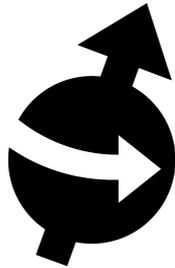


VAMPIRE

User Manual



VAMPIRE

User Manual

Software version 6.0

Manual written by Richard F. L. Evans, Daniel Meilak, Andrew Naden and Andreas Biternas.

Copyright © 2018 Department of Physics, The University of York, Heslington, York, YO10 5DD. All rights Reserved.

The VAMPIRE software package is principally developed and maintained by Richard F. L. Evans. Code contributors: Sarah Jenkins, Andrea Meo, Andrew Naden, Daniel Meilak, Samuel Morris, Matthew Ellis, Oscar David Arbeláez Echeverri, Weijia Fan, Phanwadee Chureemart, Rory Pond, Pawel Sobieszczyk, Joe Barker, Thomas Ostler, Andreas Biternas, Roy W Chantrell, Wu Hong-Ye

The entire VAMPIRE package is available under the GNU General Public License. You are free to use vampire for personal, academic and commercial research, and to modify the source code as you wish. For details of the licence, check the README file in the source code or consult www.gnu.org/copyleft/gpl.html.

The VAMPIRE source code is available from www.github.com/richard-evans/vampire. This manual, software features, tutorials and more information is available from

the VAMPIRE webpage at <http://vampire.york.ac.uk/>

Table of Contents

Table of Contents	3
Preface	12
Introducing VAMPIRE	12
1 Background theory	13
Atomistic Spin Models	13
The spin Hamiltonian	14
Spin Dynamics	14
Citations	15
2 Installation	17
System Requirements	17
Binary installation	17
Compiling from source	18
Compiling on Linux	18
Compiling on Mac OSX	18
Compiling on Windows	19
Compiling for ARCHER	19
Compiling for GPU	19
3 Running the code	20
Running on Unix/Linux and macOS	20
Running on Windows	20
Log file	20
4 Getting Started	22
Feature Overview	22
Input and Output Files	23
Sample input files	23
5 Visualization	26

Getting started	26
Atomic visualization with rasmol	27
Atomic visualization with POV-Ray	27
General Customisation options	28
slice	28
slice-void	28
slice-sphere	28
slice-cylinder	29
remove-materials	29
POV-Ray Customisation options	29
frame-start	29
frame-final	29
camera-position	29
camera-look-at	29
camera-zoom	29
background-colour	30
atom-sizes	30
arrow-sizes	30
colourmap	30
custom-colourmap	32
3D	32
vector-z	32
vector-x	33
afm	33
Micromagnetic visualization with POV-Ray	33
Visualization Movies	33
6 Unit Cell Files	34
The unit cell file format	34
Example: Simple Cubic System	36
7 Input File Command Reference	38
System Generation	38
create:full	38
create:cube	38
create:cylinder	38
create:ellipsoid	38
create:sphere	38

create:truncated-octahedron	38
create:particle	38
create:particle-array	39
create:voronoi-film	39
create:voronoi-size-variance	39
create:voronoi-row-offset	39
create:voronoi-random-seed	39
create:voronoi-rounded-grains	39
create:voronoi-rounded-grains-area	40
create:particle-parity	40
create:crystal-structure	40
create:crystal-sublattice-materials	40
create:single-spin	40
create:periodic-boundaries-x	40
create:periodic-boundaries-y	40
create:periodic-boundaries-z	40
create:periodic-boundaries	40
create:select-material-by-height	40
create:select-material-by-geometry	40
create:fill-core-shell-particles	41
create:interfacial-roughness	41
create:material-interfacial-roughness	41
create:interfacial-roughness-random-seed	41
create:interfacial-roughness-number-of-seed-points	41
create:interfacial-roughness-type	41
create:interfacial-roughness-seed-radius	41
create:interfacial-roughness-seed-radius-variance	41
create:interfacial-roughness-mean-height	41
create:interfacial-roughness-maximum-height	41
create:interfacial-roughness-height-field-resolution	41
create:alloy-random-seed	41
create:grain-random-seed	41
create:dilution-random-seed	41
create:intermixing-random-seed	42
create:spin-initialisation-random-seed	42
System dimensions	42
dimensions:unit-cell-size	42
dimensions:unit-cell-size-x	42

dimensions:unit-cell-size-y	42
dimensions:unit-cell-size-z	42
dimensions:system-size	42
dimensions:system-size-x	42
dimensions:system-size-y	42
dimensions:system-size-z	42
dimensions:particle-size	42
dimensions:particle-spacing	42
dimensions:particle-shape-factor-x	42
dimensions:particle-shape-factor-y	42
dimensions:particle-shape-factor-z	43
dimensions:particle-array-offset-x	43
dimensions:particle-array-offset-y	43
dimensions:macro-cell-size	43
Exchange calculation	43
exchange:interaction-range double default [1.0]	43
exchange:function string default [nearest-neighbour]	43
exchange:decay-multiplier double default [1.0]	44
exchange:decay-length double default [0.4]	44
exchange:decay-shift double default [0.0]	44
exchange:ucc-exchange-parameters[i][j] std::vector < std::vector <ex- change_parameters_t> > [i][j] default [1.0, 0.4, 0.0]	44
exchange:dmi-cutoff-range double default [1.0]	44
exchange:ab-initio flag default false	44
Anisotropy calculation	44
anisotropy:surface-anisotropy-threshold	44
anisotropy:surface-anisotropy-nearest-neighbour-range	44
anisotropy:enable-bulk-neel-anisotropy	44
anisotropy:neel-anisotropy-exponential-range	45
anisotropy:neel-anisotropy-exponential-factor	45
Dipole field calculation	45
dipole:solver	45
Simulation Control	47
sim:integrator	47
sim:program	48
sim:enable-dipole-fields	50
sim:enable-fmr-field	50
sim:enable-fast-dipole-fields	50

sim:dipole-field-update-rate	50
sim:time-step	50
sim:total-time-steps	50
sim:loop-time-steps	50
sim:time-steps-increment	50
sim:equilibration-time-steps	50
sim:simulation-cycles	50
sim:maximum-temperature	50
sim:minimum-temperature	51
sim:equilibration-temperature	51
sim:temperature	51
sim:temperature-increment	51
sim:cooling-time	51
sim:laser-pulse-temporal-profile	51
sim:laser-pulse-time	51
sim:laser-pulse-power	51
sim:second-laser-pulse-time	51
sim:second-laser-pulse-power	51
sim:second-laser-pulse-maximum-temperature	51
sim:second-laser-pulse-delay-time	51
sim:two-temperature-heat-sink-coupling	51
sim:two-temperature-electron-heat-capacity	51
sim:two-temperature-phonon-heat-capacity	51
sim:two-temperature-electron-phonon-coupling	52
sim:cooling-function	52
sim:applied-field-strength	52
sim:maximum-applied-field-strength	52
sim:equilibration-applied-field-strength	52
sim:applied-field-strength-increment	52
sim:applied-field-angle-theta	52
sim:applied-field-angle-phi	52
sim:applied-field-unit-vector	52
sim:demagnetisation-factor	52
sim:integrator-random-seed	53
sim:constraint-rotation-update	53
sim:constraint-angle-theta	53
sim:constraint-angle-theta-minimum	53
sim:constraint-angle-theta-maximum	53

sim:constraint-angle-theta-increment	53
sim:constraint-angle-phi	53
sim:constraint-angle-phi-minimum	53
sim:constraint-angle-phi-maximum	53
sim:constraint-angle-phi-increment	54
montecarlo:algorithm	54
montecarlo:constrain-by-grain	54
sim:checkpoint	54
sim:preconditioning-steps	55
sim:electrical-pulse-time	55
sim:electrical-pulse-rise-time	55
sim:electrical-pulse-fall-time	55
Data output	55
output:time-steps	55
output:real-time	55
output:temperature	55
output:applied-field-strength	55
output:applied-field-unit-vector	56
output:applied-field-alignment	56
output:material-applied-field-alignment	56
output:magnetisation	56
output:magnetisation-length	56
output:mean-magnetisation-length	56
output:mean-magnetisation	56
output:material-magnetisation	56
output:material-mean-magnetisation-length	56
output:material-mean-magnetisation	57
output:total-torque	57
output:mean-total-torque	57
output:constraint-phi	57
output:constraint-theta	57
output:material-mean-torque	57
output:mean-susceptibility	57
output:mean-material-susceptibility	58
output:material-standard-deviation	58
output:electron-temperature	58
output:phonon-temperature	58
output:total-energy	58

output:mean-total-energy	58
output:anisotropy-energy	58
output:mean-anisotropy-energy	58
output:exchange-energy	58
output:mean-exchange-energy	58
output:applied-field-energy	58
output:mean-applied-field-energy	58
output:magnetostatic-energy	58
output:mean-magnetostatic-energy	58
output:material-total-energy	58
output:material-mean-total-energy	59
output:mean-specific-heat	59
output:material-mean-specific-heat	59
output:fractional-electric-field-strength	59
output:mpi-timings	59
output:gnuplot-array-format	59
output:output-rate	59
output:precision	59
output:fixed-width	59
output:column-headers	60
Configuration output	60
config:atoms	60
config:atoms-output-rate	60
config:atoms-min-x	60
config:atoms-min-y	60
config:atoms-min-z	60
config:atoms-max-x	60
config:atoms-max-y	61
config:atoms-max-z	61
config:macro-cells	61
config:macro-cells-output-rate	61
config:output-format	61
config:output-mode	61
config:output-nodes	62
8 Material File Command Reference	63
Material File Parameters	63
material:num-materials	63

material:material-name	63
material:damping-constant	63
material:exchange-matrix	64
material:exchange-matrix-1st-nn	64
material:exchange-matrix-2nd-nn	64
material:exchange-matrix-3rd-nn	65
material:exchange-matrix-4th-nn	65
material:biquadratic-exchange-matrix	65
material:atomic-spin-moment	65
material:uniaxial-anisotropy-constant	66
material:second-order-uniaxial-anisotropy-constant	66
material:fourth-order-uniaxial-anisotropy-constant	66
material:cubic-anisotropy-constant	66
material:fourth-order-cubic-anisotropy-constant	66
material:uniaxial-anisotropy-direction	66
material:surface-anisotropy-constant	67
material:neel-anisotropy-constant	67
material:lattice-anisotropy-constant	67
material:lattice-anisotropy-file	68
material:voltage-controlled-magnetic-anisotropy-coefficient	68
material:relative-gamma	68
material:initial-spin-direction	68
material:material-element	68
material:geometry-file	68
material:alloy-host	69
material:alloy-fraction	69
material:minimum-height	69
material:maximum-height	69
material:core-shell-size	69
material:interface-roughness	70
material:intermixing	70
material:density	70
material:continuous	70
material:fill-space	70
material:couple-to-phononic-temperature	70
material:temperature-rescaling-exponent	71
material:temperature-rescaling-curie-temperature	71
material:non-magnetic	71

material:unit-cell-category	71
Bibliography	72

Introducing VAMPIRE

VAMPIRE is a state-of-the-art atomistic simulator for magnetic nanomaterials. This software is the culmination of several years of continuous development, with an aim to make atomistic simulation of magnetic materials routinely available to the non-specialist researcher. Before now, using atomistic models to simulate magnetic systems required in depth and technical knowledge of the underlying theoretical methods, computer programming skills and the ability to debug and understand intricate computational problems. The code is designed with ease of use in mind, and includes an extensive set of input parameters to control the simulations through a plain text input file. Subject to future funding it is also hoped to develop graphical user interfaces for macOSTM and WindowsTM which should make using the code more accessible.

The VAMPIRE project is still very much under active development, with an open development of all code features. The features are always available during the development stages from the develop branch of the code, but with the caveat that they are not always fully reliable. Feedback of any bugs or errors to the VAMPIRE developers is always welcome, as well as any feature requests or enhancements.

We hope that as the VAMPIRE project develops it will become a useful tool for the magnetics community for specialists and non-specialists alike.

1 Background theory

While the underlying theory behind the atomistic spin model is well known in the scientific literature, in the following a very brief overview of the fundamental theory is presented for the benefit of those who do not wish to study the methods in great detail. If more information is required then a comprehensive review of the methods implemented in VAMPIRE is available from the project website.

Atomistic Spin Models

Atomistic spin models form the natural limit of two distinct approaches, namely micromagnetics and ab-initio models of the electronic structure. In micromagnetics a material is discretized into small domains where the magnetization is assumed to be fully ordered within it. If the micromagnetic cell size is reduced to less than 1 nm, then the magnetization is no longer a true continuum, but a discrete entity considering localized moments on individual atoms. Similarly, when the electronic properties of the system are considered, the quantum mechanical properties can be mapped onto atomic cores in a manner similar to molecular dynamics, where the effective properties can often be treated with a classical approximation.

The advantage of the atomistic model over micromagnetics is that it naturally deals with atomic ordering and variation of local properties seen in real materials, such as interfaces, defects, roughness etc. The discrete formulation also allows the simulation of high temperatures above and beyond the Curie temperature, where the usual continuum micromagnetic approach breaks down. Such effects or often central to current problems in magnetism such as materials for spin electronics, heat assisted magnetic recording or ultrafast laser processes. Similarly for ab-initio calculations, mapping onto an effective spin model allows apply the full quantum mechanical deal of the properties to much larger systems and the consideration of dynamic effects on much longer timescales.

The Spin Hamiltonian

The basis of the atomistic spin model is the spin Hamiltonian, which describes the fundamental spin-dependent interactions at the atomic level (neglecting the effects of potential and kinetic energy and electron correlations). The spin Hamiltonian is typically defined as

$$\mathcal{H} = - \sum_{i < j} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j - k_2 \sum_i S_z^2 - \mu_S \sum_i \mathbf{B}_{\text{app}} \cdot \mathbf{S}_i$$

describing exchange, uniaxial anisotropy and applied field contributions respectively. Important parameters are the Heisenberg exchange J_{ij} , the anisotropy constant k_2 and the atomic spin moment, μ_S . \mathbf{S}_i is a unit vector which describes the orientation of the local spin moment. In most magnetic materials the exchange interactions are the dominant contribution, usually by two orders of magnitude, and gives rise to the atomic ordering of the spin directions. For ferromagnetic materials (parallel alignment of spins) $J_{ij} > 0$, while for anti-ferromagnetic materials (antiparallel alignment of spins), $J_{ij} < 0$.

While the exchange interaction determines the ordering of the spins, it is usually isotropic, and so there is no preferential orientation of all the spins in the system. Most magnetic materials are anisotropic, that is the spins have a preferred orientation in space, which arises at the atomic level due to the local crystal environment, hence its full name of magnetocrystalline anisotropy. In the model this is most commonly uniaxial anisotropy, where the spins prefer to lie along a single preferred axis, known as the easy axis. The strength of the anisotropy is determined by the anisotropy constant, in our case k_2 , where positive value prefer alignment along the z -axis, while negative values prefer alignment around the $x - y$ plane.

The last term describes the coupling of the spin system to an externally applied field, \mathbf{B}_{app} , or Zeeman field. The applied field is used to reverse the orientation of the spins, and can be used in the simulation to calculate hysteresis loops, for example.

Spin Dynamics

The spin Hamiltonian describes the energetics of the system, but says nothing about the dynamic behaviour. For that the Landau-Lifshitz-Gilbert (LLG) equation

is used to describe the dynamics of atomic spins. The LLG is given by

$$\frac{\partial \mathbf{S}_i}{\partial t} = -\frac{\gamma}{(1 + \lambda^2)} [\mathbf{S}_i \times \mathbf{B}_{\text{eff}}^i + \lambda \mathbf{S}_i \times (\mathbf{S}_i \times \mathbf{B}_{\text{eff}}^i)] \quad (1.1)$$

where \mathbf{S}_i is a unit vector representing the direction of the magnetic spin moment of site i , γ is the gyromagnetic ratio and $\mathbf{B}_{\text{eff}}^i$ 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:

$$\mathbf{B}_{\text{eff}}^i = -\frac{1}{\mu_S} \frac{\partial \mathcal{H}}{\partial \mathbf{S}_i} \quad (1.2)$$

where μ_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. The LLG is integrated numerically using the Heun numerical scheme, which allows the time evolution of the spin system to be simulated.

Citations

If you use VAMPIRE for your research, it is helpful to acknowledge the authors of the code by citing relevant papers and include a statement in the paper such as the following:

The simulations in this work made use of the VAMPIRE software package [1]

and add a footnote reading:

[1] VAMPIRE software package version 5.0 available from <https://vampire.york.ac.uk>

In addition, it is recommended for reproducibility that you include the githash for the specific version of the code, which enables someone to checkout the specific version of the code used for the simulations.

[1] VAMPIRE software package version 5.0 available from <https://vampire.york.ac.uk> (Version aa842a409c68d6724e156df6cab0bcaa172f5f41)

If you use the code, please cite the following article:

Atomistic spin model simulations of magnetic nanomaterials

R. F. L. Evans, W. J. Fan, P. Chureemart, T. A. Ostler, M. O. A. Ellis and R. W. Chantrell

J. Phys.: Condens. Matter 26, 103202 (2014)

If you use the constrained Monte Carlo method, in addition please cite:

Constrained Monte Carlo method and calculation of the temperature dependence of magnetic anisotropy

P. Asselin, R. F. L. Evans, J. Barker, R. W. Chantrell, R. Yanes, O. Chubykalo-Fesenko, D. Hinzke and U. Nowak

Phys. Rev. B. 82, 054415 (2010)

If you use the temperature rescaling method please cite:

Quantitative simulation of temperature-dependent magnetization dynamics and equilibrium properties of elemental ferromagnets

R. F. L. Evans, U. Atxitia, and R. W. Chantrell

Phys. Rev. B 91, 144425 (2015)

2 Installation

This chapter covers the requirements, installation and support for VAMPIRE on different platforms.

System Requirements

VAMPIRE is designed to be generally portable and compilable on Linux, Unix, Mac OSX and Windows with a range of different compilers. By design the software has a very minimal dependence on external libraries to aid compilation on the widest possible range of platforms without needing to first install and configure a large number of other packages. VAMPIRE is designed to be maximally efficient on high performance computing clusters and scalable to thousands of processors, and as such is the recommended platform if you have access to appropriate resources.

Hardware Requirements

VAMPIRE has been successfully tested on a wide variety of x86 and power PC processors. Memory requirements are generally relatively modest for most systems, though larger simulations will require significantly more memory. VAMPIRE is generally computationally limited, and so the faster the clock speed and number of processor cores the better.

Binary installation

Compiled binaries of the latest release version are available to download from:

<https://vampire.york.ac.uk/download/>

for Linux and MacTM OS X platforms. For the Linux and Mac OS X releases, a simple installation script `install.sh` installs the binary in `/opt/vampire/` and appends the directory to your environment path. On Windows the recommended method is to use the Linux subsystem for windows developer feature which adds

a linux subsystem that is capable of running the standard linux binary. A copy of qvoronoi is integrated into VAMPIRE for generating granular structures.

Compiling from source

The best way to get the vampire source code is using git, a distributed version control program which enables changes in the code to be tracked. Git is readily available on linux (git-core package on ubuntu) and Mac (via MacPorts). To get vampire from the Github repository checkout your own copy of the repository using:

```
git clone git://github.com/richard-evans/vampire.git
```

This way, updates to the code can be easily merged with the downloaded version. Compiling is generally as easy as running make in Unix platforms. In addition, on a multicore processor compilation can be parallelised using the `-j N_t` option, where N_t is the number of threads to use.

Compiling on Linux

In order to compile in linux, a working set of development tools are needed, which on ubuntu includes the packages `build-essential` and `g++`. VAMPIRE should compile without issue following a simple `make` command in the source directory.

For the parallel version, a working installation of `openmpi` is recommended, which must usually include a version of the development tools (`openmpi-bin` and `openmpi-dev` packages on ubuntu). Compilation is usually straightforward using `make parallel`.

Compiling on Mac OSX

With OS X, compilation from source requires a working installation of Xcode, available for free from the Mac App Store. In addition command line tools must also be installed. A working installation of MacPorts is recommended to provide access to a wide range of open source libraries and tools such as `openmpi`, `rasmol` and `povray`. For the serial version, compilation is the same as for linux, following a simple `make serial-llvm` command in the source directory.

Similarly for the parallel version, `openmpi` needs to be installed via MacPorts, and compilation is usually straightforward using `make parallel-llvm`.

Compiling on Windows

The recommended way to use vampire on Windows is to install Linux subsystem for windows 10 (see <https://docs.microsoft.com/en-us/windows/wsl/install-win10>). Older versions of windows are no longer supported. Once installed, you can download the serial linux binary as for linux and run as normal from the command line.

Compiling for ARCHER/Cray systems

ARCHER is the UK national supercomputer and includes custom compilers developed by Cray Inc. However, performance is generally better for the gnu compiler collection and so there is an optimized makefile option for compilation on the ARCHER and similar Cray XC30 systems. To compile, you need to swap the environment to the GNU compiler suite using module swap PrgEnv-cray PrgEnv-gnu. You can then compile with make parallel-archer which will compile a parallel binary.

Compiling for GPU acceleration with CUDA (beta)

The latest release includes a CUDA implementation for GPU accelerated atomistic spin dynamics. To compile the CUDA version of the code, you need to install the CUDA drivers and runtime. Once installed, compilation should be straightforward using make gcc-cuda. By default, the binary includes device code for a wide range of architectures. Depending on your device/card, you may need to modify the device code generation option in the makefile.

3 Running the code

To run the code in all version, you first need to specify an input file and material file, which must reside in the same directory where you run the code. Example files are available in the source code distribution, or from the Download section of the website (<http://vampire.york.ac.uk/download/index.html>).

Unix/Linux and macOS

In the directory including the input and material files, typing

```
./vampire-serial
```

will run the code in serial mode. For the parallel mode with openmpi,

```
mpirun -np 2 vampire-parallel
```

will run the code in parallel mode, on 2 CPUs. Increasing the `-np` argument will run on more cores.

Windows

Once you have installed Linux subsystem for Windows, you can run the code by launching bash for windows and following the instructions for Unix/Linux systems above.

Log file

When you run the program it will output some program information to screen to indicate that the program is running listing some details about the developers and some steps on the program execution. There are also options to tell the program to output simulation data to screen to see how the simulation is progressing. In addition, a log file is produced which provides more detail about the execution of

the program and progress including timestamps, and errors that may occur and certain performance metrics, such as time to update the dipole field or output a spin configuration file. A sample screen output header from the program is shown below.

```
          _
         ( )
-----
  \ \ / / _ ' | ' ' _ \ | ' _ \ | | ' _ / _ \
   \ V / ( | | | | | | | | | | | | | | | | | | | |
    \ / \ _ , _ | | | | | | | | | | | | | | | | | |
                                     | |
                                     | |
                                    | |
```

Version 5.0.0 Aug 25 2018 23:01:25

Git commit: 4377c4a8c3b2334decf6b3892542927fcaddebc7

Licensed under the GNU Public License(v2). See licence file for details.

Lead Developer: Richard F L Evans <richard.evans@york.ac.uk>

Contributors: Andrea Meo, Rory Pond, Weijia Fan,
Phanwadee Chureemart, Sarah Jenkins, Joe Barker,
Thomas Ostler, Andreas Biternas, Roy W Chantrell,
Wu Hong-Ye, Matthew Ellis, Razvan Ababei,
Sam Westmoreland, Oscar Arbelaez, Sam Morris

Compiled with: LLVM C++ Compiler
Compiler Flags:

Vampire includes a copy of the qhull library from C.B. Barber and
The Geometry Center and may be obtained via http from www.qhull.org.

4 Getting Started

VAMPIRE is a powerful software package, capable of simulating many different systems and the determination of parameters such as coercivity, Curie and Néel temperatures, reversal dynamics, statistical behaviour and more. This chapter contains an overview of the capabilities of VAMPIRE and how to use them.

Feature Overview

The features of the VAMPIRE code are split into three main categories: material parameters, structural parameters, and simulation parameters. Details of these parameters are given in the following chapters, but between them they define the parameters for a particular simulation.

Materials

Material parameters essentially define the magnetic properties of a class of atoms, including magnetic moments, exchange interactions, damping constants etc. VAMPIRE includes support for up to one hundred defined materials, and material parameters control the simulation of multilayers, random alloys, core shell particles and lithographically defined patterns.

Structures

Structural parameters define properties such as the system size, shape, particle size, or voronoi grain structures. In combination with material parameters they essentially define the system to be simulated.

Simulations

VAMPIRE includes a number of built-in simulations for determining the most common magnetic properties of a system,  for example Curie temperature,

hysteresis loops, or even a time series. Additionally the parameters for these simulations, such as applied field, maximum temperature, temperature increment, etc. can be set.

Input and Output Files

VAMPIRE requires at least two files to run a simulation, the input file and the material file. The input file defines all the properties of the simulated system, such as the dimensions or particle shape, as well as the simulation parameters and program output. The material file defines the properties of all the materials used in the simulation, and is usually given the .mat file extension. A sample material file Co.mat is included with the code which defines a minimum set of parameters for Co.

The output of the code includes a main output file, which records data such as the magnetisation, timesteps, temperature etc. The format of the output file is fully customisable, so that the amount of output data is limited to what is useful. In addition to the output file, the other main available outputs are spin configuration files, which with post-processing allow output of snapshots of the magnetic configurations during the simulation.

Sample input files

Sample input and output files are included in the source code distribution, but the files for a simple test simulation which computes the time dependence of the magnetisation of a cubic system are given here.

```
input
#-----
# Sample vampire input file to perform
# benchmark calculation for version 5.0
#
#-----

#-----
# Creation attributes:
#-----
create:crystal-structure = sc

#-----
```

```

# System Dimensions:
#-----
dimensions:unit-cell-size = 3.54 !A
dimensions:system-size-x = 7.7 !nm
dimensions:system-size-y = 7.7 !nm
dimensions:system-size-z = 7.7 !nm

#-----
# Material Files:
#-----
material:file = Co.mat

#-----
# Simulation attributes:
#-----
sim:temperature = 300.0
sim:time-steps-increment = 1000
sim:total-time-steps = 10000
sim:time-step = 1 !fs

#-----
# Program and integrator details
#-----
sim:program = benchmark
sim:integrator = llg-heun

#-----
# data output
#-----
output:real-time
output:temperature
output:magnetisation
output:magnetisation-length

screen:time-steps
screen:magnetisation-length

```

Co.mat

```

#####
# Sample vampire material file version 5
#####

#-----
# Number of Materials
#-----
material:num-materials=1
#-----
# Material 1 Cobalt Generic
#-----
material[1]:material-name = Co
material[1]:damping-constant = 1.0

```

```
material[1]:exchange-matrix[1] = 11.2e-21
material[1]:atomic-spin-moment = 1.72 !muB
material[1]:uniaxial-anisotropy-constant = 1.0e-24
material[1]:material-element = Ag
material[1]:minimum-height = 0.0
material[1]:maximum-height = 1.0
```

5 Visualization

VAMPIRE provides tools for visualising systems using external programs such as Rasmol, Jmol and POV-Ray. To compile these utilities, use the following command in the main directory of your VAMPIRE installation folder:

```
make vdc
```

The VAMPIRE data converter, or VDC, is run to produce the input files needed. To generate the input files, an output type must be specified in the command line. For example, to produce POV-Ray input files:

```
vdc --povray
```

Note that all command line parameters passed to vdc are not case sensitive.

Getting started

To generate the positions of your atoms, config:atoms must be set in the input file. This will produce .data and .meta files containing the atomic and spin configuration of your system. Their frequency can be adjusted using the config:atoms-output-rate parameter or they can be set to be output at the end of the simulation solely.

In addition, the format of the output can be text or binary format, the latter can help with particularly large systems. Files written in binary format are system specific and usually cannot be read by VDC compiled on separate hardware.

input

```
#-----  
# data output  
#-----  
config:atoms  
config:output-nodes      = 12  
config:atoms-output-rate = 1000  
config:output-format     = binary
```

Atomic visualization with rasmol

To visualise your system using Rasmol, simply run VDC in the same directory as your output with `vdc --xyz`. The `config:atoms` files must be present.

This produced a file called `crystal.xyz`, which is a chemical file format with information on the atomic positions. The format of the `.xyz` format is as follows:

```
.xyz
<number of atoms>
comment line
<element> <X> <Y> <Z>
...
```

The element in the `.xyz` file does not necessarily need to be the same as the atoms used in your system. They can instead be chosen for a different colour palette depending on the users requirements.

Atomic visualization with POV-Ray

To produce pictures of your material of publishable quality and high configurability, it is also possible to use POV-Ray. After running VDC, the file `"spins.pov"` contains all the necessary information and an image may be produced by using:

```
povray spins.pov
```

When running `povray` it is also possible to select specific snapshots or ranges to render using the following flags:

```
povray +KFF[initial frame number] +KFI[final frame number]
```

For example, to render frame 9 only, you could use:

```
povray -W3600 -H2700 +A0.3 +KFI9 +KFF9 spins.pov
```

where the `"-W"` and `"-H"` flags define the width and height of the image (the resolution), and `"+A"` is used for antialiasing.

Output from VDC can be customised in several ways, either by passing parameters to the command line (using the `'--'` notation), or by using a separate VDC input file. The VDC input file is a plain text file containing parameters and arguments which change the output behaviour. This should be helpful when many parameters, or multiple `vdc` runs with the same parameters, are needed. Comments can be included with the `#` symbol and the special characters `' ,(){}[]:=!'` can also

be used but are ignored by VDC. By default this file is called vdc-input, and is read automatically. To change the name of the VDC input file, a command line parameter can be used:

```
vdc --input-file [filename]
```

There are many options that can be used to change all visualisation outputs including Rasmol, Jmol and POV-Ray. To get help with the usage of these parameters outside of the manual, it is also possible to print help messages from the command line by using the -h or --help command line argument, followed by the name of a vdc-input file parameter:

```
vdc -h [parameter-name]
```

The help message should contain information on the parameter, as well as the type of the associated value given, the default value and example usage of the parameter.

General customisation options

It may be beneficial to use only smaller portions of your full system when generating POV-Ray or Rasmol images. This can help with large systems where rendering can become a time constraint, or systems made up of several elements which might be less relevant for the visualisation. There are several similar command line options which can be used to cut up the system in different ways:

slice = float vector(6) [0-1 : default {0,1,0,1,0,1}] The first slice type defines minimum and maximum values for each axis. Only atoms and spins inside these boundaries are included in the visualisation. The parameters passed to this argument are interpreted as fractional coordinates.

slice-void = float vector(6) [0-1 : default not set] This parameter will remove all atoms and spins inside the given borders. This can be used to create cubic hollow systems where only surface atoms are shown, removing a very high percentage of atoms in the system, which can greatly reduce rendering time for both POV-Ray and Rasmol.

slice-sphere = float vector(3) [0-1 : xfrac,yfrac,zfrac] The sphere slice is also used to remove the atoms and spins at the centre of a system. This particular parameter lends itself well to spherical systems as it removes a spherical section of atoms. Three parameters are required, instead of six. Each one defines a region, centred on the centre of the original system, along the respective axis,

equal to a fraction of the system size along that axis. As these parameters are not necessarily equal to each other, this can be used to create an ellipse of missing atoms at the centre of the system.

slice-cylinder = float vector(4) [0-1 : xfrac,yfrac,zmin,zmax] This slice parameter can be used to remove all atoms outside a cylindrical section by defining the x,y-fractional sizes as well as a fractional minimum and maximum along the z-axis.

remove-materials = int [one or more values] In some cases whole materials are not relevant for visualisation purposes and can be altogether removed. To use this command line parameter a list of material indices need to be provided. Material indices start from 1.

POV-Ray Customisation options

The following section contains a list of parameters that only affect POV-Ray output configurations. If another output type is requested, these parameters are ignored.

frame-start = int [default 0] Depending on output options used in VAMPIRE, multiple frames may be rendered by VDC. frame-start can be used to skip an initial number of frames.

frame-final = int [default 0] Depending on output options used in VAMPIRE, multiple frames may be rendered by VDC. frame-final can be used to skip later frames.

camera-position = float vector(3) [(-1,1) : default not set] POV-Ray camera position, set using fractional coordinates. Camera distance from look at point is calculated automatically however it can be changed by using camera-zoom.

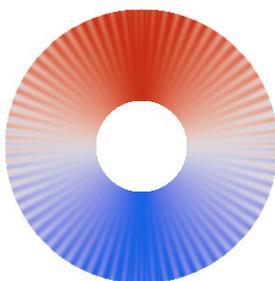
camera-look-at = float vector(3) [(-1,1) : default not set] POV-Ray camera look at position, set using fractional coordinates. The position is a location in the bounding box of the system, with centre (0,0,0).

camera-zoom = float vector(3) [0-∞ : default not set] The default distance from the camera is automatically calculated according to the size of the system. This can be increased or reduced using camera-zoom to multiply the default distance. Values less than 1.0 reduce the distance while values above 1.0 increase it.

background-colour = string [default Gray30] POV-Ray includes various predefined colours such as White, Black, Gray. Misspelled colour names will not be detected by vdc but will cause error in POV-Ray.

atom-sizes = float [one or more : default 1.2] POV-Ray atom sizes. Atoms are represented by spheres with a defined radius. Individual materials can have different atoms sizes by including a list of floats, starting from material 1.

arrow-sizes = float [one or more : default 2.0] POV-Ray arrow sizes. Individual materials can have different arrow sizes by including a list of floats, starting from material 1.



CBWR colourmap

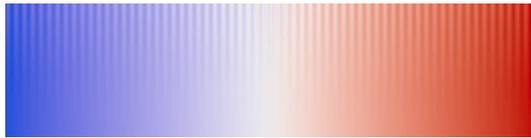
colourmap = string [default CBWR] By default, a 1D colourmap is used. Aligned along the z-axis, spins in the $\{0,0,1\}$ direction are red, while spins antiparallel to this $\{0,0,-1\}$ are blue. Between these values, the colour transitions to white around the xy-plane. This corresponds to the CBWR colourmap, a cyclic blue-white-red map, which lends itself well to 1D or 2D spin systems where there are two principle spin directions, such as antiferromagnets and ferrimagnets. Some care must be taken to align the principle spin directions with the z-axis, as this is the axis along which colour is applied. This can also be changed using the vector-z input parameter. There are several choices of possible colourmap configurations, the ones provided by default are made to be perceptually uniform and in some cases take account of colourblindness. Information on the colourmaps, the importance of perceptually uniform maps and how to adapt and use different maps can be found from "Peter Kovesi. Good Colour Maps: How to Design Them. 2015".

The C2 colourmap is also cyclic and useful for 3D magnetic systems such as vortex states. It has four principle directions of magenta, yellow, green and blue. As it is cyclic, there will be a smooth transition between colour at all angles, irrespective



C2 colourmap

of what is chosen as the zero degree spin direction. Systems which benefit from this colourmap may also use the 3D parameter which applies a brightness effect along the x-axis.



BWR colourmap

The BWR colourmap is very similar in properties to the CBWR map however it is not cyclic. This means that spins along the positive z-axis will be red with a small positive y-component and blue with a small negative y-component. There will be an immediate flip from bright red to blue as this transition occurs. This can be used to emphasise the transition between spin directions. The transition point can be changed by using vector-z.



Rainbow colourmap

The Rainbow colourmap can be used in 2D systems where spins are aligned in many different directions such as high temperature simulations. While it is still designed to be somewhat perceptually uniform, this is very difficult to do with

rainbow palettes hence its use typically loses detail when compared to other maps, however it is also one of the most vibrant.

custom-colourmap = filename A user defined colourmap can also be used. To apply a different map, a file containing 256 colours in the RGB format must be provided in the same directory that VDC is run. RGB values must be space separated, with no other information such as line numbers. The beginning of an example colourmap is shown below.

Pregenerated perceptually uniform colourmaps of various forms, including those included in vampire by default, can be found in peterkovesi.com/projects/colourmaps/index.html under the Download section.

custom_colourmap_file

```
0.000000 0.000000 0.000000
0.005561 0.005563 0.005563
0.011212 0.011219 0.011217
0.016877 0.016885 0.016883
0.022438 0.022448 0.022445
0.027998 0.028011 0.028008
0.033540 0.033554 0.033551
0.039316 0.039333 0.039329
0.044700 0.044719 0.044714
0.049695 0.049713 0.049709
0.054322 0.054343 0.054338
```

3D = bool [default false] POV-Ray images produced by VDC can have a 3D brightening effect applied. Spins which do not line only in the yz-plane have their brightness adjusted according to their x-axis spin component.

vector-z = float vector(3) [default {0,0,1}] The principle axis, along which colour is applied, is the z-axis. This determines where colours will occur depending on the colourmap being used. By default the CBWR map is used; spins along the positive z-direction are red, those along the negative z-direction are blue, and spins aligned along the xy-plane are white.

In many cases, the overall magnetic moment does not necessarily lie along the z-axis. To remedy this, a new vector-z may be defined. To redefine the z-axis, use the parameter vector-z followed by a direction vector. This does not need to be normalised.

For example, if the user defines vector-z = {1,1,1}, spins along the {1,1,1} direction will be red, {-1,-1,-1} will be blue and those perpendicular to the given axis will be white. Brackets can be omitted.

vector-x = float vector(3) [default {0,0,1}] In some cases, the colourmap may not be symmetric along the default xy-plane, such as the C2 colourmap. Here, spins along positive-y are magenta, while those antiparallel are green. This can be adjusted using a similar command line argument vector-x, however this argument cannot be used without first defining vector-z.

afm = int [one or more values] POV-Ray visualization of antiferromagnets can be difficult due to the contrast of colours of antiparallel spins. To remedy this, it is possible to define materials as antiferromagnetic. These materials will have their colours flipped so that they match neighbouring spins while their spin direction remains antiferromagnetic.

Micromagnetic visualization with POV-Ray

cell2povray

macros

customization

colouring options

Visualization Movies

6 Unit Cell Files

VAMPIRE provides a selection of built-in unit cell crystal structures which can be specified with the `create:crystal-structure` keyword. However, there are many more crystal structures and magnetic materials than the code could possibly support, and so an advanced mode is available where the user can specify any atomic structure and exchange interactions which are then imported into the code and can be used to generate large systems and/or cut into the standard geometric shapes. It is also possible to simulate complex non-periodic structures such as nanoparticles obtained from molecular dynamics simulations with the same approach. The unit cell file is specified in the main input file using the keyword:

```
material:unit-cell-file = file.ucf
```

where "file.ucf" is the filename.

The unit cell file format

The unit cell file is split into two main parts. The first part specifies the unit cell shape, the atoms in the unit cell, and their material associations and categorisation for statistics purposes. The second part specifies all atomic exchange interactions within and between neighbouring unit cells. The general plain text format is as follows:

```
1 # Unit cell size:
2 ucx ucy ucz
3 # Unit cell lattice vectors:
4 ucvxx ucvxy ucvxz
5 ucvyx ucvyy ucvyz
6 ucvzx ucvzy ucvzz
7 # Atoms
8 num_atoms_in_unit_cell number_of_materials
```

```

9 atom_id cx cy cz [mat_id cat_id hcat_id]
10 ...
11 ...
12 # Interactions
13 num_interactions [exchange_type]
14 IID i j dx dy dz | Jij
      | Jx Jy Jz
      | Jxx Jxy Jxz Jyx Jyy Jyz Jzx Jzy Jzz

```

In general this format now allows the specification of any system we want, but clearly complex multi-layered systems require large file sizes. Working through line by line:

- 1 # defines a comment line which is ignored by the parser – so these lines are optional.
- 2 ucx, ucy and ucz are the unit cell size in angstroms.
- 4 – 6 These lines define the shape of the unit cell to be replicated, for cubic cells this is the unit matrix. Note that at present only orthogonal lattice vectors are supported by the code due to complexities relating to parallelisation.
- 8 Define the number of atoms and number of materials in the unit cell. Materials allow grouping of atoms by material, and have the same parameters (ie moment, damping, etc). Material specification affects the way statistics are collected and displayed, and also allows the simple creation of ordered alloys. The list of atoms must immediately follow this line.
- 9 – 10 These lines define the atoms in each unit cell and their parameters:
 - . atom_id Number identifier of atom in unit cell, starts at 0.
 - . cx,cy,cz unit cell coordinates as a fraction of unit cell size
 - . mat_id material id of the atom, integer starting at 0
 - . cat_id category id of the atom, used for calculating properties not categorised by material, eg height or sublattice. Integer starting at 1.
 - . hcat_id Height category id used for calculating properties as a function of height

- 12 Defines the total number of interactions for the unit cell and the expected type of exchange (isotropic, vectorial, tensorial, normalized-isotropic, normalized-vectorial, normalized-tensorial). No lines are allowed between this line and the list of interactions.
- 13 These lines list all the interactions.
- . IID Interaction ID is only used for accounting purposes, starts at 0.
 - . i Atom number of atom in local unit cell
 - . j Atom number of atom in local/remote unit cell
 - . dxuc,dyuc,dzuc relative integer coordinates of unit cell for atom j
 - . Jij, Jxx... Exchange values specified in Joules.

Example: Simple Cubic System

As an example, here is a complete sample file for a simple cubic system with a single material.

```

1 # Unit cell size:
2 3.54 3.54 3.54
3 # Unit cell vectors:
4 1.0 0.0 0.0
5 0.0 1.0 0.0
6 0.0 0.0 1.0
7 # Atoms num_atoms num_materials; id cx cy cz mat cat hcat
8 1 1
9 00.00.00.0000
10 # Interactions n exctype; id i j dx dy dz Jij
11 6 isotropic
12 0
13 1
14 2
15 3
16 4
17 5

```

Here only easy axis anisotropy and isotropic exchange are defined. Since there is

only a single atom in the unit cell, all i - j pairs are 0-0, but over the neighbouring unit cells ± 1 in all directions. This generally leads to a large number of interactions for increasing numbers of atoms in the unit cell, and in future I will write a program to generate some different lattices and interaction lists.

7 Input File Command Reference

The input file can accept a large number of commands, and this chapter gives a comprehensive list of all the options and what they do. Commands are in the form category:keyword=value, where value can be optional depending on the keyword.

System Generation

The following commands control generation of the simulated system, including dimensions, crystal structures etc.

create:full Uses the entire generated system without any truncation or consideration of the create:particle-size parameter. create:full should be used when importing a complete system, such as a complete nanoparticle and where a further definition of the system shape is not required. This is the default if no system truncation is defined.

create:cube Cuts a cuboid particle of size $l_x = l_y = l_z = \text{create:particle-size}$ from the defined crystal lattice.

create:cylinder Cuts a cylindrical particle of diameter create:particle-size from the defined crystal lattice. The height of the cylinder extends to the whole extent of the system size create:system-size-z in the z -direction.

create:ellipsoid Cuts an ellipsoid particle of diameter create:particle-size with fractional diameters of dimensions:particle-shape-factor-x, dimensions:particle-shape-factor-y, dimensions:particle-shape-factor-z from the defined crystal lattice.

create:sphere Cuts a spherical particle of diameter create:particle-size from the defined crystal lattice.

create:truncated-octahedron Cuts a truncated octahedron particle of diameter create:particle-size from the defined crystal lattice.

create:particle Defines the creation of a single particle at the centre of the defined system. If create:particle-size is greater than the system dimensions then

the outer boundary of the particle is truncated by the system dimensions.

create:particle-array Defines the creation of a two-dimensional array of particles on a square lattice. The particles are separated by a distance `create:particle-spacing`. If the system size is insufficient to contain at least a single entire particle of size `create:particle-size` then no atoms will be generated and the program will terminate with an error.

create:voronoi-film Generates a two-dimensional voronoi structure of particles, with a mean grain size of `create:particle-size` and variance `create:voronoi-size-variance` as a fraction of the grain size. If `create:voronoi-size-variance=0` then hexagonal shaped grains are generated. The spacing between the grains (defined by the initial voronoi seed points) is controlled by `create:particle-spacing`. The pseudo-random pattern uses a predefined random seed, and so the generated structure will be the same every time. A different structure can be generated by setting a new random seed using the `create:voronoi-random-seed` parameter. Depending on the desired edge structure, the first row can be shifted using the `create:voronoi-row-offset` flag which changes the start point of the voronoi pattern. The `create:voronoi-rounded-grains` parameter generates a voronoi structure, but then applies a grain rounding algorithm to remove the sharp edges.

create:voronoi-size-variance=[float] Controls the randomness of the voronoi grain structure. The voronoi structure is generated using a hexagonal array of seed points appropriately spaced according to the particle size and particle spacing. The seed points are then displaced in x and y according to a gaussian distribution of width `create:voronoi-size-variance` times the particle size. The variance must be in the range 0.0-1.0. Typical values for a realistic looking grain structure are less than 0.2, and larger values will generally lead to oblique grain shapes and a large size distribution.

create:voronoi-row-offset flag [default false] Offsets the first row of hexagonal points to generate a different pattern, e.g. 2,3,2 grains instead of 3,2,3 grains.

create:voronoi-random-seed = int Sets a different integer random seed for the voronoi seed point generation, and thus produces a different random grain structure.

create:voronoi-rounded-grains flag [default false] Controls the rounding of voronoi grains to generate more realistic grain shapes. The algorithm works by expanding a polygon from the centre of the grain, until the total volume bounded by the edges of the grain is some fraction of the total grain area, defined by `create:voronoi-rounded-grains-area`. This generally leads to the removal of sharp

edges.

create:voronoi-rounded-grains-area = float [0.0-1.0, default 0.9] Defines the fractional grain area where the expanding polygon is constrained, in the range 0.0-1.0. Values less than 1.0 will lead to truncation of the voronoi grain shapes, and very small values will generally lead to circular grains. A typical value is 0.9 for reasonable voronoi variance.

create:particle-centre-offset Shifts the origin of a particle to the centre of the nearest unit cell.

create:crystal-structure = string [sc, fcc, bcc, hcp, heusler, kagome, rocksalt, spinel; default sc] Defines the default crystal lattice to be generated. The code supports the basic metallic crystal types simple cubic (sc), body-centred-cubic (bcc), face-centred-cubic (fcc) and hexagonal close-packed (hcp). The code also supports important magnetic structures such as Heusler alloys (heusler), rock-salt such as NiO (rocksalt) spinels such as magnetite (spinel) and kagome lattices.

create:crystal-sublattice-materials = flag [true, false]; default false When set or defined as true, simple crystals with more than one atom per unit cell (bcc, bcc110, fcc, hpc, and kagome) will allocate each atom in the unit cell to a different material. The material allocation to atomic sites can then be done using the usual material:unit-cell-category flags in the material file, in much the same way as for complex crystals.

create:single-spin flag Overrides all create options and generates a single isolated spin.

create:periodic-boundaries-x flag Creates periodic boundaries along the x -direction.

create:periodic-boundaries-y flag Creates periodic boundaries along the y -direction.

create:periodic-boundaries-z flag Creates periodic boundaries along the z -direction.

create:periodic-boundaries flag Creates periodic boundaries along user-defined directions. If left empty this sets periodic boundaries along the xyz -directions. If set to any combination of x, y, z it sets the commensurate directions. For example create:periodic-boundaries = yz will set the periodic boundary conditions along the y and z directions.

create:select-material-by-height Specifies that materials are preferentially assigned by their height specification.

create:select-material-by-geometry Specifies that materials are preferentially assigned by their geometric specification (eg in core-shell systems).

create:fill-core-shell-particles

create:interfacial-roughness Specifies that a global roughness is applied to the material height specification (eg from a non-flat substrate).

create:material-interfacial-roughness Specifies that a material-specific roughness is applied to the material height specification (eg from differences in local deposition rate).

create:interfacial-roughness-random-seed Specifies the random seed for generating the roughness pattern, where different numbers generate different random patterns. Number should ideally be large and around 2,000,000,000.

create:interfacial-roughness-number-of-seed-points Determines the undulation for the roughness, where more points gives a larger undulation.

create:interfacial-roughness-type Determines whether the roughness is applied as peaks or troughs in the material-specific material heights. Valid options are "peaks" or "troughs".

create:interfacial-roughness-seed-radius

create:interfacial-roughness-seed-radius-variance

create:interfacial-roughness-mean-height

create:interfacial-roughness-maximum-height

create:interfacial-roughness-height-field-resolution

create:alloy-random-seed integer [default 683614233] Sets the random seed for the psuedo random number generator for generating random alloys. Simulations use a predictable sequence of psuedo random numbers to give repeatable results for the same simulation. The seed determines the actual sequence of numbers and is used to generate a different alloy distribution. Note that different numbers of cores will change the structure that is generated.

create:grain-random-seed integer [default 1527349271] Sets the random seed for the psuedo random number generator for generating random grain structures.

create:dilution-random-seed integer [default 465865253] Sets the random seed for the psuedo random number generator for diluting the atoms, leading to a different realization of a dilute material. Note that different numbers of cores will change the structure that is generated.

create:intermixing-random-seed integer [default 100181363] Sets the random seed for the psuedo random number generator for calculating intermixing of materials. A different seed will lead to a different realization of a dilute material. Note that different numbers of cores will change the structure that is generated.

create:spin-initialisation-random-seed = integer [default 123456] Sets the random seed for the psuedo random number generator for initialising spin directions. Note that different numbers of cores will change the spin positions that are generated.

System dimensions

The commands here determine the dimensions of the generated system.

dimensions:unit-cell-size = float [0.1 Å- 10 μ m, default 3.54 Å] Defines the size of the unit cell.

dimensions:unit-cell-size-x Defines the size of the unit cell if asymmetric.

dimensions:unit-cell-size-y Defines the size of the unit cell if asymmetric.

dimensions:unit-cell-size-z Defines the size of the unit cell if asymmetric.

dimensions:system-size Defines the size of the symmetric bulk crystal.

dimensions:system-size-x Defines the total size if the system along the x -axis.

dimensions:system-size-y Defines the total size if the system along the y -axis.

dimensions:system-size-z Defines the total size if the system along the z -axis.

dimensions:particle-size = float Defines the size of particles cut from the bulk crystal.

dimensions:particle-spacing Defines the spacing between particles in particle arrays or voronoi media.

dimensions:particle-shape-factor-x = float [0.001-1, default 1.0] Modifies the default particle shape to create elongated particles. The selected particle shape is modified by changing the effective particle size in the x direction. This property scales the as a fraction of the particle-size along the x -direction.

dimensions:particle-shape-factor-y = float [0.001-1, default 1.0] Modifies the default particle shape to create elongated particles. The selected particle shape

is modified by changing the effective particle size in the y direction. This property scales the as a fraction of the particle-size along the y -direction.

dimensions:particle-shape-factor-z = float [0.001-1, default 1.0] Modifies the default particle shape to create elongated particles. The selected particle shape is modified by changing the effective particle size in the z direction. This property scales the as a fraction of the particle-size along the z -direction.

dimensions:particle-array-offset-x [0-10⁴ Å] Translates the 2-D particle array the chosen distance along the x-direction.

dimensions:particle-array-offset-y Translates the 2-D particle array the chosen distance along the y-direction.

dimensions:double macro-cell-size determines the macro cell size for calculation of the demagnetizing field and output of the magnetic configuration. Finer discretisation leads to more accurate results at the cost of significantly longer run times. The cell size should always be less than the system size, as highly asymmetric cells will lead to significant errors in the demagnetisation field calculation.

Exchange calculation

The following commands control the calculation of built-in exchange interactions for the system.

exchange:interaction-range Determines the cutoff range exchange interactions for built-in crystal structures in terms of the nearest neighbour range. Larger ranges will enable more interactions via an exchange function which can include 2nd-10th nearest neighbour interaction shells or exponential functions. Note that longer ranged interactions are slower to calculate. In shell mode the computed interaction shells are printed in the log file.

exchange:function Determines the type of interaction to be used in the spin Hamiltonian. The default nearest-neighbour option forces nearest neighbour interactions only. The shell option groups neighbours at the same interaction distance into shells which can then be assigned different exchange constants. The exponential option implements an exponential decay that is useful for simulating spin glasses and systems such as NdFeB where there are no well-defined neighbour shells. The material-exponential function is similar to exponential, however it allows different exponential exchange functions to be defined for different

inter-material type (for materials as defined in the unit-cell module) interactions e.g. in NdFeB Nd-Fe interactions can have a different function defined vs Fe-Fe interactions.

exchange:decay-multiplier Determines the value of A to be used in $A \exp -r/B + C$ for exchange:function = exponential.

exchange:decay-length Determines the value of B to be used in $A \exp -r/B + C$ for exchange:function = exponential.

exchange:decay-shift Determines the value of C to be used in $A \exp -r/B + C$ for exchange:function = exponential.

exchange:ucc-exchange-parameters[i][j] This is used in conjunction with exchange:function = material-exchange. i and j represent the unit cell category (material as per unit cell module) of the interacting atoms that the user wishes to set the exponential exchange function for. This variable is set to the three comma separated values: A , exchange:decay-multiplier; B , exchange:decay-length; C , exchange:decay-shift in this order.

exchange:dmi-cutoff-range Determines the cutoff range for i-j-k interactions for the built-in DMI in VAMPIRE.

exchange:ab-initio Interprets exchange constants in the ab-initio sense and applies a factor 2 increase in the strength of the exchange constants.

Anisotropy calculation

The following commands control the calculation of the magnetic anisotropy energy for the system.

anisotropy:surface-anisotropy-threshold = integer [default native] Determines minimal number of neighbours to classify as surface atom. The default value is the number of neighbours specified by the crystal or unit cell file. You can set this as a lower threshold.

anisotropy:surface-anisotropy-nearest-neighbour-range = float [default ∞] Sets the interaction range for the nearest neighbour list used for the surface anisotropy calculation.

anisotropy:enable-bulk-neel-anisotropy = bool [default false] Enables calculation of the Néel pair anisotropy in the bulk, irrespective of the number of

neighbours, enabling the effect of localised spin-orbit interactions. Internally this sets a large threshold, and so specifying `anisotropy:surface-anisotropy-threshold` will override this flag.

anisotropy:neel-anisotropy-exponential-range = float [default 2.5] Enables an exponentially range dependent Néel pair anisotropy so that lattice distortions and strains change the magnetoelastic coupling strength. In the usual form the method only takes into account the symmetry ($L_{ij}(r) = \text{const}$). The value should be set to the typical lattice parameter otherwise the total anisotropy will be significantly higher or lower than expected. The functional form of the range dependence is

$$L_{ij}(r_{ij}) = L_0 \exp\left(-F \frac{r_{ij} - r_0}{r_0}\right) \quad (7.1)$$

where r_{ij} is the pair separation, r_0 is the exponential range, F is the exponential factor and L_0 is the usual Néel anisotropy constant. The functional form assures that at the first neighbour distance the value of the Néel anisotropy constant is the same as would be without the range-dependent form.

anisotropy:neel-anisotropy-exponential-factor = float [default 5.52] Enables an exponentially range dependent Néel pair anisotropy so that lattice distortions and strains change the magnetoelastic coupling strength. In the usual form the method only takes into account the symmetry ($L_{ij}(r) = \text{const}$). The prefactor controls the falloff with increasing range.

Dipole field calculation

The following commands control the calculation of the dipole-dipole field. By default the dipole fields are disabled for performance reasons, but for large systems (> 10 nm) the interactions can become important. The VAMPIRE code implements several different solvers balancing accuracy and performance. The default in V5+ is the tensor method, which approximates the dipole-dipole interactions at the macrocell level but calculating a dipole-dipole tensor which is exact if the magnetic moments in each cell are aligned.

dipole:solver = exclusive string [default tensor] Declares the solver to be used for the dipole calculation. Available options are:

macrocell

tensor

atomistic

HAMR calculation

hamr:laser-FWHM-x = float [default 20.0 nm] Defines the full width at half maximum of the Gaussian temperature profile in x-direction in the program hamr-simulation with default units of Angstrom and a default value of 20 nm.

hamr:laser-FWHM-y = float [default 20.0 nm] Defines the full width at half maximum of the Gaussian temperature profile in y-direction in the program hamr-simulation with default units of Angstrom and a default value of 20 nm.

hamr:head-speed = float [default 30.0 m/s] Defines the speed of the head sweeping over the medium in the program hamr-simulation with default units of Angstrom/second and a default value of 30 m/s.

hamr:head-field-x = float [default 20.0 nm] Defines the full width of the box in x-direction where the magnetic field is applied in the program hamr-simulation with default units of Angstrom and a default value of 20 nm.

hamr:head-field-y = float [default 20.0 nm] Defines the full width of the box in y-direction where the magnetic field is applied in the program hamr-simulation with default units of Angstrom and a default value of 20 nm.

hamr:field-rise-time = float [default 1 ps] Defines the field linear rise time in the program hamr-simulation with default units of seconds and a default value of 1 ps.

hamr:field-fall-time = float [default 1 ps] Defines the field linear fall time in the program hamr-simulation with default units of seconds and a default value of 1 ps.

hamr:NPS = float [default 0.0 nm] Defines the shift between the centre of the temperature pulse and the centre of the box defined by hamr:head-field-x and hamr:head-field-y in the program hamr-simulation with default units of Angstrom and a default value of 0 nm. The parameter can be also parsed via the key hamr:NFT-to-pole-spacing.

hamr:bit-size = float [default 0.0 nm] Defines the size of the bit along x (down-track) in the program hamr-simulation with default units of Angstrom

and a default value of 0 nm. The parameter can be also parsed via the key hamr:bit-length.

hamr:track-size = float [default 0.0 nm] Defines the size of the bit along y (cross-track), i.e. the track size of the bit pattern, in the program hamr-simulation with default units of Angstrom and a default value of 0 nm. The parameter can be also parsed via the key hamr:track-width.

hamr:track-padding = float [default 0.0 nm] Defines the spacing between the edges of the system along y (cross-track) and the written bit pattern in the program hamr-simulation with default units of Angstrom and a default value of 0 nm.

hamr:number-of-bits = int [default 0] Defines the number of bits to be written in total in the program hamr-simulation with default value of 0. If the system is too small for the number of bits requested, the sequence is truncated to adapt it to the system.

hamr:bit-sequence-type = exclusive string [default text] Specifies the format type of bit sequence to be simulated in the program hamr-simulation. Available options are:

- single-tone-predefined
- user-defined

If "single-tone-predefined" is given, a single tone adapted to the system size will be generated and hamr:bit-sequence is ignored.

hamr:bit-sequence = int vector Specifies the bit sequence to be simulated in the program hamr-simulation. Acceptable values are -1 (opposite to field direction), 0 (zero field) and 1 (along field direction) and by default the vector is empty.

Simulation Control

The following commands control the simulation, including the program, maximum temperatures, applied field strength etc.

sim:integrator = exclusive string [default llg-heun] Declares the integrator to be used for the simulation. Available options are:

- llg-heun
- monte-carlo

llg-midpoint

constrained-monte-carlo

hybrid-constrained-monte-carlo

sim:program = exclusive string Defines the simulation program to be used.

sim:program = benchmark Program which integrates the system for 10,000 time steps and exits. Used primarily for quick performance comparisons for different system architectures, processors and during code performance optimisation.

sim:program = time-series Program to perform a single time series typically used for switching calculations, ferromagnetic resonance or to find equilibrium magnetic configurations. The system is usually simulated with constant temperature and applied field. The system is first equilibrated for `sim:equilibration-time-steps` time steps and is then integrated for `sim:time-steps` time steps.

sim:program = hysteresis-loop Program to simulate a dynamic hysteresis loop in user defined field range and precision. The system temperature is fixed and defined by `sim:temperature`. The system is first equilibrated for `sim:equilibration-time-steps` time steps at `sim:maximum-applied-field-strength` applied field. For normal loops `sim:maximum-applied-field-strength` should be a saturating field. After equilibration the system is integrated for `sim:loop-time-steps` at each field point. The field increments from `+sim:maximum-applied-field-strength` to `=sim:maximum-applied-field-strength` in steps of `sim:applied-field-increment`, and data is output after each field step.

sim:program = static-hysteresis-loop Program to perform a hysteresis loop in the same way as a normal hysteresis loop, but instead of a dynamic loop the equilibrium condition is found by minimisation of the torque on the system. For static loops the temperature must be zero otherwise the torque is always finite. At each field increment the system is integrated until either the maximum torque for any one spin is less than the tolerance value (10^{-6} T), or if `sim:loop-time-steps` is reached. Generally static loops are computationally efficient, and so `sim:loop-time-steps` can be large, as many integration steps are only required during switching, i.e. near the coercivity.

sim:program = curie-temperature Simulates a temperature loop to determine the Curie temperature of the system. The temperature of the system is increased stepwise, starting at `sim:minimum-temperature` and ending at `sim:maximum-temperature` in steps of `sim:temperature-increment`. At each temperature the system is first equilibrated for `sim:equilibration-steps` time steps and then a

statistical average is taken over `sim:loop-time-steps`. In general the Monte Carlo integrator is the optimal method for determining the Curie temperature, and typically a few thousand steps is sufficient to equilibrate the system. To determine the Curie temperature it is best to plot the mean magnetization length at each temperature, which can be specified using the `output:mean-magnetisation-length` keyword. Typically the temperature dependent magnetization can be fitted using the function

$$m(T) = \langle \sqrt{\sum_i \mathbf{S}_i} \rangle = \left(1 - \frac{T}{T_C}\right)^\beta \quad (7.2)$$

where T is the temperature, T_C is the Curie temperature, and $\beta \sim 0.34$ is the critical exponent.

sim:program = field-cooling

sim:program = temperature-pulse

sim:program = electrical-pulse Simulates the effect of an electrical pulse through either spin-transfer (STT) or spin-orbit (SOT) torques, or through the spin-transport circuit theory model. The system is first equilibrated at constant temperature with zero voltage. A trapezium shaped electrical pulse is applied, linearly increasing from zero voltage to that defined in the code, held constant, then linearly decreased back to zero. In the case of direct STT and SOT simulations, the effective fields are scaled in direct proportion with the applied voltage. The pulse duration is controlled by the parameter `sim:electrical-pulse-time` with a rise time of `sim:electrical-pulse-rise-time` and fall time of `sim:electrical-pulse-fall-time`. The default pulse time is 1 ns, and default fall and rise times are 0, reproducing a square pulse. The time dependence of the fractional voltage can be printed in the output file with the parameter `output:fractional-electric-field-strength`.

sim:program = cmc-anisotropy Iterates through a series of angles at which the global magnetisation is constrained, allowing individual spins to vary, but preventing the system from reaching a true equilibrium. This allows for the examination of magnetocrystalline anisotropy energy and restoring torques.

sim:program = hamr-simulation Simulates a heat assisted magnetic recording (HAMR) writing process with a head sweeping across the medium at a speed defined by the input parameter `hamr:head-speed`, generating an external magnetic field of maximum magnitude `sim:maximum-applied-field-strength` with rise time `hamr:field-rise-time` and fall time `hamr:field-fall-time` within the

region underneath the head defined by the parameters `hamr:head-field-x` and `hamr:head-field-y`. The head also generates a heat pulse with Gaussian profile in the `xy`-plane and uniform along `z` defined by `FWHM` in `x` and `y` direction `hamr:laser-FWHM-x` and `hamr:laser-FWHM-y` respectively, minimum and maximum values of the Gaussian `sim:minimum-temperature` and `sim:maximum-temperature`, respectively. The desired number of bits to be written and bit sequence are defined via `hamr:number-of-bits`, `hamr:bit-sequence-type` and `hamr:bit-sequence` parameters, whereas `hamr:bit-size/hamr:bit-length` and `hamr:track-size/hamr:track-width` set the bit dimension in down-track and cross-track respectively. The margin between the edge of the system and the written tracks in cross-track is specified via `hamr:track-padding`, while `hamr:NPS/hamr:NFT-to-pole-spacing` set the shift between the centre of application of the external field and temperature pulse.

sim:enable-dipole-fields flag Enables calculation of the demagnetising field.

sim:enable-fmr-field

sim:enable-fast-dipole-fields = Bool [default false] Enables fast calculation of the demag field by pre calculation of the interaction matrix.

sim:dipole-field-update-rate = integer [default 1000] Number of timesteps between recalculation of the demag field. Default value is suitable for slow calculations, fast dynamics will generally require much faster update rates.

sim:time-step The timestep for the evolution of the system, determines how long a simulation will take.

sim:total-time-steps The total number of time steps the program will run for.

sim:loop-time-steps The number of time steps that statistics are taken over, including the mean-magnetisation and material-standard-deviation. This takes place after `sim:equilibration` time steps have passed in simulations such as `program:curie-temperature`.

sim:time-steps-increment

sim:equilibration-time-steps The number of simulation time steps that the system is allowed to equilibrate for at each temperature. Statistics are not taken over this range.

sim:simulation-cycles

sim:maximum-temperature The maximum temperature in a simulation over a temperature series, such as `sim:program = curie-temperature`.

sim:minimum-temperature The minimum temperature in a simulation over a temperature series, such as `sim:program = curie-temperature`.

sim:equilibration-temperature The temperature at which a simulation equilibrates, for example, prior to the temperature pulse in `sim:program = temperature-pulse`.

sim:temperature The temperature of the simulation.

sim:temperature-increment The temperature step size in a simulation over a temperature series, such as `sim:program = curie-temperature`.

sim:cooling-time

sim:laser-pulse-temporal-profile The shape of the laser temperature pulse in time, used in `sim:program = temperature-pulse`.

square

two-temperature

double-pulse-two-temperature

double-pulse-square

sim:laser-pulse-time The length of the laser temperature pulse in time, used in `sim:program = temperature-pulse`.

sim:laser-pulse-power The fluence of the laser temperature pulse, used in `sim:program = temperature-pulse`.

sim:second-laser-pulse-time

sim:second-laser-pulse-power

sim:second-laser-pulse-maximum-temperature

sim:second-laser-pulse-delay-time

sim:two-temperature-heat-sink-coupling

sim:two-temperature-electron-heat-capacity The heat capacity of the electrons in the system, used in `sim:program = temperature-pulse`.

sim:two-temperature-phonon-heat-capacity The heat capacity of the phonons in the system, used in `sim:program = temperature-pulse`.

sim:two-temperature-electron-phonon-coupling Dictates the heat exchange coupling between the electrons and the phonons in the system, used in sim:program = temperature-pulse.

sim:cooling-function Dictates the shape of the cooling curve in sim:program = field-cool simulations. Choose from:

exponential

gaussian

double-gaussian

linear

sim:applied-field-strength The strength of the applied external field acting on the system.

sim:maximum-applied-field-strength The maximum strength of the applied external field acting on the system, in a magnetisation field series simulation such as sim:program = hysteresis-loop. In this simulation, this maximum is the maximum magnitude, and dictates both the maximum and minimum (\pm) magnetisation in the target direction.

sim:equilibration-applied-field-strength The strength of the applied external field the system equilibrates in, in a magnetisation field series simulation such as sim:program = hysteresis-loop.

sim:applied-field-strength-increment The increment in the strength of the applied external field acting on the system, in a magnetisation field series simulation such as sim:program = hysteresis-loop.

sim:applied-field-angle-theta

sim:applied-field-angle-phi

sim:applied-field-unit-vector

sim:demagnetisation-factor = float vector [default (000)] Vector describing the components of the demagnetising factor from a macroscopic sample. By default this is disabled, and specifying a demagnetisation factor adds an effective field, such that the total field is given by:

$$\mathbf{H}_{\text{tot}} = \mathbf{H}_{\text{ext}} + \mathbf{H}_{\text{int}} - \mathbf{M} \cdot \mathbf{N}_d$$

where \mathbf{M} is the magnetisation of the sample and \mathbf{N}_d is the demagnetisation factor of the macroscopic sample. The components of the demagnetisation factor must sum to 1. In general the demagnetisation factor should be used without the dipolar field, as this results in counting the demagnetising effects twice. However, the possibility of using both is not prevented by the code.

sim:integrator-random-seed = integer [default 12345] Sets a seed for the psuedo random number generator. Simulations use a predictable sequence of psuedo random numbers to give repeatable results for the same simulation. The seed determines the actual sequence of numbers and is used to give a different realisation of the same simulation which is useful for determining statistical properties of the system.

sim:constraint-rotation-update

sim:constraint-angle-theta = float (default 0) When a constrained integrator is used in a normal program, this variable controls the angle of the magnetisation of the whole system from the x-axis [degrees]. In constrained simulations (such as cmc anisotropy) this has no effect.

sim:constraint-angle-theta-minimum float (default 0) The minimum angle of theta that the global magnetisation is constrained to in a constrained angle series, used in `sim:program = cmc-anisotropy`.

sim:constraint-angle-theta-maximum The maximum angle of theta that the global magnetisation is constrained to in a constrained angle series, used in `sim:program = cmc-anisotropy`.

sim:constraint-angle-theta-increment = float [0.001-360, default 5] Incremental Change of the angle of global magnetisation from z-direction in constrained simulations. Controls the resolution of the angular sweep.

sim:constraint-angle-phi When a constrained integrator is used in a normal program, this variable controls the angle of the magnetisation of the whole system from the x-axis [degrees]. In constrained simulations (such as cmc anisotropy) this has no effect.

sim:constraint-angle-phi-minimum The minimum angle of phi that the global magnetisation is constrained to in a constrained angle series, used in `sim:program = cmc-anisotropy`.

sim:constraint-angle-phi-maximum The maximum angle of phi that the global magnetisation is constrained to in a constrained angle series, used in `sim:program`

= cmc-anisotropy.

sim:constraint-angle-phi-increment Incremental Change of the angle of global magnetisation from z-direction in constrained simulations. Controls the resolution of the angular sweep.

montecarlo:algorithm Selects the trial move algorithm for use with the Monte Carlo solver. The following options are available:

adaptive (default)

spin-flip

uniform

angle

hinzke-nowak

The adaptive move performs a gaussian move with a tuned trial width to attempt to maintain a 50% acceptance probability, and is the most efficient method in most cases. A spin flip flips the direction of the spin 180° and can be used to perform Ising-type simulations for a uniform starting configuration. Uniform moves a spin to a random location on the unit sphere. Angle performs a gaussian move with a parametric estimate of the optimal width. Hinzke-Nowak performs a random combination of spin-flip, uniform and angle type-moves.

montecarlo:constrain-by-grain Applies a local constraint in granular systems so that the magnetisation within individual grains is conserved along the global constraint directions `sim:constrain-phi` and `sim:constraint-theta`. Without this additional constraint, the system will tend to demagnetise and form a demagnetised state (with zero torque). With this parameter defined it is possible to determine grain-level properties and distributions of the Curie temperature and anisotropy.

sim:checkpoint flag [default false] Enables checkpointing of spin configuration at end of the simulation. The options are:

`sim:save-checkpoint=end`

`sim:save-checkpoint=continuous`

`sim:save-checkpoint-rate=1`

`sim:load-checkpoint=restart`

`sim:load-checkpoint=continue`

sim:preconditioning-steps = integer [default 0] Defines a number of preconditioning steps to thermalise the spins at `sim:equilibration-temperature` prior to the main simulation starting. The preconditioner uses a Monte Carlo algorithm to develop a Boltzmann spin distribution prior to the main program starting. The method works in serial and parallel mode and is especially efficient for materials with low Gilbert damping. The preconditioning steps are applied after loading a checkpoint, allowing you to take a low temperature starting state and thermally equilibrate it.

sim:electrical-pulse-time = float [default 1.0 ns] Defines the pulse time in the program `electrical-pulse` with default units of seconds and a default pulse time of 1 ns.

sim:electrical-pulse-rise-time = float [default 0.0 ns] Defines the pulse linear rise time in the program `electrical-pulse` with default units of seconds and a default pulse time of 0, i.e. an instantaneous turning on of the current.

sim:electrical-pulse-fall-time = float [default 0.0 ns] Defines the pulse linear fall time in the program `electrical-pulse` with default units of seconds and a default pulse time of 0, i.e. an instantaneous turning off of the current.

Data output

The following commands control what data is output to the output file. The order in which they appear is the order in which they appear in the output file. Most options output a single column of data, but some output multiple columns, particularly vector data or parameters related to materials, where one column per material is output. Note that this means that for vector data, one set of columns per material is output.

output:time-steps Outputs the number of time steps (or Monte Carlo steps) completed during the simulation so far.

output:real-time Outputs the simulation time in seconds. The real time is given by the number of time steps multiplied by `sim:time-step` (default value is 1.0×10^{-15} s). The real time has no meaning for Monte Carlo simulations.

output:temperature Outputs the instantaneous system temperature in Kelvin.

output:applied-field-strength Outputs the strength of the applied field in Tesla. For hysteresis simulations the sign of the applied field strength changes along a fixed axis and is represented in the output by a similar change in sign.

output:applied-field-unit-vector Outputs a unit vector in three columns $\hat{h}_x, \hat{h}_y, \hat{h}_z$ indicating the direction of the external applied field.

output:applied-field-alignment Outputs the dot product of the net magnetization direction of the system with the external applied field direction $\hat{\mathbf{m}} \cdot \hat{\mathbf{H}}$.

output:material-applied-field-alignment Outputs the dot product of the net magnetization direction of each material defined in the material file with the external applied field direction $[\hat{\mathbf{m}}_1 \cdot \hat{\mathbf{H}}], [\hat{\mathbf{m}}_2 \cdot \hat{\mathbf{H}}] \dots [\hat{\mathbf{m}}_n \cdot \hat{\mathbf{H}}]$.

output:magnetisation Outputs the instantaneous magnetization of the system. The data is output in four columns $\hat{m}_x, \hat{m}_y, \hat{m}_z, |m|$ giving the unit vector direction of the magnetization and normalized length of the magnetization respectively. The normalized length of the magnetization $|m| = |\sum_i \mu_i S_i| / \sum \mu_i$ is given by the sum of all moments in the system assuming ferromagnetic alignment of all spins. Note that the localized spin moments μ_i are taken into account in the summation.

output:magnetisation-length Outputs the instantaneous normalized magnetization length $|m| = |\sum_i \mu_i S_i| / \sum \mu_i$, where the saturation value is defined by ferromagnetic alignment of all spins in the system. Note that the localized spin moments μ_i are taken into account in the summation.

output:mean-magnetisation-length Outputs the time-averaged normalized magnetization length $\langle |m| \rangle$.

output:mean-magnetisation Outputs the time-averaged normalized magnetization vector $\langle |\mathbf{m}| \rangle$.

output:material-magnetisation Outputs the instantaneous normalized magnetization for each material in the simulation. The data is output in blocks of four columns, with one block per material defined in the material file, e.g.

$$[\hat{m}_1^x, \hat{m}_1^y, \hat{m}_1^z, |m_1|], [\hat{m}_2^x, \hat{m}_2^y, \hat{m}_2^z, |m_2|] \dots [\hat{m}_n^x, \hat{m}_n^y, \hat{m}_n^z, |m_n|]$$

Note that obtaining the actual macroscopic magnetization length from this data is not trivial, since it is necessary to know how many atoms of each material are in the system. This information is contained within the log file (giving the fraction of atoms which make up each material). However it is usual to also output the total normalized magnetization of the system to give the relative ordering of the entire system.

output:material-mean-magnetisation-length Outputs the time-averaged normalized magnetization length for each material, e.g. $\langle |m_1| \rangle, \langle |m_2| \rangle \dots \langle |m_n| \rangle$.

output:material-mean-magnetisation Outputs the time-averaged normalized magnetization length for each material, e.g. $\langle |\mathbf{m}_1| \rangle, \langle |\mathbf{m}_2| \rangle \dots \langle |\mathbf{m}_n| \rangle$.

output:total-torque Outputs the instantaneous components of the torque on the system $\tau = \sum_i \mu_i \mathbf{S}_i \times \mathbf{H}_i$ in three columns τ_x, τ_y, τ_z (units of Joules). In equilibrium the total torque will be close to zero, but is useful for testing convergence to an equilibrium state for zero temperature simulations.

output:mean-total-torque Outputs the time average of components of the torque on the system $\langle \tau \rangle = \langle \sum_i \mu_i \mathbf{S}_i \times \mathbf{H}_i \rangle$ in three columns $\langle \tau_x \rangle, \langle \tau_y \rangle, \langle \tau_z \rangle$. In equilibrium the total torque will be close to zero, but the average torque is useful for extracting effective anisotropies or exchange using constrained Monte Carlo simulations.

output:constraint-phi Outputs the current angle of constraint from the z -axis for constrained simulations using either the Lagrangian Multiplier Method (LMM) or Constrained Monte Carlo (CMC) integration methods.

output:constraint-theta Outputs the current angle of constraint from the x -axis for constrained simulations using either the Lagrangian Multiplier Method (LMM) or Constrained Monte Carlo (CMC) integration methods.

output:material-mean-torque Outputs the time average of components of the torque on the each material system $\langle \tau \rangle$ in blocks of three columns, with one block for each material defined in the material file e.g.

$$[\langle \tau_1^x \rangle, \langle \tau_1^y \rangle, \langle \tau_1^z \rangle], [\langle \tau_2^x \rangle, \langle \tau_2^y \rangle, \langle \tau_2^z \rangle] \dots [\langle \tau_n^x \rangle, \langle \tau_n^y \rangle, \langle \tau_n^z \rangle]$$

Computing the torque on each material is particularly useful for determining equilibrium properties of multi-component systems with constrained Monte Carlo simulations. In certain cases the components of a system (different materials) can exert equal and opposite torques on each other, giving a total system torque of zero. The decomposition of the torques for each material allows the determination of internal torques in the system.

output:mean-susceptibility Outputs the components of the magnetic susceptibility χ . The magnetic susceptibility is defined by

$$\chi_\alpha = \frac{\sum_i \mu_i}{k_B T} (\langle m_\alpha^2 \rangle - \langle m_\alpha \rangle^2)$$

where $\alpha = x, y, z, m$ giving the directional components of the magnetization in x, y and z respectively as well as the longitudinal susceptibility χ_m . The data is

output in four columns χ_x , χ_y , χ_z , and χ_m in units of Tesla^{-1} . The susceptibility is useful for identifying the critical temperature for a system as well as atomistic parameterisation of the micromagnetic Landau-Lifshitz-Bloch (LLB) equation.

output:material-mean-susceptibility Outputs the components of the magnetic susceptibility χ for each defined material in the system. The data is output in sets of four columns χ_x , χ_y , χ_z , and χ_m for each material. In multi-sublattice systems the susceptibility of each sublattice can be different.

output:material-standard-deviation Outputs the standard deviation in the components of the instantaneous normalized magnetization for each material in the simulation. The data is output in blocks of four columns, with one block per material defined in the material file, e.g.

$$\left[\hat{\sigma}_1^x, \hat{\sigma}_1^y, \hat{\sigma}_1^z, \sigma_{|m_1|} \right], \left[\hat{\sigma}_2^x, \hat{\sigma}_2^y, \hat{\sigma}_2^z, \sigma_{|m_2|} \right] \dots \left[\hat{\sigma}_n^x, \hat{\sigma}_n^y, \hat{\sigma}_n^z, \sigma_{|m_n|} \right]$$

The statistic is taken over the range of values gathered during the loop-time-steps after equilibration.

output:electron-temperature Outputs the instantaneous electron temperature as calculated from the two temperature model.

output:phonon-temperature Outputs the instantaneous phonon (lattice) temperature as calculated from the two temperature model.

output:total-energy

output:mean-total-energy

output:anisotropy-energy

output:mean-anisotropy-energy

output:exchange-energy

output:mean-exchange-energy

output:applied-field-energy

output:mean-applied-field-energy

output:magnetostatic-energy

output:mean-magnetostatic-energy

output:material-total-energy Outputs the total energy of each material in the system.

output:material-mean-total-energy Outputs the mean total energy of each material in the system.

output:mean-specific-heat Outputs the mean total specific heat C_v defined by:

$$C_v = \frac{(\langle U^2 \rangle - \langle U \rangle^2)}{k_B T^2}$$

where U is the internal (total) energy, T is the system temperature and k_B is the Boltzmann constant. The outputted value has units of k_B per spin. For classical magnets the specific heat tends to a non-zero constant approaching zero temperature. When using spin temperature rescaling the specific heat tends to zero as expected for a quantum system.

output:material-mean-specific-heat Outputs the mean specific heat for each defined material in the system in units of k_B per spin. The data is formatted as one column per material.

output:fractional-electric-field-strength Outputs the fractional electric field strength (or voltage) during an electrical-pulse simulation.

output:mpi-timings

output:gnuplot-array-format

output:output-rate = integer [default 1] Controls the number of data points written to the output file or printed to screen. By default VAMPIRE calculates statistics once every `sim:time-steps-increment` number of time steps. Usually you want to output the updated statistic (e.g. magnetization) every time, which is the default behaviour. However, sometimes you may want to plot the time evolution of an average, where you want to collect statistics much more frequently than you output to the output file, which is controlled by this keyword. For example, if `output:output-rate = 10` and `sim:time-steps-increment = 10` then statistics (and average values) will be updated once every 10 time steps, and the new statistics will be written to the output file every 100 time steps.

output:precision = integer [default 6] Controls the number of digits to be used for data written to the output file or printed to screen. The default value is 6 digits of precision.

output:fixed-width = flag [default false] Controls the formatting to be used for data written to the output file or printed to screen. The default is false which ignores trailing zeros in the output.

output:column-headers= flag [default false] Controls the headers at the top of output columns in the output file. The default is false which writes no headers.

Configuration output

These options enable the output of spin configuration snapshots during the simulation. The configurations can then be visualised using povray or other software generated with the vampire data converter (vdc) utility.

config:atoms flag [default false] Enables the output of atomic spin configurations either at the end of the simulations or during the simulation. The options are:

config:atoms - to output continuously during the simulation

config:atoms=continuous - to output continuously during the simulation (same as previous option)

config:atoms=end - to output at the end of the simulation

config:atoms-output-rate = int [0+, default 1000] Determines the rate configuration files are outputted as a multiple of sim:time-steps-increment. It is considered only if config:atoms=continuous or is empty.

The following options allow a cubic slice of the total configuration data to be output to the configuration file. This is useful for reducing disk usage and processing times, especially for large scale simulations.

config:atoms-minimum-x = float [0.0 - 1.0] Determines the minimum x value (as a fraction of the total system dimensions) of the data slice to be outputted to the configuration file.

config:atoms-minimum-y Determines the minimum y value (as a fraction of the total system dimensions) of the data slice to be outputted to the configuration file.

config:atoms-minimum-z Determines the minimum z value (as a fraction of the total system dimensions) of the data slice to be outputted to the configuration file.

config:atoms-maximum-x Determines the maximum x value (as a fraction of the total system dimensions) of the data slice to be outputted to the configuration file.

config:atoms-maximum-y Determines the maximum y value (as a fraction of the total system dimensions) of the data slice to be outputted to the configuration file.

config:atoms-maximum-z Determines the maximum z value (as a fraction of the total system dimensions) of the data slice to be outputted to the configuration file.

config:macro-cells flag [default false] Enables the output of macro cell spin configurations either at the end of the simulations or during the simulation. The options are:

config:macro-cells - to output continuously during the simulation

config:macro-cells=continuous - to output continuously during the simulation (same as previous option)

config:macro-cells=end - to output at the end of the simulation

config:macro-cells-output-rate Determines the rate configuration files are outputted as a multiple of sim:time-steps-increment. It is considered only if config:macro-cells = continuous or is empty

config:output-format = exclusive string [default text] Specifies the format of the configuration data. Available options are:

text

binary

The text option outputs data files as plain text, allowing them to be read by a wide range of applications and hence the highest portability. There is a performance cost to using text mode and so this is recommended only if you need portable data and will not be using the vampire data converter (vdc) utility. The binary option outputs the data in binary format and is typically 100 times faster than text mode. This is important for large-scale simulations on large numbers of processors where the data output can take a significant amount of time. Binary files are generally not compatible between operating systems and so the vdc tools generally needs to be run on the same system which generated the files.

config:output-mode = exclusive string [default file-per-node] Specifies how configuration data is outputted to disk. Available options are:

file-per-node

file-per-process

mpi-io

Using this option is important for obtaining good performance on Tier-0 (European) and Tier-1 (National) supercomputers for simulations typically using more than 1000 cores. Large scale supercomputers have high performance parallel file systems with a peak bandwidth typically over 10 GB/s. VAMPIRE supports three different modes of data output: file-per-node (FPN), file-per-process (FPP) and mpi-io. Note that high performance requires `config:output-format = binary` to be set to output the data in binary format, but this is not default behaviour for easier data analysis and portability for the casual user.

The first (default) option of file-per-node collates data from different processes onto a defined number of `config:output-nodes` before outputting to disk, with the total data spread out with a different file per output node. This has good performance for medium-scale simulations and above (>100 cores) with a reasonable number of output nodes (typically 1 per physical node). Small simulations (<100 cores) benefit from a larger number of output processes. to maximise bandwidth fro the independent write operations. For typical simulations with > 40,000 atoms per core striping of the parallel file system improves performance, while for less atoms striping can be detrimental and should be disabled. This option is also best for distributed file systems, typical on local resources such as university clusters.

The file-per-process option means every process in the simulation outputs its own data to disk independent of all others. The option is available for advanced tuning but is generally not recommended for typical simulations due to the large number of small files generated, complicating data analysis and generally having very poor performance.

The mpi-io option uses the MPI library routines to output a large, single shared file. Enabling this option automatically forces binary data output. In general this option gives good performance for large systems with a single file for each configuration snapshot. In general this gives worse performance than the file-per-node option except for the largest system sizes.

config:output-nodes = int [default 1] Specifies the number of files to be generated per snapshot. For typical small scale simulations (on a single physical node) the default value of 1 is fine. For larger scale simulations more output nodes are beneficial to achieve maximum performance, with one output node per physical node being a sensible choice, but this can be specified up to the maximum number of processes in the simulation.

8 Material File Command Reference

The material file defines all the magnetic properties of the materials used in the simulation, including exchange, anisotropy, damping etc. Material properties are defined by an index number for each material, starting at one. Material properties are then defined as follows:

```
material[index]:keyword = value !unit
```

followed by a carriage return, so that each property is defined on a separate line. The defined keywords are listed below. The material file is largely free-format, apart from the first line which must specify the number of materials for the simulation. The material properties can be defined in any order, and if omitted the default value will be used. When the same property for a particular material is defined in the file, the last definition (reading top to bottom) will be used. Comments can be added to the file using the # character, which moves the file parser to the next line.

Material File Parameters

material:num-materials = int [1-100; default 1] Defines the number of materials to be used in the simulation, and must be the first uncommented line in the file. If more than n materials are defined, then only the first n materials are actually used. The maximum number of different materials is currently limited to 100. If using a custom unit cell then the number of materials in the unit cell must match the number of materials here, otherwise the code will produce an error.

material:material-name = string [default material#n] Defines an identifying name for the material with a maximum length of xx characters. The identifying name is only used in the output files and does not affect the running of the code.

material:damping-constant = float [0.0-10.0; default 1.0] Defines the phenomenological relaxation rate (damping) in dynamic simulations using the LLG

equation. For equilibrium properties the damping should be set to 1 (critical damping), while for realistic dynamics the damping should be representative of the material. Typical values range from 0.005 to 0.1 for most materials.

material:exchange-matrix[index] = float [default 0.0 J/link] Defines the pairwise exchange energy between atoms of type index and neighbour-index. The pairwise exchange energy is independent of the coordination number, and so the total exchange integral will depend on the number of nearest neighbours for the crystal lattice. The exchange energy must be defined between all material pairs in the simulation, with positive values representing ferromagnetic coupling, and negative values representing anti ferromagnetic coupling. For a ferromagnet with nearest neighbour exchange, the pairwise exchange energy can be found from the Curie temperature by the meanfield expression:

$$J_{ij} = \frac{3k_B T_C}{\epsilon z}$$

where J_{ij} is the exchange energy, k_B is the Boltzmann constant, T_C is the Curie temperature, z is the coordination number (number of nearest neighbours) and ϵ is a correction factor to account for spin wave fluctuations in different crystal lattices. If a custom unit cell file with non-normalised exchange interactions is used the exchange values defined here are ignored.

material:exchange-matrix-1st-nn[index] = float [default 0.0 J/link] Defines the pairwise exchange energy between atoms of type index and neighbour-index for the first nearest neighbour shell when using the built in exchange functions. This is exactly the same as the usual parameter exchange-matrix[index] but with a more specific syntax including the shell number of 1.

material:exchange-matrix-2nd-nn[index] = float [default 0.0 J/link] Defines the pairwise exchange energy between atoms of type index and neighbour-index for the second nearest neighbour shell when using the built in exchange functions. If you are using the generic crystal structures available in VAMPIRE, then it is possible to define a longer ranged Hamiltonian with next-nearest up to 10th nearest neighbour interactions. The interaction shells refer to sets of neighbours with the same interaction range from a target atom. To define a longer range Hamiltonian you need to define the input file parameters textitexchange:interaction-range = R and exchange:function = shell, where R is the interaction range as a multiple of the nearest neighbour distance. When these options are defined, the number of computed shells is printed in the log file along

with their distance and coordination number. For the common crystal structures the shell coordinations are well known. Note that longer ranged Hamiltonians will naturally be slower due to the larger number of computed exchange interactions. For more than ten shells or asymmetric crystals with different interactions along different crystal directions the unit cell file is still required.

material:exchange-matrix-3rd-nn[index] = float [default 0.0 J/link] Defines the pairwise exchange energy between atoms of type index and neighbour-index for the third nearest neighbour shell when using the built in exchange functions. See 2nd neighbour description for more details on using this feature.

material:exchange-matrix-(4th-10th)-nn[index] = float [default 0.0 J/link] Defines the pairwise exchange energy between atoms of type index and neighbour-index for the fourth nearest neighbour shell when using the built in exchange functions. See 2nd neighbour description for more details on using this feature.

material:biquadratic-exchange-matrix[index] = float [default 0.0 J/link] Defines the pairwise biquadratic exchange energy between atoms of type index and neighbour-index. The pairwise exchange energy is independent of the coordination number, and so the total exchange integral will depend on the number of neighbours for the crystal lattice. The exchange energy must be defined between all material pairs in the simulation, with positive values representing ferromagnetic coupling, and negative values representing anti ferromagnetic coupling. The biquadratic exchange Hamiltonian is given by the expression

$$E^{\text{BQ}} = \sum_{i < j} J_{ij}^{\text{BQ}} (\mathbf{S}_i \cdot \mathbf{S}_j)^2$$

where J_{ij}^{BQ} is the biquadratic exchange energy.

material:atomic-spin-moment = float [0.01+ μ_B , default 1.72 μ_B] Defines the local effective spin moment for each atomic site. Atomic moments can be found from ab-initio calculations or derived from low temperature measurements of the saturation magnetisation. The atomic spin moments are related to the macroscopic magnetisation by the expression:

$$\mu_S = \frac{M_S a^3}{n}$$

where a is the lattice constant, n is the number of atoms per unit cell, and

M_S is the saturation magnetisation in units of $J/T/m^3$ (A/m). Note that unlike micromagnetic simulations, atomistic simulations always use zero-K values of the spin moments, since thermal fluctuations of the magnetisation are provided by the model. Small values ($< 1\mu_B$) will typically lead to integration problems for the LLG unless sub-femtosecond time steps are used.

material:uniaxial-anisotropy-constant = float [default 0.0 J/atom] Defines the local second order single-ion magnetocrystalline anisotropy constant at each atomic site. The anisotropy energy is given by the expression

$$E_i = -k_2(\mathbf{S}_i \cdot \mathbf{e}_i)^2$$

where \mathbf{S}_i is the local spin direction and \mathbf{e}_i is the easy axis unit vector. Positive values of k_2 give a preferred easy axis orientation, and negative values give a preferred easy plane orientation of the spin perpendicular to the easy axis direction.

material:second-order-uniaxial-anisotropy-constant = float [default 0.0 J/atom] Has the same meaning and is the preferred form for material:uniaxial-anisotropy-constant.

material:fourth-order-uniaxial-anisotropy-constant = float [default 0.0 J/atom] Implements fourth order uniaxial anisotropy as implemented with spherical harmonics.

material:cubic-anisotropy-constant = float [default 0.0 J/atom] Defines the local cubic magnetocrystalline anisotropy constant at each atomic site. The anisotropy energy is given by the expression

$$E_i = +\frac{k_c}{2}(S_x^4 + S_y^4 + S_z^4)$$

where $S_{x,y,z}$ are the components of the local spin direction and k_c is the cubic anisotropy constant. Positive values of k_c give a preferred easy axis orientation along the [001] directions, medium-hard along the [110] directions and hard along the [111] directions. Negative values give a preferred easy direction along the [111] directions, medium hard along the [110] directions and hard along the [100] directions.

material:fourth-order-cubic-anisotropy-constant = float [default 0.0 J/atom] Has the same meaning and preferred form for material:cubic-anisotropy-constant.

material:uniaxial-anisotropy-direction = float vector [default (0,0,1)] A unit

vector \mathbf{e}_i describing the magnetic easy axis direction for uniaxial anisotropy. The vector is entered in comma delimited form - For example:

```
material[1]:uniaxial-anisotropy-direction = 0,0,1
```

The unit vector is self normalising and so the direction can be expressed in standard form (with length $r = 1$) or in terms of crystallographic directions, e.g. [111].

Random anisotropy can be specified using

```
material[1]:uniaxial-anisotropy-direction = random
```

which allocates each atom in the material a random anisotropy vector in three dimensions. This may be applicable to truly amorphous Rare earth alloys where no local anisotropy direction is preferred.

Random anisotropy for each grain in the simulation (specified by generating a voronoi structure or particle array) can be specified using

```
material[1]:uniaxial-anisotropy-direction = random-grain
```

which allocates each atom in the material a random anisotropy vector in three dimensions, giving all atoms of the material in each grain a different anisotropy vector. Both random anisotropies are compatible with higher order uniaxial anisotropy but not lattice anisotropy or cubic anisotropy.

material:surface-anisotropy-constant = float default 0.0 (J/atom) Describes the surface anisotropy constant in the Néel pair anisotropy model. The anisotropy is given by a summation over nearest neighbour atoms given by

$$E_i = +\frac{1}{2} \sum_j^{\text{nn}} k_s (\mathbf{S} \cdot \mathbf{r}_{ij})^2$$

where k_s is the surface anisotropy constant between atoms i and j and \mathbf{r}_{ij} is a unit vector between sites i and j .

neel-anisotropy-constant[index] = float [default 0.0 J] Has the same meaning and is the preferred form for material:surface-anisotropy-constant.

lattice-anisotropy-constant = float [default 0.0 J/atom] Defines anisotropy arising from temperature dependent lattice expansion such as in RETM alloys. The temperature dependence of the lattice anisotropy is defined with a user defined function specified in the parameterlattice-anisotropy-file.

lattice-anisotropy-file = string Defines a file containing the temperature dependence of anisotropy arising from lattice expansion. The first line of the file specifies the number of points, followed by a list of pairs indicating the temperature (in K) and normalised lattice anisotropy, typically of order 1 at $T = 0K$. The specified points are linearly interpolated by the code and so excessively high resolution is not required for high accuracy, and 1 K resolution is typically sufficient for most problems.

voltage-controlled-magnetic-anisotropy-coefficient = float [default 0.0 J/V]
Defines the material-dependent magnetic anisotropy induced by an applied voltage. Here the easy axis is always assumed along z , and so a positive coefficient will add uniaxial anisotropy along the z -direction. The coefficient is multiplied by spin-transport:applied-voltage and so has units of J/V (or Coulombs), with typical values in the range $0.1 - 10 \times 10^{-21}$ J / volt.

material:relative-gamma float [default 1] Defines the gyromagnetic ratio of the material relative to that of the electron $\gamma_e = 1.76 \text{ T}^{-1}\text{s}^{-1}$. Valid values are in the range 0.01 - 100.0. For most materials $\gamma_r = 1$.

material:initial-spin-direction float vector /bool [default (001) / false] Determines the initial direction of the spins in the material. Value can either be a unit vector defining a direction in space, or a boolean which initialises each spin to a different random direction (equivalent to infinite temperature). As with other unit vectors, a normalised value or crystallographic notation (e.g. [110]) may be used.

material:material-element string [default "Fe"] Defines a purely descriptive chemical element for the material, which gives visual contrast in a range of interactive atomic structure viewers such as jmol, rasmol etc. In rasmol, Fe is a gold colour, H is white, Li is a deep red, O is red, B is green and Ag is a medium grey. This parameter has no relevance to the simulation at all, and only appears when outputting atomic coordinates, which can be post-processed to be viewable in rasmol. The contrast is particularly useful in inspecting the generated structures, particularly ones with a high degree of complexity.

material:geometry-file string [default ""] Specifies a filename containing a series of connected points in space which is used to cut a specified shape from the material, in a process similar to lithography. The first line defines the total number of points, which must be in the range 3-100 (A minimum of three points is required to define a polygon). The points are normalised to the sample size, and so all points are defined as x, y pairs in the range 0-1, with one point per line. The last point is automatically connected first, so need not be defined twice.

material:alloy-host flag [default off] Scans over all other materials to replace the desired fraction of host atoms with alloy atoms. This is primarily used to create random alloys of materials with different properties (such as FeCo, NiFe) or disordered ferrimagnets (such as GdFeCo).

material:alloy-fraction[index] = float [0-1 : default 0.0] Defines the fractional number of atoms of the host material to be replaced by atoms of material *index*.

material:minimum-height = float [0-1 : default 0.0] Defines the minimum height of the material as a fraction of the total height z of the system. By defining different minimum and maximum heights it is easy to define a multilayer system consisting of different materials, such as FM/AFM, or ECC recording media.

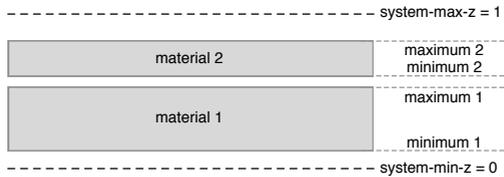


Figure 8.1: Schematic diagram showing definition of a multilayer system consisting of two materials. The minimum-height and maximum-height are defined as a fraction of the total z -height of the system.

The heights of the material are applied when the crystal is generated, and so in general further geometry changes can also be applied, for example cutting a cylinder shape or voronoi granular media, while preserving the multilayer structure. The code will also print a warning if materials overlap their minimum/maximum ranges, since such behaviour is usually (but not always) undesirable.

material:maximum-height = float [0-1 : default 1.0] Defines the maximum height of the material as a fraction of the total height z of the system. See `material:minimum-height` for more details.

material:core-shell-size = float [0-1 : default 1.0] Defines the radial extent of a material as a fraction of the particle radius. This parameter is used to generate core-shell nanoparticles consisting of two or more distinct layers.

The core-shell-size is compatible with spherical, ellipsoidal, cylindrical, truncated octahedral and cuboid shaped particles. In addition when particle arrays are generated all particles are also core-shell type. This option is also comparable with the minimum/maximum-height options, allowing for partially filled or coated

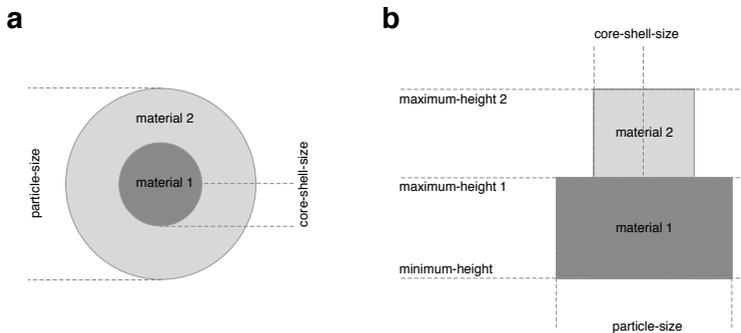


Figure 8.2: (a) Schematic diagram showing definition of a nanoparticle with two materials with different radii. core-shell-size is defined as a fraction of the particle radius (particle-size/2). (b) Schematic diagram showing side-on view of a cylinder, consisting of two materials with different core-shell-size and different maximum heights. Part of the core material is exposed, while the other part is covered with the other material.

nanoparticles.

material:interface-roughness = float [0-1 : default 1.0] Defines interfacial roughness in multilayer systems.

material:intermixing[index] = float [0-1 : default 1.0] Defines intermixing between adjacent materials in multilayer systems. The intermixing is defined as a fraction of the total system height, and so small values are usually used. The intermixing defines the mixing of material *index* into the host material, and can be asymmetric ($a \rightarrow b \neq b \rightarrow a$).

material:density = float [0-1 : default 1.0] Defines the fraction of atoms to remove randomly from the material (density).

material:continuous = flag [default off] Defines materials which ignore granular CSG operations, such as particles, voronoi media and particle arrays.

material:fill-space = flag [default off] Defines materials which obey granular CSG operations, such as particles, voronoi media and particle arrays, but in-fill the void created. This is useful for embedded nanoparticles and recording media with dilute interlayer coupling.

material:couple-to-phononic-temperature = flag [default off] Couples the spin system of the material to the phonon temperature instead of the electron temperature in pulsed heating simulations utilising the two temperature model.

Typically used for rare-earth elements.

material:temperature-rescaling-exponent = float [0-10 : default 1.0] Defines the exponent when rescaled temperature calculations are used. The higher the exponent the flatter the magnetisation is at low temperature. This parameter must be used with `temperature-rescaling-curie-temperature` to have any effect.

material:temperature-rescaling-curie-temperature = float [0-10,000 : default 0.0] Defines the Curie temperature of the material to which temperature rescaling is applied.

material:non-magnetic flag [default remove] Defines atoms of that material as being non-magnetic. Non-magnetic atoms by default are removed from the simulation and play no role in the simulation. If configuration output is specified then the positions of the non-magnetic atoms are saved and processed by the `vdc` utility. This preserves the existence of non-magnetic atoms when generating visualisations but without needing to simulate them artificially. The "keep" option preserves the non-magnetic atoms in the simulation for parallelization efficiency but instructs the dipole field solver to ignore them for improved accuracy.

material:unit-cell-category = int [default 0] Allocates different materials to different atoms in the unit cell. In complex crystals such as spinel and rocksalt, the material allocations of different atoms in the unit cell are defined by the structure. For example, in the rocksalt structure there are two distinct types of atoms. Material 1 could be allocated to this site using `material[1]:unit-cell-category = 1`, and a second material could be allocated to the second site using `material[2]:unit-cell-category = 2`. The default value of this variable is 0, and so for complex crystals only a single site is defined and the other sites will not be generated unless a material is attached to it in this way.

This keyword also works with the `create:crystal-sublattice-materials` flag to allocate different materials to different sites in the simple crystals `bcc`, `fcc`, `hcp` and `kagome`. This feature is especially useful for simulating simple antiferromagnets and materials with different kinds of defects or site specific alloying.

Bibliography