M, Marrows C H, Al-Jawad M, Hickey B J, Misra A, Nowak U and Usadel K-D 2003 *Phys. Rev. B* **68** 214420
- [16] Garanin D and Kachkachi H 2003 *Phys. Rev. Lett.* **90** 065504
- [17] Yanes R, Chubykalo-Fesenko O, Kachkachi H, Garanin D, Evans R and Chantrell R 2007 *Phys. Rev. B* **76** 064416
- [18] Ma P-W, Dudarev S L and Woo C H 2012 *J. Appl. Phys.* **111** 07D114
- [19] Evans R, Nowak U, Dorfbauer F, Schrefl T, Mryasov O, Chantrell R W and Grochola G 2006 *J. Appl. Phys.* **99** 08G703
- [20] Evans R, Dorfbauer F, Mryasov O, Chubykalo-Fesenko O, Schrefl T and Chantrell R 2007 *IEEE Trans. Magn.* **43** 3106
- [21] Ho P, Evans R F L, Chantrell R W, Han G, Chow G-M and Chen J 2011 *IEEE Trans. Magn.* **47** 2646
- [22] Chureemart P, Evans R F L and Chantrell R W 2011 *Phys. Rev. B* **83**
- [23] Asselin P, Evans R F L, Barker J, Chantrell R W, Yanes R, Chubykalo-Fesenko O, Hinzke D and Nowak U 2010 *Phys. Rev. B* **82** 054415
- [24] Barker J, Evans R F L, Chantrell R W, Hinzke D and Nowak U 2010 *Appl. Phys. Lett.* **97** 192504
- [25] Yanes R, Chubykalo-Fesenko O, Evans R F L and Chantrell R W 2010 *J. Phys. D: Appl. Phys.* **43** 474009
- [26] Bauer D S G, Mavropoulos P, Lounis S and Blügel S 2011 *J. Phys.: Condens. Matter* **23** 394204
- [27] Fal T J, Plumer M L, Whitehead J P, Mercer J I, van Ek J and Srinivasan K 2013 *Appl. Phys. Lett.* **102** 202404
- [28] Victora R H and Huang P-W 2013 *IEEE Trans. Magn.* **49** 751
- [29] Mryasov O, Nowak U, Guslienko K and Chantrell R W 2005 *Europhys. Lett.* **69** 805
- [30] Szunyogh L, Udvardi L, Jackson J, Nowak U and Chantrell R W 2011 *Phys. Rev. B* **83**
- [31] Szunyogh L, Lazarovits B, Udvardi L, Jackson J and Nowak U 2009 *Phys. Rev. B* **79**
- [32] Sandratskii L M and Mavropoulos P 2011 *Phys. Rev. B* **83** 174408
- [33] Kazantseva N, Hinzke D, Nowak U, Chantrell R, Atxitia U and Chubykalo-Fesenko O 2008 *Phys. Rev. B* **77** 184428
- [34] Atxitia U, Hinzke D, Chubykalo-Fesenko O, Nowak U, Kachkachi H, Mryasov O, Evans R F L and Chantrell R W 2010 *Phys. Rev. B* **82** 134440
- [35] Jourdan T, Marty A and Lançon F 2008 *Phys. Rev. B* **77** 224428
- [36] Garcia Sanchez F, Chubykalo-Fesenko O, Mryasov O and Chantrell R W 2006 *Physica B* **372** 328
- [37] <http://math.nist.gov/oommf>
- [38] Scholz W, Fidler J, Schrefl T, Suess D, Dittich R, Forster H and Tsiantos V 2003 *Comput. Mater. Sci.* **28** 366
- [39] Fischbacher T, Franchin M, Bordignon G and Fangohr H 2007 *IEEE Trans. Magn.* **43** 2896
- [40] Mistral Q, van Kampen M, Hrkac G, Kim J-V, Devolder T, Crozat P, Chappert C, Lagae L and Schrefl T 2008 *Phys. Rev. Lett.* **100** 257201
- [41] Dobin A Y and Richter H J *Appl. Phys. Lett.* **89** 062512
- [42] Schrefl T, Fidler J, Kirk K J and Chapman J N 1997 *J. Magn. Magn. Mater.* **175** 193
- [43] Kryder M H, Gage E C, McDaniel T W, Challener W A, Rottmayer R E, Ju G, Hsia Y-T and Erden M F 2008 *Proc. IEEE* **96** 1810
- [44] Beaupaire E, Merle J C, Daunois A and Bigot J Y 1996 *Phys. Rev. Lett.* **76** 4250
- [45] Stanciu C, Tsukamoto A, Kimel A, Hansteen F, Kirilyuk A, Itoh A and Rasing T 2007 *Phys. Rev. Lett.* **99** 217204
- [46] O'Grady K, Fernandez-Outon L E and Vallejo-Fernandez G 2010 *J. Magn. Magn. Mater.* **322** 883
- [47] Ikeda S, Miura K, Yamamoto H, Mizunuma K, Gan H D, Endo M, Kanai S, Hayakawa J, Matsukura F and Ohno H 2010 *Nature Mater.* **9** 721
- [48] Jamet M, Wernsdorfer W, Thirion C, Mailly D, Dupuis V, Mélinon P and Pérez A 2001 *Phys. Rev. Lett.* **86** 4676
- [49] Klemmer T, Shukla N, Liu C, Wu X, Svedberg E, Mryasov O, Chantrell R, Weller D, Tanase M and Laughlin D 2002 *Appl. Phys. Lett.* **81** 2220
- [50] Garanin D 1997 *Phys. Rev. B* **55** 3050
- [51] Evans R F L, Hinzke D, Atxitia U, Nowak U, Chantrell R W and Chubykalo-Fesenko O 2012 *Phys. Rev. B* **85** 014433

- [52] Atxitia U, Nieves P and Chubykalo-Fesenko O 2012 *Phys. Rev. B* **86** 104414
- [53] McDaniel T W 2012 *J. Appl. Phys.* **112** 093920
- [54] Atxitia U and Chubykalo-Fesenko O 2011 *Phys. Rev. B* **84** 144414
- [55] Sultan M, Atxitia U, Melnikov A, Chubykalo-Fesenko O and Bovensiepen U 2012 *Phys. Rev. B* **85** 184407
- [56] Atxitia U, Ostler T, Barker J, Evans R F L, Chantrell R W and Chubykalo-Fesenko O 2013 *Phys. Rev. B* **87** 224417
- [57] Hovorka O, Devos S, Coopman Q, Fan W J, Aas C J, Evans R F L, Chen X, Ju G and Chantrell R W 2012 *Appl. Phys. Lett.* **101** 052406
- [58] Details available from <http://vampire.york.ac.uk>
- [59] Dorfbauer F, Schrefl T, Kirschner M, Hrkac G, Suess D, Ertl O and Fidler J 2006 *J. Appl. Phys.* **99** 08G706
- [60] Jiles D 1991 *Introduction to Magnetism and Magnetic Materials* (London: Chapman and Hall)
- [61] Schwarz K, Mohn P, Blaha P and Kubler J 1984 *J. Phys. F: Met. Phys.* **14** 2659
- [62] Liechtenstein A, Katsnelson M, Antropov V and Gubanov V 1987 *J. Magn. Magn. Mater.* **67** 65
- [63] Udvardi L, Szunyogh L, Palotas K and Weinberger P 2003 *Phys. Rev. B* **68** 104436
- [64] Antal A, Lazarovits B, Udvardi L, Szunyogh L, Ujfalussy B and Weinberger P 2008 *Phys. Rev. B* **77** 174429
- [65] Ebert H and Mankovsky S 2009 *Phys. Rev. B* **79** 045209
- [66] Drautz R and Fähnle M 2004 *Phys. Rev. B* **69** 104404
- [67] Drautz R and Fähnle M 2005 *Phys. Rev. B* **72** 212405
- [68] Deák A, Szunyogh L and Ujfalussy B 2011 *Phys. Rev. B* **84** 224413
- [69] Skomski R 2012 *Simple Models of Magnetism* (Oxford: Oxford University Press)
- [70] Bruno P 1993 Physical origins and theoretical models of magnetic anisotropy *Magnetismus von Festkörpern und grenzflächen* vol 24, ed P H Dederichs, P Grünberg and W Zinn (Jülich: Forschungszentrums Jülich) pp 24.1–24.28
- [71] Hohenberg P and Kohn W 1964 *Phys. Rev. B* **136** 864
- [72] Kohn W and Sham L J 1965 *Phys. Rev. A* **140** 1133
- [73] von Barth U and Hedin L 1972 *J. Phys. C: Solid State Phys.* **5** 1629
- [74] Kresse G and Furthmüller J 1996 *Phys. Rev. B* **54** 11169
- [75] Segall M D, Lindan P J D, Probert M J, Pickard C J, Hasnip P J, Clark S J and Payne M C 2002 *J. Phys.: Condens. Matter.* **14** 2717
- [76] Payne M C, Teter M P, Allan D C, Arias T A and Joannopoulos J D 1992 *Rev. Mod. Phys.* **64** 1045
- [77] Soler J M, Artacho E, Gale J D, García A, Junquera J, Ordejón P and Sánchez-Portal D 2002 *J. Phys.: Condens. Matter.* **14** 2745
- [78] Zabloudil J, Hammerling R, Szunyogh L and Weinberger P 2005 *Electron Scattering in Solid Matter* (Berlin: Springer)
- [79] Ebert H, Ködderitzsch D and Minár J 2011 *Rep. Prog. Phys.* **74** 096501
- [80] Katsnelson M I and Lichtenstein A I 2000 *Phys. Rev. B* **61** 8906
- [81] Pajda M, Kudrnovský J, Turek I, Drchal V and Bruno P 2001 *Phys. Rev. B* **64** 174402
- [82] Sandratskii L M 1986 *Phys. Status Solidi b* **136** 167
- [83] Halilov S V, Eschrig H, Perlov A Y and Oppeneer P M 1998 *Phys. Rev. B* **58** 293
- [84] Ležaić M, Mavropoulos P, Bihlmayer G and Blügel S 2013 *Phys. Rev. B* **88** 134403
- [85] Uhl M and Siberchicot B 1999 *J. Phys.: Condens. Matter.* **7** 4227
- [86] Skubic B, Peil O, Hellsvik J, Nordblad P, Nordström L and Eriksson O 2009 *Phys. Rev. B* **79** 024411
- [87] Aas C J, Hasnip P, Cuadrado R, Plotnikova E, Szunyogh L, Udvardi L and Chantrell R W 2013 *Phys. Rev. B* **88** 174409
- [88] Sato K *et al* 2010 *Rev. Mod. Phys.* **82** 1633
- [89] Ležaić M, Mavropoulos P, Blügel S and Ebert H 2011 *Phys. Rev. B* **83** 094434
- [90] Danan H 1968 *J. Appl. Phys.* **39** 669
- [91] Garanin D A 1996 *Phys. Rev. B* **53** 11593
- [92] Wysin G M 2000 *Phys. Rev. B* **62** 3251
- [93] Colarieti-Tosti M, Simak S, Ahuja R, Nordström L, Eriksson O, Åberg D, Edvardsson S and Brooks M 2003 *Phys. Rev. Lett.* **91** 157201
- [94] Fan W J, Evans R F L, Hancock Y and Chantrell R W 2011 *J. Appl. Phys.* **109** 07B752
- [95] Schröder C, Schmidt H-J, Schnack J and Luban M 2005 *Phys. Rev. Lett.* **94** 207203
- [96] Schröder C, Nojiri H, Schnack J, Hage P, Luban M and Kögerler P 2005 *Phys. Rev. Lett.* **94** 017205
- [97] Dorfbauer F, Evans R, Kirschner M, Chubykalo-Fesenko O, Chantrell T and Schrefl T 2007 *J. Magn. Magn. Mater.* **316** 791
- [98] Landau LD and Lifshitz E M 1935 *Phys. Z. Sow.* **8** 153
- [99] Gilbert T 1955 *Phys. Rev.* **100** 1243
- [100] Karakurt S, Chantrell R W and Nowak U 2007 *J. Magn. Magn. Mater.* **316** e280
- [101] Fähnle M, Seib J and Illg C 2010 *Phys. Rev. B* **82** 144405
- [102] Dobin A Y and Victora R H 2004 *Phys. Rev. Lett.* **92** 257204
- [103] Ellis M, Ostler T A and Chantrell R W 2012 *Phys. Rev. B* **86** 174418
- [104] Brown W F Jr 1979 *IEEE Trans. Magn.* **15** 1196
- [105] García-Palacios J L and Lázaro F J 1998 *Phys. Rev. B* **58** 14937
- [106] Lyberatos A, Berkov D V and Chantrell R W 1999 *J. Phys.: Condens. Matter.* **5** 8911
- [107] Nowak U, Mryasov O, Wieser R, Guslienko K and Chantrell R 2005 *Phys. Rev. B* **72**
- [108] Matsumoto M and Nishimura T 1998 *ACM Trans. Model. Comput. Simul.* **8** 3
- [109] Marsaglia G and Tsang W-W 2000 *J. Stat. Softw.* **5** 1–7 ([www.jstatsoft.org/v05/i08](http://www.jstatsoft.org/v05/i08))
- [110] Atxitia U, Chubykalo-Fesenko O, Chantrell R, Nowak U and Rebei A 2009 *Phys. Rev. Lett.* **102** 057203
- [111] Schellekens A J and Koopmans B 2013 *Phys. Rev. Lett.* **110** 217204
- [112] Grossmann B and Rancourt D 1996 *Phys. Rev. B* **54** 12294
- [113] Ma P-W, Dudarev S L and Woo C H 2012 *Phys. Rev. B* **85** 184301
- [114] Battiatto M, Carva K and Oppeneer P M 2010 *Phys. Rev. Lett.* **105** 027203
- [115] Schellekens A J, Verhoeven W, Vader T N and Koopmans B 2013 *Appl. Phys. Lett.* **102** 252408
- [116] Berkov D V and Gorn N L 2002 *J. Phys.: Condens. Matter.* **14** L281
- [117] d'Aquino M, Serpico C and Miano G 2005 *J. Comput. Phys.* **209** 730
- [118] Mentink J, Tretyakov M, Fasolino A, Katsnelson M and Rasing T 2010 *J. Phys.: Condens. Matter.* **22** 176001



- [119] d'Aquino M, Serpico C, Coppola G, Mayergoz I D and Bertotti G 2006 *J. Appl. Phys.* **99** 08B905
- [120] Kloeden P E and Platen E 1995 *Numerical Solution of Stochastic Differential Equations* (Berlin: Springer)
- [121] Hannay J D 2001 Computational simulations of thermally activated magnetisation dynamics at high frequencies *PhD Thesis* School of Informatics, University of Wales
- [122] Metropolis N, Rosenbluth A W, Rosenbluth M N, Teller A H and Teller E 1953 *J. Chem. Phys.* **21** 1087
- [123] Hinzke D and Nowak U 1999 *Comput. Phys. Commun.* **121** 334
- [124] <http://www.ctcms.nist.gov/~rdm/mumag.org>
- [125] Stoner E C and Wohlfarth E P 1948 *Philos. Trans. R. Soc. A* **240** 599
- [126] Binder K 1969 *Phys. Lett. A* **30** 273
- [127] Yuan S W and Bertram H N 1992 *IEEE Trans. Magn.* **28** 2031
- [128] Berkov D V, Ramstöck K and Hubert A 1993 *Phys. Status Solidi a* **137** 207
- [129] Kurzak J 2005 *J. Parallel. Distr. Com.* **65** 870
- [130] Aharoni A 1998 *J. Appl. Phys.* **83** 3432
- [131] Osborn J A 1945 *Phys. Rev.* **67** 351
- [132] Hinzke D and Nowak U 2011 *Phys. Rev. Lett.* **107** 027205
- [133] Brown W F 1945 *Rev. Mod. Phys.* **17** 15