

# *monteswitch*: A package for evaluating solid–solid free energy differences via lattice-switch Monte Carlo

T. L. Underwood<sup>a,\*</sup>, G. J. Ackland<sup>b</sup>

<sup>a</sup>*Department of Physics, University of Bath, Bath, BA2 7AY, UK*

<sup>b</sup>*School of Physics and Astronomy, SUPA, The University of Edinburgh, Edinburgh, EH9 3JZ, UK*

## Abstract

Lattice-switch Monte Carlo (LSMC) is a method for evaluating the free energy between two given solid phases. LSMC is a general method, being applicable to a wide range of problems and interatomic potentials. Furthermore it is extremely efficient, ostensibly more efficient than other existing general methods. Here we introduce a package, *monteswitch*, which can be used to perform LSMC simulations. The package can be used to evaluate the free energy differences between pairs of solid phases, including multicomponent phases, via LSMC for atomic (i.e., non-molecular) systems in the NVT and NPT ensembles. It could also be used to evaluate the free energy cost associated with interfaces and defects. Regarding interatomic potentials, *monteswitch* currently supports various commonly-used pair potentials, including the hard-sphere, Lennard-Jones, and Morse potentials, as well as the embedded atom model. However the main strength of the package is its versatility: it is designed so that users can easily implement their own potentials.

*Keywords*: free energy, phase transition, MPI, embedded atom model

## PROGRAM SUMMARY

*Manuscript Title*: monteswitch: A package for evaluating solid–solid free energy differences via lattice-switch Monte Carlo

*Authors*: T. L. Underwood, G. J. Ackland

*Program Title*: monteswitch

*Journal Reference*:

*Catalogue identifier*:

*Licensing provisions*: MIT

*Programming language*: Fortran 95, MPI

*Computer*: Any

*Operating system*: Unix-like

*RAM*: Depends on the nature of the problem.

*Number of processors used*: Any

*Keywords*: free energy, Monte Carlo, solid, phase transition, defect, MPI, embedded atom model, multicanonical

*Classification*: 7.7 Other Condensed Matter inc. Simulation of Liquids and Solids

*External routines/libraries*: To perform parallel simulations MPI is required (but MPI is not required to perform serial simulations).

*Nature of problem*: Calculating the free energy difference between two solid systems.

*Solution method*: Lattice-switch Monte Carlo (LSMC) [1] is a versatile and efficient method for evaluating

\*Corresponding author

the free energy difference between two solid phases. The package presented here allows LSMC simulations to be performed for a variety of interatomic potentials, including commonly-used pair potentials and the embedded atom model. Furthermore the package is designed so that users can easily implement their own potentials. The package supports LSMC simulations in the NVT and NPT ensembles, and can treat multicomponent systems. A version of the main program is included which is parallelised using MPI. This program parallelises the LSMC calculation by simulating multiple replicas of the system in parallel.

*Restrictions:* monteswitch cannot treat molecular systems, i.e., systems in which the particles exhibit rotational degrees of freedom, and is restricted to systems which can be represented within an orthorhombic supercell. Furthermore, the interatomic potential is ‘hard-coded’ in the sense that implementing a different potential requires that the package be recompiled.

*Additional comments:* monteswitch includes programs to assist with the creation of input files and the post-processing of output files created by the main Monte Carlo programs. A user manual, a suite of test cases, a worked example, and a collection of plug-ins to implement various commonly-used interatomic potentials are also included with the package.

*Running time:* Depends on the nature of the problem and the underlying computing platform. For the Zr EAM example in the manuscript one iteration (i.e., one 160,000-sweep weight-function-generation simulation and one 700,000-sweep production simulation) took a wall-clock time of approximately 1.9 hours on a desktop machine (an iMac14,2 with a 3.2GHz Intel Core i5-4570 processor) exploiting 4 cores for the 384-atom system, and 17.7 hours for the 1296-atom system. For each ensemble in the hard-sphere example the 18,000,000-sweep weight-function-generation simulation and two 125,000,000-sweep production simulations took a total of approximately 11 hours exploiting 16 cores on one node of a HPC cluster.

[1] A. D. Bruce, N. B. Wilding & G. J. Ackland, Phys. Rev. Lett. 79 3002 (1997)

## 1. Introduction

The stable phase under given conditions is that with the lowest free energy. For this reason, efficiently calculating free energies is one of the most fundamental problems in theoretical materials science. A plethora of different methods have been developed to this end, each designed with a particular problem in mind (see, e.g., Ref. [1]). Unfortunately however, commonly-used methods for calculating free energies of solid phases often cannot achieve the accuracy required for practical applications: an intractable amount of computational effort would be required. This problem is by no means limited to ‘complicated’ models of particle interactions, but persists even when simple models are used. For instance it was only relatively recently demonstrated that the fcc phase is favoured over the hcp phase in the hard-sphere solid – an archetype of a simple system [2, 3, 4].

*Lattice-switch Monte Carlo* (LSMC) [2, 3]<sup>1</sup> is a method which can be used to efficiently evaluate the free energy difference between two solid phases. It has been applied to a wide range of systems [2, 5, 6, 3, 4, 7, 8, 9, 10, 11, 12, 13, 14], beginning with the hard-sphere solid [2, 5, 6, 3], where it was used to resolve the aforementioned hcp–fcc problem [2, 3]. The method was later applied to soft interatomic potentials [4, 11, 14], systems containing multiple particle species [7, 8], and molecular systems [9, 10, 12, 13]. LSMC has also inspired *phase switch Monte Carlo*, a method for calculating the free energy difference between a solid and a fluid phase [15], which has also seen some use [15, 16, 17, 18, 19, 20, 21]. As well as being versatile, LSMC is an accurate method: it is ‘exact’ in the sense that it relies upon no approximations other than those present in the model of particle interactions it is used in conjunction with. Moreover for the purposes of evaluating the free energy difference between pairs of solid phases LSMC is

---

<sup>1</sup>The reader should be aware that LSMC has also been referred to as *lattice-switching* Monte Carlo.

ostensibly more efficient than other existing general methods [11, 9].<sup>2</sup><sup>3</sup> However, despite its strengths, LSMC has unfortunately yet to have gained widespread popularity. This stems in part from the lack of an LSMC code which is both widely available and applicable to a wide range of systems.

With this in mind we have developed a package, *monteswitch*, which implements the LSMC method. The package, written in Fortran 95, can be used to evaluate free energy differences between pairs of solid phases in the NVT and NPT ensembles. Furthermore the package contains a version of the main executable which is parallelised using MPI for HPC applications. Note that the two ‘phases’ under consideration need not necessarily be homogeneous crystals; an interesting prospect is to use *monteswitch* to evaluate free energy costs associated with interfaces and defects – the former is something which has been done previously using LSMC [5].<sup>4</sup> Furthermore *monteswitch* can treat systems containing multiple species of particles. However it should be noted that *monteswitch* can only treat ‘atomic’ systems (i.e., ‘non-molecular’ systems: those in which the constituent particles do not have rotational degrees of freedom), and pairs of phases which can be represented by orthorhombic unit cells.

While steps have been recently taken to implement LSMC in an existing general-purpose code,<sup>5</sup> we believe that *monteswitch* will fulfil an important ‘gap in the market’ for the foreseeable future because it was designed from the outset to be highly-customisable with regards to the interatomic potential. By contrast general-purpose codes tend to have a fixed set of interatomic potentials to draw upon. In *monteswitch* all of the procedures pertaining to the interatomic potential are housed within a single Fortran module. It is intended that users write their own version of this module which implements the interatomic potential they are interested in.<sup>6</sup> (A similar scheme is utilised in the molecular dynamics program MOLDY [27]). Templates are provided with *monteswitch* to assist with this. Furthermore modules are included with *monteswitch* which correspond to some commonly-used interatomic potentials, which can serve as examples. Of course these modules can also be used within *monteswitch* to perform LSMC calculations.

Here we provide an introduction to *monteswitch*. Note that much of what follows is elaborated upon in *monteswitch*’s user manual (included with the package), where we direct interested readers for more details. The layout of this work is as follows. In the next section we describe the theory which underpins *monteswitch*. In Section 3 we provide an overview of what is included in the *monteswitch* package. In Section 4 we describe how interatomic potentials are implemented in *monteswitch*, list the various interatomic potentials included with *monteswitch*, and describe how users can implement their own potentials. In Section 5 we describe the main Monte Carlo programs within *monteswitch*. In Section 6 we describe the various utility programs included with *monteswitch* for the creation of input files and post-processing of output files. In Section

---

<sup>2</sup>We do not include methods rooted in the harmonic approximation (including the quasi-harmonic approximation [22, 23]) within the class of ‘general methods’ mentioned here: these methods are not ‘general’ in the sense that they break down in the anharmonic regime.

<sup>3</sup>To elaborate, in Refs. [11, 9] LSMC was shown to significantly outperform thermodynamic integration (TI) [1, 24]. However the claim that LSMC outperforms TI has proved contentious [5, 25]. Of course, like-for-like comparisons between the two methods are difficult, since different implementations of LSMC or TI may be more or less efficient than other implementations. We believe that the claim that LSMC is *at least* as efficient as TI reflects the findings of studies up to the present time.

<sup>4</sup>We describe how LSMC can be used to evaluate interfacial free energies in Section 8.

<sup>5</sup>Specifically, LSMC is earmarked for inclusion in the general-purpose Monte Carlo code *DLMONTE* [26].

<sup>6</sup>Of course, in doing this the user’s module is free to interface with ‘external’ modules, or even external programs.

7 we provide two examples to elucidate how *monteswitch* could be used in practice. In the first example we apply *monteswitch* to the hard-sphere solid, and test *monteswitch* against known LSMC results for this system; in the second we use *monteswitch* to determine the hcp–bcc transition temperature and related quantities for an embedded atom model of Zr. Finally in Section 8 we present our conclusions and outlook.

## 2. Theoretical background

### 2.1. Calculating free energy differences

Consider a system which is free to visit two phases 1 and 2 (and only phases 1 and 2). The equilibrium phase is that with the lower free energy  $\mathcal{F}$ , where  $\mathcal{F}$  is the Helmholtz free energy in the NVT ensemble and the Gibbs free energy in the NPT ensemble. It is the free energy difference between the phases  $\Delta\mathcal{F} \equiv \mathcal{F}_1 - \mathcal{F}_2$  which we wish to evaluate, where  $\mathcal{F}_1$  and  $\mathcal{F}_2$  denote the free energies of phases 1 and 2. It can be shown that

$$\Delta\mathcal{F} = \beta^{-1} \ln\left(\frac{p_2}{p_1}\right), \quad (1)$$

where  $p_1$  and  $p_2$  denote the probability of the system being in phase 1 and 2 respectively,  $\beta = 1/(k_B T)$ ,  $k_B$  denotes Boltzmann’s constant, and  $T$  denotes the temperature of the system. For a simulation which samples the ensemble under consideration, e.g., molecular dynamics,  $p_2/p_1$  can be determined: measure the relative time  $t_1$  and  $t_2$  which the system spends in each phase 1 and 2 during the simulation, and substitute these quantities into the above equation, bearing in mind that  $t_2/t_1 = p_2/p_1$  for a sufficiently long simulation. Hence  $\Delta\mathcal{F}$  can in principle be obtained from such a simulation via the above equation. However, this method is usually intractable in practice for two solid phases, because the time taken for the system to transition between the two phases is too long to allow a reasonable estimate of  $p_2/p_1$  to be deduced in a reasonable simulation time. It may even be the case that, regardless of the phase in which the simulation is initialised, the system *never* transitions to the ‘other’ phase during the course of the simulation. The problem is that, while the regions of phase space corresponding to phase 1 and phase 2 both correspond to probable states of the system at thermodynamic equilibrium, these regions are separated by a *free energy barrier* – a region of phase space associated with states which are very improbable at thermodynamic equilibrium. This barrier inhibits transitions between the regions of phase space associated with phase 1 and phase 2.

This problem can in principle be circumvented with the Monte Carlo method. In the original incarnation of Monte Carlo, which we refer to as *canonical Monte Carlo*[28] (which we contrast with *multicanonical* Monte Carlo later), the system is evolved during the simulation as follows. Each time step we generate a trial state of the system  $\sigma'$ , and attempt to change the system to the trial state from its current state  $\sigma$ . The traditional approach for NVT ensembles is to perform a ‘particle move’ to generate a trial state. Here, one particle in  $\sigma$  is moved to yield  $\sigma'$ . In NPT ensembles particle moves are supplemented by ‘volume moves’, in which the volume, and potentially the shape, of the entire system is altered, along with a commensurate rescaling of the particle positions. We accept the change of state from  $\sigma$  to  $\sigma'$  with a probability  $p_{\sigma \rightarrow \sigma'}$ , which is a function of the energies of the states  $\sigma$  and  $\sigma'$  in the NVT ensemble, and the enthalpies and volumes of the states  $\sigma$  and  $\sigma'$  in the NPT ensemble. The function also depends on the specific scheme used to generate state  $\sigma'$  from  $\sigma$  (see, e.g., Ref. [1]). The end result is that each state

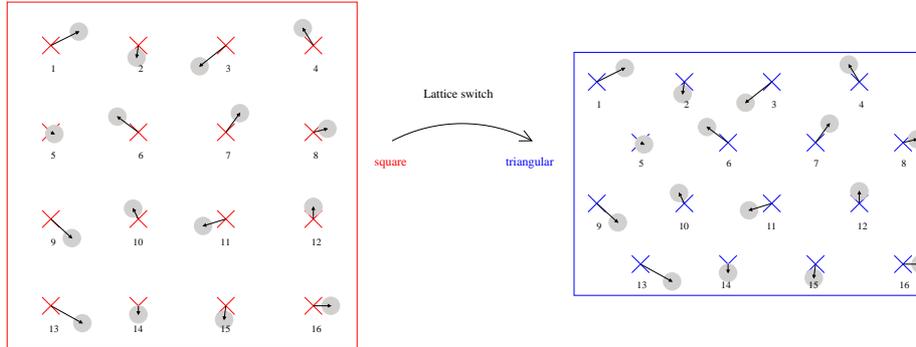


Figure 1: Schematic diagram illustrating a lattice switch from a state in the square phase of a notional two-dimensional system to a state in the triangular phase. In the lattice switch the underlying square lattice (red crosses) is transformed into a triangular lattice (blue crosses), while the displacements of the particles from their lattice sites (black arrows) is unchanged – the displacement for each particle  $n$  is the same before and after the lattice switch. Note also that the lattice switch here transforms the shape of the system: the red box is transformed into the blue box.

$\sigma$  is sampled with a probability  $p_\sigma$  which reflects the underlying ensemble, e.g., for the NVT ensemble:

$$p_\sigma \propto e^{-\beta E_\sigma}, \quad (2)$$

where  $E_\sigma$  denotes the energy of state  $\sigma$ . However, in canonical Monte Carlo one has considerable freedom as to how trial states are generated; one is by no means limited to the aforementioned ‘traditional’ move set. The prospect therefore exists of generating trial states in a manner which results in the system traversing a path in phase space which allows  $\Delta\mathcal{F}$  to be calculated in a reasonable simulation time. Such a path would involve frequent transitions between both phases 1 and 2 by ‘tunnelling’ through the free energy barrier separating them.

## 2.2. Lattice-switch moves

In LSMC a new type of move, a *lattice switch*, is introduced to supplement the traditional move set mentioned above. A lattice switch move takes the system *directly* from one phase to the other, bypassing any free energy barriers separating the phases. Every time a lattice switch is accepted, the system transitions to the ‘other’ phase. The salient feature of the move is that the underlying ‘lattice’ which characterises the current phase is ‘switched’ for a lattice which characterises the other phase, while the *displacements* of all particles from their associated lattice sites are preserved. This is illustrated in Fig. 1 for the square and triangular phases of a notional two-dimensional system.

More formally, we can characterise a given state of the system as belonging to a solid phase  $\alpha$  if the positions of the particles ‘approximately’ form a lattice characteristic of  $\alpha$ . Let  $\{\mathbf{R}_i^{(\alpha)}\}$  denote the positions of the sites on this lattice, and let  $\{\mathbf{r}_i\}$  denote the positions of the particles, where  $i$  ranges from 1 to the number of particles in the system. The position  $\mathbf{r}_i$  of particle  $i$  can be expressed as follows:

$$\mathbf{r}_i = \mathbf{R}_i^{(\alpha)} + \mathbf{u}_i, \quad (3)$$

where  $\mathbf{u}_i$  is the displacement of  $i$  from that lattice site. Note that the displacements  $\{\mathbf{u}_i\}$  are necessarily small since the particle positions form an approximate  $\alpha$  lattice (and we have chosen to label particles and lattice sites in a ‘sensible’ manner: such that  $\mathbf{R}_i$  is the closest lattice site to

$\mathbf{r}_i$ ). Now, in a lattice switch from phase  $\alpha$  to the ‘other’ phase  $\alpha'$  we transform the underlying lattice  $\{\mathbf{R}_i^{(\alpha)}\}$  to  $\{\mathbf{R}_i^{(\alpha')}\}$ , while keeping the particle displacements  $\{\mathbf{u}_i\}$  unchanged. The result is that the trial state belongs to phase  $\alpha'$ : the positions of the particles in the trial state form an approximate  $\alpha'$  lattice.

Of course, the above description of a lattice switch is not a complete account of how lattice switches are implemented in *monteswitch* – which supports lattice switches which change the shape and size of the supercell, as well as the species of the particle. Details of how lattice switches are implemented in *monteswitch* can be found in the user manual.

### 2.3. Multicanonical Monte Carlo

One might expect that by regularly making lattice switches, the system will regularly transition between phases, and hence allow  $\Delta\mathcal{F}$  to be efficiently evaluated as described above. Unfortunately using canonical Monte Carlo one finds that lattice switches are too rarely accepted for this approach to be useful. The problem is that the trial state  $\sigma'$  generated by a lattice switch is almost always of much higher energy than the current state  $\sigma$ , and hence will almost always be rejected.<sup>7</sup> The solution to this problem is to use *multicanonical Monte Carlo*[29, 30, 31] instead of canonical Monte Carlo. Multicanonical Monte Carlo can be regarded as canonical Monte Carlo, but if the energy for each state  $\sigma$  were

$$\tilde{E}_\sigma = E_\sigma - \eta_\sigma/\beta \quad (4)$$

instead of  $E_\sigma$ , where  $\eta_\sigma$ , known as the *weight function*, is chosen according to the aims of the simulation. Note that if  $\eta_\sigma > 0$  then state  $\sigma$  is sampled more frequently than would be the case for the ensemble of interest; and if  $\eta_\sigma < 0$  then  $\sigma$  is sampled less frequently. The strength of this approach is that through judicious choice of the weight function, one can ‘control’ the path the system traverses through phase space.

Of course, in a multicanonical simulation the states are no longer sampled with probabilities corresponding to the true ensemble in question – which is the case for canonical Monte Carlo. Accordingly the time average of some physical quantity  $X$  throughout a long multicanonical Monte Carlo simulation is not equivalent to the equilibrium value of  $X$  for the true ensemble, as it is in a canonical Monte Carlo simulation. Nevertheless one can obtain the equilibrium value of  $X$  from a multicanonical simulation by exploiting the fact that, since the weight function is known, then so is the degree of over- or under-sampling of each state. To elaborate, the equilibrium value of  $X$  in a multicanonical Monte Carlo simulation is given by

$$\langle X \rangle \approx \frac{\sum_{t=1}^{\tau} e^{-\eta(t)} X(t)}{\sum_{t=1}^{\tau} e^{-\eta(t)}}, \quad (5)$$

where  $X(t)$  denotes the quantity  $X$  corresponding to the state sampled at timestep  $t$ , and  $\tau$  denotes the total number of timesteps.

---

<sup>7</sup>The situation is slightly more complicated for lattice switches which change the system volume. In this case the extent to which the volume of the system is expanded/contracted influences how likely the lattice switch is to be successful. Accordingly the order parameter defined later in Eqn. (7) for ‘selecting’ gateway states (defined in a moment) takes a slightly different form for volume-altering lattice switches in *monteswitch* – see the *monteswitch* user manual for more details. Aside from that the forthcoming discussion applies generally: to both volume-altering and volume-preserving lattice switches.

#### 2.4. Multicanonical Monte Carlo in LSMC

How does multicanonical Monte Carlo resolve the problem that lattice switch moves are too rarely accepted to be useful? Recall that lattice switches are usually rejected because they result in a trial state with a much higher energy. There are, however, a small number of states ‘close’ to those realised at equilibrium for each phase from which a lattice switch yields a trial state  $\sigma'$  which is of comparable energy to  $\sigma$ . From such states a lattice switch has a good chance of being accepted. We refer to such states as *gateway states*, since they provide the key to jumping between both phases. It is these states which we wish to over-sample, and we set the weight function accordingly. The idea is that by over-sampling these states, lattice switches are accepted reasonably often, enabling both phases to be explored in a reasonable simulation time. This in turn allows us to determine  $p_1$  and  $p_2$ , and hence  $\Delta\mathcal{F}$  via Eqn. (1). Specifically,  $p_1$  and  $p_2$  are obtained via Eqn. (5):

$$p_\alpha = \langle \theta_\alpha \rangle \approx \frac{\sum_{t=1}^{\tau} e^{-\eta(t)} \theta_\alpha(t)}{\sum_{t=1}^{\tau} e^{-\eta(t)}}, \quad (6)$$

where  $\theta_\alpha(t)$  takes the value 1 if the system is in phase  $\alpha$  at timestep  $t$  and 0 otherwise.

How should the weight function be engineered such that gateway states are over-sampled? Consider a state  $\sigma$ , and let  $\sigma'$  denote the state which results from a lattice switch performed from  $\sigma$ . Let us define the state-dependent quantity

$$M_\sigma = \begin{cases} (E_\sigma - E_{\sigma'}) & \text{if } \sigma \text{ belongs to phase 1} \\ -(E_\sigma - E_{\sigma'}) & \text{if } \sigma \text{ belongs to phase 2.} \end{cases} \quad (7)$$

This quantity provides a practical means for resolving gateway states, states corresponding to equilibrium (for the true ensemble under consideration) for phase 1, and states corresponding to equilibrium for phase 2. Accordingly we refer to  $M$  as the *order parameter*. Consider first gateway states. Above we illustrated that gateway states correspond to the condition  $E_\sigma \approx E_{\sigma'}$ . As can be seen from the above this corresponds to states with  $M_\sigma \approx 0$ . By contrast, if  $|M_\sigma| \gg 0$ , then the two states have significantly different values of  $E$ . In this case, while switching from the state with the higher value of  $E$  to that with the lower value of  $E$  is guaranteed, the converse is not: the two states are not concordant with switching *to and from* both phases.  $|E_\sigma|$  therefore provides a measure of how ‘un-gateway-like’ state  $\sigma$  is, with zero corresponding to ‘very gateway-like’. Consider now phase-1-equilibrium states. From such states we generally expect a lattice switch to be unsuccessful, and hence  $E_{\sigma'} \gg E_\sigma$ . Therefore for such states  $M_\sigma \ll 0$ . Finally consider phase-2-equilibrium states. Similarly we generally expect a lattice switch from such states to be unsuccessful, and hence  $E_{\sigma'} \gg E_\sigma$ . However this time  $M_\sigma \gg 0$ . We therefore have three regimes:  $M_\sigma \ll 0$  corresponds to phase-1-equilibrium states;  $M_\sigma \approx 0$  corresponds to gateway states; and  $M_\sigma \gg 0$  corresponds to phase-2-equilibrium states.

With this in mind, if we choose the weight function  $\eta_\sigma$  to take the same value  $\eta_M$  for all states with the same  $M$  and also choose  $\eta_M$  to be peaked at  $M = 0$  and to decay monotonically with  $|M|$ , then the weight function corresponds to a ‘force’ which drives the system towards gateway states, allowing the system to transition between the phase-1 and phase-2 regions of phase space, corresponding to  $M \ll 0$  and  $M \gg 0$  respectively, in a reasonable simulation time. This is, of course, just a *qualitative* description of a form for  $\eta_M$  which is sufficient for our purposes. As

one might expect, the quantitative details of the weight function  $\eta_M$  strongly affect the efficiency of the path traversed through phase space with regards to calculating  $\Delta\mathcal{F}$ ; a ‘bad’ weight function might result in the system getting stuck in one phase, or an unimportant region of phase space, for a long time. Furthermore, it is not obvious *a priori* what a suitable weight function for a given system should be. Hence one must *generate* a weight function which leads to an efficient sampling of phase space. After this *weight function generation simulation*, the resulting weight function can be used in a *production simulation* to calculate  $\Delta\mathcal{F}$  as described earlier.

However in practice one cannot treat  $M$  as an unbounded continuous variable as above; one cannot define an arbitrary weight function in computer memory via this scheme. Hence in practice one considers a finite range of  $M$  which is divided into  $N_{\text{macro}}$  bins, each corresponding to a distinct range of ‘ $M$ -space’. Each bin itself corresponds to a macrostate: the macrostate is the collection of states corresponding to the range of  $M$ -space covered by the bin. We will henceforth explicitly take this discretisation of  $M$  into account, and let  $\mathcal{M}$  denote the macrostate corresponding to the  $M$ th bin, where  $\mathcal{M} = 1, 2, \dots, N_{\text{macro}}$ . Accordingly let  $\eta_{\mathcal{M}}$  denote the weight function for macrostate  $\mathcal{M}$ .

## 2.5. Weight function generation

The weight function can be generated in many different ways, some of which are more efficient than others. We now list the methods implemented in *monteswitch*. All of these methods share the same notion of the ‘ideal’ weight function  $\eta_{\mathcal{M}}^*$ , which leads to all macrostates within the considered order parameter range to be sampled with equal probability in the multicanonical Monte Carlo simulation.

### 2.5.1. Visited states method

The *visited states method* (see Ref. [31] and references therein) is arguably the simplest method for generating the weight function. In the visited states method, the simulation consists of a number of ‘blocks’, which themselves consist of a large number of Monte Carlo sweeps. Multicanonical sampling is used throughout, and the weight function is updated at the end of each block. The weight function is different – closer to the ideal – in each subsequent block, and the number of visits to all macrostates during each block is used to inform the weight function to be used in the next block. Eventually the weight function converges on the ideal: it provides a ‘flat’ macrostate histogram; the weight function is such that all macrostates are sampled with equal probability. Specifically, the following scheme is used to update the weight function at the end of each block:

$$\eta_{\mathcal{M}}^{(n+1)} = \eta_{\mathcal{M}}^{(n)} - \ln \left\{ \frac{C_{\mathcal{M}}^{(n)} + 1}{\sum_{\mathcal{M}'} (C_{\mathcal{M}'}^{(n)} + 1)} \right\} + k, \quad (8)$$

where where  $C_{\mathcal{M}}^{(n)}$  denotes the number of states belonging to macrostate  $\mathcal{M}$  visited during block  $n$ ;  $\eta_{\mathcal{M}}^{(n)}$  denotes the weight function for block  $n$ ; the summation over  $\mathcal{M}'$  on the denominator of the fraction is over all macrostates  $1, 2, \dots, N_{\text{macro}}$ ; and  $k$  is an inconsequential arbitrary constant, which we choose such that the minimum value of  $\eta_{\mathcal{M}}^{(n+1)}$  over all  $\mathcal{M}$  is 0.

### 2.5.2. Transition matrix method

A more sophisticated method than the visited states method, which is significantly more efficient, is the *transition matrix method* [31, 32]. This method exploits the fact that the ideal

weight function  $\eta_M^*$  is related to the *canonical* probability  $p_M$  of the system being in macrostate  $M$  via the equation

$$\eta_M^* = A - \ln p_M, \quad (9)$$

where  $A$  is an arbitrary constant.  $p_M$  in turn can be determined from the *macrostate transition probability matrix*  $\mathcal{T}_{MM'}$ , which describes the probability that the system, currently in macrostate  $M$ , transitions to macrostate  $M'$  in the *canonical ensemble*. In the transition matrix method we determine  $\mathcal{T}_{MM'}$ , and then use this to obtain  $p_M$ , and finally the ideal weight function  $\eta_M^*$  via Eqn. (9).

$\mathcal{T}_{MM'}$  is determined as follows. During the simulation we keep track of the number of transitions between all pairs of macrostates, which we store in a matrix  $C_{MM'}$  – where  $C_{MM'}$  denotes the number of transitions from macrostate  $M$  to macrostate  $M'$ . We then use  $C_{MM'}$  to obtain an estimate for  $\mathcal{T}_{MM'}$  via the equation

$$\mathcal{T}_{MM'} \approx \frac{C_{MM'} + 1}{\sum_{M''} (C_{MM''} + 1)}. \quad (10)$$

However  $C_{MM'}$  is not simply the number of *observed* transitions from  $M$  to  $M'$  during the simulation, but rather the number of *inferred* transitions. To elaborate, consider a trial state  $\sigma'$  generated from a state  $\sigma$  which, if accepted, would take the system from macrostate  $M$  to macrostate  $M'$ , and let the *canonical* probability of the move being accepted be  $p_{\sigma \rightarrow \sigma'}$ . Instead of performing the update  $C_{MM'} \rightarrow C_{MM'} + 1$  if the move is accepted and  $C_{MM'} \rightarrow C_{MM'}$  if it is not – which would result in  $C_{MM'}$  being the number of observed transitions from  $M$  to  $M'$  – we perform the update

$$\begin{aligned} C_{MM'} &\rightarrow C_{MM'} + p_{\sigma \rightarrow \sigma'} \\ C_{MM} &\rightarrow C_{MM} + 1 - p_{\sigma \rightarrow \sigma'} \end{aligned} \quad (11)$$

regardless of whether it is accepted or not. Note that the canonical quantity  $p_{\sigma \rightarrow \sigma'}$  is always used in the update procedure, which leads to  $C_{MM'}$  being the inferred number of canonical transitions between  $M$  and  $M'$ . Because of this one can use any method for exploring  $M$ -space: canonical, multicanonical, or something else. We elaborate on this point in a moment.

Having determined  $\mathcal{T}_{MM'}$ , our task is to now calculate  $p_M$ . It can be shown that the macrostates obey the following detailed balance condition [31]:

$$\mathcal{T}_{M'M} p_M = \mathcal{T}_{MM'} p_{M'}. \quad (12)$$

Setting  $M' = M + 1$  and rearranging the above gives

$$p_{(M+1)} = \frac{\mathcal{T}_{M(M+1)}}{\mathcal{T}_{(M+1)M}} p_M. \quad (13)$$

Using this equation,  $p_M$  can be obtained from the matrix  $\mathcal{T}_{MM'}$  via the following procedure. Firstly, one chooses some arbitrary value for  $p_1$ .<sup>8</sup> With this  $p_2$  can be obtained from the above equation ( $M = 1$  in Eqn. (13)). This in turn can be used to obtain  $p_3$  ( $M = 2$  in Eqn. (13)),

---

<sup>8</sup>In this section  $p_1$  and  $p_2$  denote the probability of the system being in macrostates 1 and 2, *not* the probabilities of the system being in phases 1 and 2.

which in turn can be used to obtain  $p_4$ , etc., until  $p_{N_{\text{macro}}}$  is obtained. Finally, one normalises the resulting function  $p_M$  such that

$$\sum_{M=1}^{N_{\text{macro}}} p_M = 1, \quad (14)$$

as is required. The final step is to use  $p_M$  to obtain an estimate for the ideal weight function. This is done simply by substituting  $p_M$  into Eqn. (9).

### 2.5.3. Methods for exploring $M$ -space

As alluded to above, since the updates to  $C_{MM'}$  always use the canonical probabilities of transitioning between states, with the transition matrix method one can *choose* how  $M$ -space is explored. *monteswitch* supports a number of ways of doing this.

The first method is to use multicanonical sampling to explore  $M$ -space with a continuously evolving weight function, where the weight function at a given time is the current estimate for the ideal weight function derived from the current  $C_{MM'}$  as described above. This is the ‘natural’ way of applying the transition state method.

The second method is to use what we refer to as *artificial dynamics* to force the system to explore all macrostates in a reasonable amount of time. In this method, the system is first locked into a macrostate for a certain period of time. After that period of time has elapsed, the ‘barriers’ preventing the system from moving into an adjacent macrostate is moved such that the system is free to transition into an adjacent macrostate. Once this occurs, the system is locked into this new macrostate, and the procedure starts again. There is of course the question of which adjacent macrostate to ‘open’ to the system. Assuming we are not in macrostate  $M = 1$  or  $N_{\text{macro}}$ , then there are two options:  $(M + 1)$  and  $(M - 1)$ . In *monteswitch* one can specify whether to select the new macrostate at random [33], or whether to sweep through the macrostates systematically, e.g., to explore macrostates  $3, 4, 5, \dots, (N_{\text{macro}} - 1), N_{\text{macro}}, (N_{\text{macro}} - 1), \dots, 3, 2, 1, 2, 3, \dots$ . This method is faster than the ‘natural’ method just described because one does not have to wait for the weight function to evolve such that it pushes the system to explore macrostates which are unlikely to be visited in the canonical ensemble.

## 3. Package structure

The *monteswitch* package consists of a number of programs, as well as a user manual, a suite of test cases, a worked example, and a suite of Fortran modules corresponding to different interatomic potentials. The programs are:

- monteswitch
- monteswitch\_mpi
- monteswitch\_post
- lattices\_in\_hcp\_fcc
- lattices\_in\_bcc\_fcc
- lattices\_in\_bcc\_hcp

`monteswitch` and `monteswitch_mpi` are the key programs of the package: they perform Monte Carlo simulations. By contrast `monteswitch_post`, `lattices_in_hcp_fcc`, `lattices_in_bcc_fcc` and `lattices_in_bcc_hcp` are utility programs: `monteswitch_post` is for post-processing one of the output files created by the main programs; and `lattices_in_hcp_fcc`, `lattices_in_bcc_fcc` and `lattices_in_bcc_hcp` are for generating one of the input files for the main programs. We elaborate upon these programs in later sections.

#### 4. Interatomic potentials

The file `interactions.f95` in the main directory of the package contains the Fortran module, named `interactions_mod`, which determines the interatomic potential to be utilised by the *monteswitch* programs `monteswitch`, `monteswitch_mpi` and `monteswitch_post` after the package is compiled. By default `interactions.f95` corresponds to the embedded atom model (EAM) [34]; to implement a specific interatomic potential one must copy the corresponding `interactions.f95` file to `interactions.f95` in the main directory of the package, and then compile the package.

##### 4.1. Structure of `interactions_mod`

The module `interactions_mod` contains the following procedures which interface with the main *monteswitch* programs `monteswitch` and `monteswitch_mpi`:

- `initialise_interactions`, which initialises the variables within the module, possibly by reading variables from one or more input files, for ‘new’ simulations;
- `export_interactions`, which exports the module variables to a file for the purposes of checkpointing the simulation;
- `import_interactions`, which imports the module variables from the aforementioned file to resume a checkpointed simulation;
- `after_accepted_part_interactions`, `after_accepted_vol_interactions` and `after_accepted_lattice_interactions`, which perform any housekeeping tasks required by the module (e.g., updating neighbour lists) after, respectively, a particle, volume and lattice switch move has been accepted;
- `after_all_interactions`, which performs any housekeeping tasks required by the module after all moves, including rejected moves;
- `calc_energy_scratch`, which calculates the energy of the system ‘from scratch’ for a specified state;
- `calc_energy_part_move`, which calculates the energy of the system given that one particle has moved.

Users wishing to write their own versions of `interactions.f95` to implement their own interatomic potentials must write their own versions of each of the above procedures. To assist with this, two templates for `interactions.f95` are provided with *monteswitch*. These can be found in the directory `Interactions` within the package. The file `interactions_TEMPLATE_minimal.f95` contains a ‘bare’ version of `interactions.f95`,

i.e., it leaves the above procedures empty to be filled in by the user. The file `interactions_TEMPLATE_pair.f95` allows quick generation of `interactions.f95` files for pair potentials, something which we elaborate upon below. Both templates contain comments which provide guidance to the user.

Note that there is freedom in how the module variables are initialised for a new simulation – via the procedure `initialise_interactions`. Normally `initialise_interactions` would read the variables – which parametrise the potential under consideration – from one or more input files. We emphasise that these input files, and their formats, depend on the specific version of `interactions.f95` which `monteswitch` is used in conjunction with.

#### 4.2. Potentials included with `monteswitch`

The Interactions directory within the package contains a number of versions of `interactions.f95`, which correspond to various interatomic potentials. The most important of these are as follows.

##### 4.2.1. Embedded atom model

The file `interactions_EAM.f95` implements the embedded atom model (EAM) [34] for metals (but not alloys). Here the energy of the system is given by

$$E = \frac{1}{2} \sum_{i,j \neq i} \phi(r_{ij}) + \sum_i F(\rho_i), \quad (15)$$

$$\rho_i = \sum_{j \neq i} \rho(r_{ij}), \quad (16)$$

where  $r_{ij}$  is the separation between particles  $i$  and  $j$ , and  $\phi$ ,  $F$  and  $\rho$  are functions which must be specified and constitute the parametrisation of the EAM potential. If `interactions_EAM.f95` is used then one input file, named `interactions_in`, is required by the programs `monteswitch` and `monteswitch_mpi` to input the potential for new simulations. This file must be a description of the EAM potential to be used in DYNAMO/LAMMPS ‘setfl’ format [35].

##### 4.2.2. Soft pair potentials

Table 4.2.2 gives a list of soft pair potentials included with `monteswitch`, along with the name of the corresponding `interactions.f95` file in Interactions. All of these potentials are implemented in the same way. Firstly, the pair potential  $\phi(r)$  is assumed to be 0 for inter-particle separations  $r$  greater than some cut-off distance  $r_c$ . In other words the potential is *truncated* at  $r_c$ . Secondly, only pairs of particles within a distance  $r_{\text{list}}$  of each other at the start of the simulation interact with each other throughout the entire simulation. To clarify the difference between  $r_c$  and  $r_{\text{list}}$ : the former is the distance at which the potential is truncated, while the latter determines which particles are in each others’ *neighbour list*, which remains constant throughout the simulation.

As for the EAM potential just described, for the soft potentials an input file `interactions_in`, is required by the programs `monteswitch` and `monteswitch_mpi` to input the potential for new simulations. The format of this file is as follows: each variable which parametrises the potential corresponds to a specific line in `interactions_in`, and each line must contain a string (we recommend the name of the variable followed immediately by an ‘=’ character with no spaces), followed by whitespace, followed by the value of the variable. The

final column of Table 4.2.2 gives the order of variables as they should appear, one per line, in `interactions_in` for each potential. Note that in all cases  $r_c$  and  $r_{\text{list}}$  are included as input variables. Furthermore there is a variable  $M_{\text{list}}$ , which determines the size of the array used in the program to store the neighbour lists. Details regarding this variable can be found in the user manual. To illustrate the above, here is an example `interactions_in` file for the Lennard-Jones potential (`interactions_LJ.f95`):

```
lj_epsilon= 1.0
lj_sigma= 1.0
lj_cutoff= 1.5
list_cutoff= 1.00000001
list_size= 14
```

Table 1: Soft interatomic pair potentials included with *monteswitch*

Potential	File name	Expression for potential	Order of variables in <code>interactions_in</code>
12-10	<code>interactions_12-10.f95</code>	$A/r^{12} - B/r^{10}$	$A, B, r_c, r_{\text{list}}, N_{\text{list}}$
12-6	<code>interactions_12-6.f95</code>	$A/r^{12} - B/r^6$	$A, B, r_c, r_{\text{list}}, N_{\text{list}}$
Buckingham	<code>interactions_Buckingham.f95</code>	$A \exp(-r/\rho) - C/r^6$	$A, \rho, C, r_c, r_{\text{list}}, N_{\text{list}}$
Gaussian	<code>interactions_Gaussian.f95</code>	$-A \exp(-Br^2)$	$A, B, r_c, r_{\text{list}}, N_{\text{list}}$
Lennard-Jones	<code>interactions_LJ.f95</code>	$4\epsilon[(\sigma/r)^{12} - (\sigma/r)^6]$	$\epsilon, \sigma, r_c, r_{\text{list}}, N_{\text{list}}$
9-6 Lennard-Jones	<code>interactions_LJ_9-6.f95</code>	$4\epsilon[(\sigma/r)^9 - (\sigma/r)^6]$	$\epsilon, \sigma, r_c, r_{\text{list}}, N_{\text{list}}$
Morse	<code>interactions_Morse.f95</code>	$E_0 \left\{ \left[ 1 - \exp(-k(r - r_0)) \right]^2 - 1 \right\}$	$E_0, k, r_0, r_c, r_{\text{list}}, N_{\text{list}}$
<i>n-m</i>	<code>interactions_n-m.f95</code>	$[E_0/(n - m)][m(r_0/r)^n - n(r_0/r)^m]$	$n, m, r_0, r_c, r_{\text{list}}, N_{\text{list}}$
Yukawa	<code>interactions_Yukawa.f95</code>	$A \exp(-kr)/r$	$A, k, r_c, r_{\text{list}}, N_{\text{list}}$

#### 4.2.3. User-defined pair potentials

As mentioned above, the file `interactions.TEMPLATE_pair.f95` is a template which can be used to easily create `interactions.f95` files for user-defined pair potentials. Instructions are provided in the file regarding how the file should be modified to realise the user's potential of interest. In fact this template was used to create the files for almost all of the pair potentials described above.

#### 4.2.4. Hard/penetrable spheres

The file `interactions_HS_multi.f95` implements the penetrable (including hard) spheres model, where the sphere diameter is allowed to vary with particle species. Here, the pair potential between two particles belonging to species  $s$  and  $t$  is given by

$$\phi_{st}(r) = \begin{cases} \epsilon & \text{if } r < \frac{1}{2}(\sigma_s + \sigma_t) \\ 0 & \text{otherwise,} \end{cases} \quad (17)$$

where  $\sigma_s$  denotes the diameter of spheres belonging to species  $s$ . Again, the input file for the new simulations in the programs `monteswitch` and `monteswitch_mpi` is named `interactions_in`, and its format is similar to the `interactions_in` files for the soft pair potentials described above. In this case however the order of the variables is  $\epsilon$ ,  $N_{\text{species}}$ ,  $\sigma$ ,  $r_{\text{list}}$ , and  $N_{\text{list}}$ , where  $N_{\text{species}}$  is the number of species to consider, and  $\sigma$  is an  $N_{\text{species}}$ -dimensional vector (to be specified on a single line `interactions_in`) containing the diameters for each species  $1, 2, \dots, N_{\text{species}}$ . Note that hard spheres correspond to the limit  $\beta\epsilon \rightarrow \infty$ , and hence hard spheres can be implemented by setting  $\epsilon$  to a high value.

## 5. Monte Carlo simulation programs

The key programs in the package are `monteswitch` and `monteswitch_mpi`. The main purpose of these programs is to perform LSMC simulations, though they can also be used to perform 'conventional' Monte Carlo simulations. `monteswitch_mpi` is the MPI-parallelised analogue of `monteswitch`. While in `monteswitch` one simulation is performed, in `monteswitch_mpi` multiple simulations are performed in parallel. These simulations are identical except for the seed used for the random number generator. Hence `monteswitch_mpi` is simply a convenient means to exploit parallelisation in order to obtain results quickly. In terms of usage and simulation methodology `monteswitch` and `monteswitch_mpi` are almost identical. For this reason we focus on `monteswitch` below, where it should be assumed that what is said for `monteswitch_mpi` also applies for `monteswitch_mpi` unless otherwise stated.

### 5.1. Overview of functionality

`monteswitch` can treat phases which can be represented by orthorhombic supercells in the NVT and NPT ensembles, where in the NPT ensemble both isotropic volume moves and volume moves which allow the shape of the system to alter are supported. Multicomponent systems are allowed, however `monteswitch` cannot treat 'molecular' systems in which the particles have orientational degrees of freedom: rotational Monte Carlo moves are not implemented in `monteswitch`.

`monteswitch` supports both canonical and multicanonical sampling, where the multicanonical sampling is performed over the LSMC order parameter described in Section 2.4. Regarding

weight function generation, `monteswitch` supports all of the methods described in Section 2.5. Note that the option to perform canonical sampling enables conventional Monte Carlo simulations to be performed. Note also that, like all other LSMC implementations the authors are aware of, `monteswitch` explicitly keeps track of two states of the system during a simulation, one corresponding to phase 1 and one corresponding to phase 2. At any point during the simulation only one of these is the ‘actual’ state of the system  $\sigma$ , while the other is the state  $\sigma'$  which would result if a lattice switch were performed from  $\sigma$ . For instance if the system were in phase 1, then  $\sigma$  would belong to phase 1 and  $\sigma'$  would belong to phase 2. During a simulation the actual state  $\sigma$  is evolved via particle moves, and additionally volume moves for the case of the NPT ensemble, in the conventional manner. However the ‘other’ state  $\sigma'$  is evolved contemporaneously with  $\sigma$  so that  $\sigma'$  is always what would result if a lattice switch were performed from  $\sigma$ . It is necessary to continuously track  $\sigma'$  in this manner because the LSMC order parameter for  $\sigma$  depends on the energy, and possibly also the volume, of  $\sigma'$  (see Section 2.4 for details). Regarding lattice switch moves, if such a move is accepted, then the actual state is simply relabelled from  $\sigma$  to  $\sigma'$ , while  $\sigma$  is relabelled as the ‘other’ state.

`monteswitch` also supports on-the-fly evaluation of physical quantities and their uncertainties during the simulation via block averaging (see, e.g., Ref. [1]). It also supports the ability to check whether or not the system has ‘melted’, i.e., whether or not one or more of the particles have strayed ‘too far’ from their lattice sites. In a similar vein, it is possible in `monteswitch` to perform simulations in the centre-of-mass reference frame, which provides a means of suppressing spurious melting due to ‘drift’ in the centre-of-mass of the system during the simulation.

The random number generator utilised by `monteswitch` is the Mersenne Twister (MT19937) [36].

### 5.2. Command-line argument usage

The command-line arguments passed to `monteswitch` determine the nature of the invoked simulation. Usage of `monteswitch` is as follows:

```
monteswitch [-seed <seed>] -new [-wf]
monteswitch [-seed <seed>] (-resume|-reset)
```

The function of each of these arguments is described below.

### 5.3. Seeding the random number generator

The command-line argument `-seed` allows the user to specify the seed for the forthcoming simulation explicitly. If the argument `-seed` is absent then a seed is generated using the system clock.

### 5.4. Running a new simulation

The command-line argument `-new` invokes a new simulation ‘from scratch’. In this case the simulation is initialised using information contained in input files located in the current directory. The input files required by `monteswitch` will depend on the specific version of `interactions.f95` utilised when compiling `monteswitch` (see Section 4). However it is only the input files which contain information pertaining to the interatomic potential which are version-dependent; the remaining information used to initialise a simulation are read from input files which are universal to all versions of `monteswitch`, namely `lattices.in`, `params.in` and `wf.in`. The first two of these are compulsory: they are read by all new simulations. By contrast `wf.in` is optional, only being read if the command-line argument `-wf` is present.

#### 5.4.1. Input file: *lattices\_in*

We now describe these files, beginning with `lattices_in`. This file contains specifications for two states (i.e., supercell dimensions and particle positions), one for each phase. These two states serve two purposes. Firstly, they act as prospective initial state for the simulation: if the system is to be initialised in phase 1, then the phase-1 state will be used as the initial state, and similarly if the system is to be initialised in phase 2. Secondly, they determine the nature of the lattice switch. The fractional position of particle  $i$  for phase  $\alpha$  specified in `lattices_in` is used as the lattice site for  $i$  in phase  $\alpha$  during lattice switches. Furthermore the supercell dimensions specified for each phase in `lattices_in` determine how the supercell is transformed during lattice switches. Specifically, if  $L_x^{(\alpha)}$ ,  $L_y^{(\alpha)}$  and  $L_z^{(\alpha)}$  denote the supercell dimensions specified for phase  $\alpha$  in `lattices_in`, then the  $x$ -,  $y$ - and  $z$ -dimensions of the supercell are multiplied by factors  $L_x^{(2)}/L_x^{(1)}$ ,  $L_y^{(2)}/L_y^{(1)}$  and  $L_z^{(2)}/L_z^{(1)}$  respectively during a lattice switch from phase 1 to phase 2, and by  $L_x^{(1)}/L_x^{(2)}$ ,  $L_y^{(1)}/L_y^{(2)}$  and  $L_z^{(1)}/L_z^{(2)}$  during a lattice switch from phase 2 to phase 1. Below is a pedagogical example of a `lattices_in` file, which corresponds to phase 1 being an 8-particle bcc supercell and phase 2 being an 8-particle hcp supercell, where the phase-1 state consists entirely of particles belonging to species '1' and the phase-2 state consists of a mixture of species '1' and '2'. Note that in this case the species of some of the particles is transformed during lattice switches.

```

bcc-hcp, rho = 0.5, nx,ny,nz = 1, 1, 2 # Comment line
8 # Number of particles
2.2449241 # x-dimension for phase 1
1.5874012 # y-dimension for phase 1
4.4898482 # z-dimension for phase 1
0.0000000 0.0000000 0.0000000 1 # Coords and species for phase 1
0.5000000 0.5000000 0.0000000 1
0.5000000 0.0000000 0.2500000 1
0.0000000 0.5000000 0.2500000 1
0.0000000 0.0000000 0.5000000 1
0.5000000 0.5000000 0.5000000 1
0.5000000 0.0000000 0.7500000 1
0.0000000 0.5000000 0.7500000 1
2.4494897 # x-dimension for phase 2
1.4142136 # y-dimension for phase 2
4.6188021 # z-dimension for phase 2
0.0000000 0.0000000 0.0000000 1 # Coords and species for phase 2
0.5000000 0.5000000 0.0000000 2
0.3333333 0.0000000 0.2500000 1
0.8333333 0.5000000 0.2500000 2
0.0000000 0.0000000 0.5000000 1
0.5000000 0.5000000 0.5000000 2
0.3333333 0.0000000 0.7500000 1
0.8333333 0.5000000 0.7500000 2

```

#### 5.4.2. Input file: *params\_in*

The second compulsory input file for a new simulation is `params_in`. This file contains the variables which determine the nature of the simulation. Each variable corresponds to a specific

single line in the file, and each line must consist of an arbitrary string (we recommend the name of the variable followed immediately by an '=' character with no spaces), followed by whitespace, followed by the value of the variable. To illustrate this, below is an excerpt from a `params.in` file:

```
init_lattice=          1
M_grid_size=          100
M_grid_min=           -82.0
M_grid_max=            48.0
enable_multicanonical= T
beta=                  9.403
P=                     0.0
enable_lattice_moves= T
enable_part_moves=     T
enable_vol_moves=      T
part_select=           "rand"
part_step=              0.3
enable_COM_frame=      T
vol_dynamics=          "UVM"
vol_freq=              1
vol_step=              0.03
stop_sweeps=           160000
equil_sweeps=          0
enable_melt_checks=    T
melt_sweeps=           100
melt_threshold=        3.0
melt_option=           "zero_current"
```

A full list of the variables which must appear in a `params.in`, as well as a description of their function, is provided in the user manual included with the package. The key variables are listed in Table 5.4.2. Numerous examples of `params.in` files are included with *monteswitch* which serve as templates for users.

Table 2: Important control variables for `monteswitch` to be specified in the `params.in` input file.

Variable	Type	Description
<b>init_lattice</b>	INTEGER	Starting phase for the simulation (1 or 2).
<b>M_grid_size</b>	INTEGER	Number of macrostates to divide the considered order parameter range ( <b>M_grid_min</b> to <b>M_grid_max</b> ) into.
<b>M_grid_min</b>	REAL	Minimum of considered order parameter range.
<b>M_grid_max</b>	REAL	Maximum of considered order parameter range.
<b>enable_multicanonical</b>	LOGICAL	T enables multicanonical sampling using the current weight function; F enables canonical sampling.
<b>beta</b>	REAL	Inverse temperature: $\beta = 1/(k_B T)$ .
<b>P</b>	REAL	Pressure (only relevant in NPT ensemble simulations).
<b>enable_lattice_moves</b>	LOGICAL	T enables lattice switch moves (performed after every particle and volume move).
<b>enable_part_moves</b>	LOGICAL	T enables particle moves.
<b>enable_vol_moves</b>	LOGICAL	T enables volume moves and selects the NPT ensemble; F selects the NVT ensemble. A volume move will be attempted on average <b>vol_freq</b> times per sweep.
<b>part_select</b>	CHARACTER(30)	Flag determining how the next particle to move is selected: "cycle" selects particles sequentially, "rand" selects particles at random.
<b>part_step</b>	REAL	Particle move maximum size; particles are moved according to a random walk, with a maximum move size of <b>part_step</b> in any Cartesian direction.
<b>enable_COM_frame</b>	LOGICAL	T performs the simulation in the centre-of-mass reference frame; F uses the lab frame. Using the centre-of-mass frame prevents 'drift' in the centre-of-mass, which is convenient because it keeps particles close to their lattice sites.
<b>vol_dynamics</b>	CHARACTER(30)	Flag determining which type of volume moves are performed: "FVM" (fixed volume move) keeps the supercell shape unchanged during a volume move; "UVM" (unconstrained volume move) allows the x-, y- and z-dimensions to move independently.
<b>vol_freq</b>	INTEGER	Number of volume moves performed per sweep on average if volume moves are enabled. We recommend that this be set to 1.
<b>vol_step</b>	REAL	Volume move maximum step size; the volume is moved according to a random walk in 'ln(V)-space', with a maximum move size of <b>vol_step</b> .
<b>stop_sweeps</b>	INTEGER	Total number of Monte Carlo sweeps to perform in the simulation.

<b>equil_sweeps</b>	INTEGER	Number of sweeps to disregard before the system is considered to be equilibrated; statistics are not gathered during these sweeps for block averaging (see below).
<b>enable_melt_checks</b>	LOGICAL	T enables periodic checks of whether the system has ‘melted’, i.e., if one or more of the particles has moved more than a distance of <b>melt_threshold</b> from its lattice site in any Cartesian direction then the system is considered to have ‘melted’.
<b>melt_sweeps</b>	INTEGER	Period (sweeps) to check for melting.
<b>melt_threshold</b>	REAL	See <b>enable_melt_checks</b> .
<b>melt_option</b>	CHARACTER(30)	Flag determining what the simulation does if the system has ‘melted’: "zero_1" and "zero_2" move the system to the zero-displacement states in phases 1 and 2, respectively; "zero_current" does the same but for the current phase; "stop" stops the simulation. For "zero_1", "zero_2", "zero_current" the system is allowed to re-equilibrate before statistics are gathered for block averaging. Also, the current block is disregarded for the purposes of block averaging.
<b>enable_divergence_checks</b>	LOGICAL	T enables periodic checks of whether the energy of the system is correct, given that during particle moves the energy of the system is <i>amended</i> , as opposed to being calculated from scratch every move. If the energy of the system differs from its ‘true’ energy by an amount <b>divergence_tol</b> then the simulation is stopped.
<b>divergence_sweeps</b>	INTEGER	Period (sweeps) to perform energy checks as just mentioned.
<b>divergence_tol</b>	REAL	See <b>enable_divergence_checks</b> .
<b>output_file_period</b>	INTEGER	Period (sweeps) at which information about the simulation is output to the file <code>data</code> .
<b>output_stdout_period</b>	INTEGER	Period (sweeps) at which information about the simulation is output to <code>stdout</code> .
<b>checkpoint_period</b>	INTEGER	Period (sweeps) at which the simulation is checkpointed, i.e., how often all simulation variables are output to the file <code>state</code> .
<b>update_eta</b>	LOGICAL	T results in the weight function being periodically updated every <b>update_eta_sweeps</b> sweeps, according to the method specified in <b>update_eta_method</b> ; F results in the weight function not being updated – it remains frozen at its current state.
<b>update_eta_sweeps</b>	INTEGER	Period (sweeps) at which the weight function is updated.
<b>update_trans</b>	LOGICAL	T results in the transition matrix being updated; F results in it not being updated.
<b>update_eta_method</b>	CHARACTER(30)	Method used to update the weight function: "VS" uses the visited states method; "shooting" uses the transition matrix method.
<b>enable_barriers</b>	LOGICAL	T enables artificial dynamics; for F the system is free to explore any macrostate, but is constrained to reside within the considered order parameter range ( <b>M_grid_min</b> to <b>M_grid_max</b> ).

<b>barrier_dynamics</b>	CHARACTER(30)	Flag determining how the macrostate barriers will evolve during artificial dynamics. All methods lock the system into a single macrostate for <b>lock_moves</b> moves, before unlocking an adjacent macrostate. Once the system has moved into the adjacent macrostate, the system is then locked into that macrostate, and the procedure starts again. "random" evolves the macrostate the system is locked into via a random walk: the next macrostate is decided with equal probability to be that above or that below the current macrostate. "pong_up" moves to increasingly higher macrostates until the upper limit of the supported order parameter range is encountered, at which point it reverses direction and proceeds to increasingly lower macrostates until it reaches the lower limit of the order parameter range, at which point it reverses direction, etc. "pong_down" instead moves initially to increasingly lower macrostates.
<b>lock_moves</b>	INTEGER	The number of moves to lock the system into one macrostate for if artificial dynamics is used.
<b>calc_equil_properties</b>	LOGICAL	T enables internal calculation of various physical quantities, and associated uncertainties, via block averaging.
<b>block_sweeps</b>	INTEGER	The number of sweeps which comprise a 'block' which will be used to evaluate physical quantities and their associated uncertainties via block averaging.

### 5.4.3. Input file: *wf\_in*

The command-line argument `-wf` allows one to specify the initial weight function to be used in the simulation: if `-wf` is present, then the initial weight function is read from the file `wf_in`. If `-wf` is absent then `wf_in` is not read, and the weight function is initialised to 0 for all macrostates. `wf_in` must contain **M\_grid\_size** lines (where **M\_grid\_size** is specified in `params_in`), each containing two tokens (extra lines and tokens are ignored), which both should be of type REAL. The first token on each line is ignored, while the second tokens are treated as the weight function: the value of the weight function for macrostate  $i$  is initialised to the value of the second token on line  $i$  in `wf_in`. Note that the format of `wf_in` is analogous to that output by the program `monteswitch.post` in conjunction with the `-extract_wf` argument – see Section 6.2.

### 5.5. Simulation output

During a simulation information is periodically output to a file `data` and (optionally) `stdout`. Exactly what information is output is controlled by flags in the input file `params_in`. `data` can be used to deduce how the system evolves with time during the simulation. The format of this file is transparent: each line contains a simulation variable (e.g., energy, volume), followed by the sweep number, followed by the value of the variable. To illustrate this, below is an excerpt from a `data` file:

```
Lx:      250   20.506880155055160      22.375541637220159
Ly:      250   21.983483940969180      19.585057663268277
Lz:      250   20.142474063147095      20.720990682970093
V:       250   9080.4825242547813
lattice:      250           2
E:       250  -2420.3817246101821
M:       250   12.808541584236991
eta:      250   9.9865126253137664
barrier_macro_low: 250           84
Lx:      500   20.470022712854778      22.335325610873863
Ly:      500   21.965084799027686      19.568665891297709
Lz:      500   20.294516764376276      20.877400237511747
V:       500   9124.9380258152232
lattice:      500           2
E:       500  -2409.3469532112986
M:       500   1.1974997526594968
eta:      500   54.025947498280964
barrier_macro_low: 500           84
Lx:      750   20.472567199702674      22.338101960615028
Ly:      750   21.653559219100391      19.291128151472556
Lz:      750   20.503721368631055      21.092613455213190
V:       750   9089.3805950278784
lattice:      750           1
```

In addition to `data`, a file `state` is also created by the program periodically throughout a simulation. This file contains a snapshot of all the simulation variables, and can be used for checkpointing (discussed in a moment), or to extract the ‘results’ of the simulation, e.g., equilibrium quantities, the current weight function, the number of accepted vs. rejected Monte Carlo

moves of a certain type. The format of this file is transparent. Each line pertains to a simulation variable, and the line itself contains the name of the simulation variable followed by the variables value. To illustrate this, below is an except of a state file:

```
output_stdout_sigma_equil_L_2= T
checkpoint_period= 2000
M_grid_size= 100
n_part= 384
Lx= 20.828469299691541 22.726435154004815
Ly= 21.821584954680326 19.440822063915213
Lz= 20.313660823118358 20.897094137158106
V= 9232.7662932888543
lattice= 2
E_1= -2398.4613158910693
E_2= -2417.6552334642806
E= -2417.6552334642820
M= 19.193917573211365
macro= 87
eta= 0.0000000000000000
switchscalex= 1.0911236359717214
switchscaley= 0.89089871814033939
switchscalez= 1.0287212294780348
sweeps= 200
moves= 154044
moves_lattice= 77022
accepted_moves_lattice= 1
moves_part= 76800
```

Table 5.5 provides a list of simulation variables, not already covered by Table 5.4.2, which can be found in state and could be of interest to the user.

Table 3: Useful variables found in the monteswitch output file state.

Variable	Fortran type	Description
<b>n_part</b>	INTEGER	Number of particles in the system.
<b>Lx</b>	REAL(2)	Dimension of supercells in x-direction: the first value pertains to phase 1 while the second pertains to phase 2.
<b>Ly</b>	REAL(2)	Dimension of supercells in y-direction: the first value pertains to phase 1 while the second pertains to phase 2.
<b>Lz</b>	REAL(2)	Dimension of supercells in z-direction: the first value pertains to phase 1 while the second pertains to phase 2.
<b>V</b>	REAL	Current volume of the system.
<b>lattice</b>	INTEGER	Current phase of the system (1 or 2).
<b>E_1</b>	REAL	Energy of phase 1 for the current displacements.
<b>E_2</b>	REAL	Energy of phase 2 for the current displacements.
<b>E</b>	REAL	Current energy of the system. This is <b>E_1</b> if we are in phase 1 and <b>E_2</b> if we are in phase 2.
<b>M</b>	REAL	Current order parameter of the system.
<b>macro</b>	REAL	Current macrostate of the system.
<b>eta</b>	REAL	The value of the weight function associated with the current macrostate of the system.
<b>switchscalex</b>	REAL	Scaling of the supercell in the x-dimension when performing a lattice switch from phase 1 to phase 2. The reciprocal of this is the scaling when performing a lattice switch from phase 2 to phase 1.
<b>switchscaley</b>	REAL	Scaling of the supercell in the y-dimension when performing a lattice switch from phase 1 to phase 2. The reciprocal of this is the scaling when performing a lattice switch from phase 2 to phase 1.
<b>switchscalez</b>	REAL	Scaling of the supercell in the z-dimension when performing a lattice switch from phase 1 to phase 2. The reciprocal of this is the scaling when performing a lattice switch from phase 2 to phase 1.
<b>sweeps</b>	INTEGER	Number of sweeps performed so far, including over previous simulations if we have used the <code>-resume</code> argument.

<b>moves</b>	INTEGER	Total number of moves performed so far in total, including over previous simulations if we have used the <code>-resume</code> argument.
<b>moves_lattice</b>	INTEGER	Number of lattice moves performed so far, including over previous simulations if we have used the <code>-resume</code> argument.
<b>accepted_moves_lattice</b>	INTEGER	Number of accepted lattice moves so far, including over previous simulations if we have used the <code>-resume</code> argument.
<b>moves_part</b>	INTEGER	Number of particle moves performed so far, including over previous simulations if we have used the <code>-resume</code> argument.
<b>accepted_moves_part</b>	INTEGER	Number of accepted particle moves so far, including over previous simulations if we have used the <code>-resume</code> argument.
<b>moves_vol</b>	INTEGER	Number of volume moves performed so far, including over previous simulations if we have used the <code>-resume</code> argument.
<b>accepted_moves_vol</b>	INTEGER	Number of accepted volume moves so far, including over previous simulations if we have used the <code>-resume</code> argument.
<b>melts</b>	INTEGER	The number of times the system has melted.
<b>barrier_macro_low</b>	INTEGER	The macrostate number corresponding to the lowest currently allowed macrostate (relevant only when artificial dynamics is enabled).
<b>barrier_macro_high</b>	INTEGER	The macrostate number corresponding to the highest currently allowed macrostate (relevant only when artificial dynamics is enabled).
<b>block_counts</b>	INTEGER	The total number of blocks considered so far for block averaging.
<b>equil_DeltaF</b>	REAL	The free energy difference between the phases ( $F_1 - F_2$ ; extensive) evaluated using block averaging.
<b>sigma_equil_DeltaF</b>	REAL	The uncertainty in <b>equil_DeltaF</b> evaluated using block averaging..
<b>equil_H_1</b>	REAL	The energy (for NVT simulations) or enthalpy (for NPT simulations) of phase 1 evaluated using block averaging.
<b>equil_H_2</b>	REAL	The energy (for NVT simulations) or enthalpy (for NPT simulations) of phase 2 evaluated using block averaging.
<b>sigma_equil_H_1</b>	REAL	The uncertainty in <b>equil_H_1</b> .
<b>sigma_equil_H_2</b>	REAL	The uncertainty in <b>equil_H_2</b> .
<b>equil_V_1</b>	REAL	The volume of phase 1 evaluated using block averaging.
<b>equil_V_2</b>	REAL	The volume of phase 2 evaluated using block averaging.
<b>sigma_equil_V_1</b>	REAL	The uncertainty in <b>equil_V_1</b> .

<b>sigma_equil_V_2</b>	REAL	The uncertainty in <b>equil_V_2</b> .
<b>R_1</b>	REAL(n_part,3)	The current lattice vectors for phase 1.
<b>R_2</b>	REAL(n_part,3)	The current lattice vectors for phase 2.
<b>u</b>	REAL(n_part,3)	The current displacement vectors.
<b>M_grid</b>	REAL(M_grid_size)	Array containing the minimum order parameter for each macrostate: macrostate $n$ corresponds to order parameters between <b>M_grid</b> ( $n$ ) and <b>M_grid</b> ( $n + 1$ ).
<b>M_counts_1</b>	INTEGER(M_grid_size)	<b>M_counts_1</b> ( $n$ ) is the number of times macrostate $n$ has been visited while the system was in phase 1 so far, including over previous simulations if we have used the <code>-resume</code> argument.
<b>M_counts_2</b>	INTEGER(M_grid_size)	<b>M_counts_2</b> ( $n$ ) is the number of times macrostate $n$ has been visited while the system was in phase 2 so far, including over previous simulations if we have used the <code>-resume</code> argument.
<b>eta_grid</b>	REAL(M_grid_size)	<b>eta_grid</b> ( $n$ ) is the value of the weight function for macrostate $n$ .
<b>trans</b>	REAL(M_grid_size,M_grid_size)	<b>trans</b> ( $m, n$ ) is the number of inferred transitions from macrostate $m$ to macrostate $n$ ; it is the matrix $C_{MM'}$ in Section 2.5.2.
<b>equil_umsd_1</b>	REAL(n_part)	<b>equil_umsd_1</b> ( $n$ ) is the mean-squared displacement of particle $n$ from its lattice site in phase 1, evaluated using block averaging.
<b>equil_umsd_2</b>	REAL(n_part)	<b>equil_umsd_2</b> ( $n$ ) is the mean-squared displacement of particle $n$ from its lattice site in phase 2, evaluated using block averaging.
<b>sigma_equil_umsd_1</b>	REAL(n_part)	<b>sigma_equil_umsd_1</b> ( $n$ ) is the uncertainty in <b>equil_umsd_1</b> ( $n$ ).
<b>sigma_equil_umsd_2</b>	REAL(n_part)	<b>sigma_equil_umsd_2</b> ( $n$ ) is the uncertainty in <b>equil_umsd_2</b> ( $n$ ).
<b>equil_L_1</b>	REAL(3)	The 1st, 2nd and 3rd values in the array <b>equil_L_1</b> are the mean x-, y- and z-dimensions of the supercell in phase 1, evaluated using block averaging.
<b>equil_L_2</b>	REAL(3)	The 1st, 2nd and 3rd values in the array <b>equil_L_2</b> are the mean x-, y- and z-dimensions of the supercell in phase 2, evaluated using block averaging.
<b>sigma_equil_L_1</b>	REAL(3)	<b>sigma_equil_L_1</b> ( $n$ ) is the uncertainty in <b>equil_L_1</b> ( $n$ ).
<b>sigma_equil_L_2</b>	REAL(3)	<b>sigma_equil_L_2</b> ( $n$ ) is the uncertainty in <b>equil_L_2</b> ( $n$ ).

### 5.6. Resuming a checkpointed simulation

The command-line argument `-resume` continues an ‘old’ simulation, whose variables are contained in the file `state` in the current directory. The ‘resumed’ simulation is run for the number of Monte Carlo sweeps specified in the variable `stop_sweeps` in `state`. By default this is the number of sweeps which were performed in the old simulation, though one of course this can be manually altered if one wants the resumed simulation to be of a different length to the old simulation. For a simulation invoked using the argument `-resume`, the file `data` is amended: the resumed simulation does not overwrite the data file; all information from the old simulation is retained in it.

The command-line argument `-reset` invokes a simulation from an old `state` file similarly to `-resume`, except that it resets all ‘counter variables’ to zero. This has the effect of starting a ‘new’ simulation whose nature corresponds to the old simulation, but instead uses the state of the system specified in `state`. By contrast, the argument `-new` initialises the state to be such that the particles form a perfect crystal lattice, which usually does not correspond to an equilibrated state. By ‘counter variables’ we mean those such as variables describing the number of moves performed for each move type, the number of accepted moves for each move type, and variables pertaining to equilibrium quantities. For a simulation invoked using the argument `-reset`, the file `data` is overwritten, i.e., the information from the ‘old’ simulation is not retained.

### 5.7. MPI simulations

As mentioned at the beginning of this chapter, the program `monteswitch_mpi` is the MPI-parallelised analogue of `monteswitch`. To build upon what was already said, `monteswitch_mpi` is identical to the program, except that instead of a single simulation for `stop_sweeps` Monte Carlo sweeps, `monteswitch_mpi` runs  $n$  simulations – replicas – in parallel using MPI, each being approximately `stop_sweeps/n` sweeps in length. Accordingly, during the simulation multiple `data-` and `state-`format files are created – one for each replica. These are named `state_0`, `state_1`, `state_2`, etc., and similarly for the `data-`format files. At the completion of all replicas, the results of block averaging from all replicas are combined and stored in the file `state`. The state stored in `state` corresponds to the state in `state_0`.

Further details of the parallelisation are as follows. All replicas are always initialised to be in the same state. For a new simulation this is determined by the `init_lattice` variable in `params.in` similarly to `monteswitch`. For a resumed simulation this is the state contained in the `state` file from which the simulation is to be resumed. We emphasise that *all replicas of the system are initialised with the same state* when the `-resume` argument is used with `monteswitch_mpi`. The files `data_n` are therefore always overwritten by `monteswitch_mpi` when the `-resume` argument is invoked since, given the nature of the parallelisation, there is no continuity between the replicas in subsequent simulations. The exception is replica ‘0’, whose state is always stored in the file `state`, as well as `state_0`.

## 6. Utility programs

As mentioned earlier, `monteswitch` contains a number of utility programs which assist with post-processing of the data and the construction of input files. We now describe these programs.

### 6.1. Programs for generating *lattices\_in* files

The programs `lattices_in_hcp_fcc`, `lattices_in_bcc_fcc` and `lattices_in_bcc_hcp` create `lattices_in` files containing reference states corresponding to, respectively, hcp-fcc, bcc-fcc and bcc-hcp lattice switches. Usage of these programs is as follows:

```
lattices_in_hcp_fcc <rho> <nx> <ny> <nz>
lattices_in_bcc_fcc <rho> <nx> <ny> <nz>
lattices_in_bcc_hcp <rho> <nx> <ny> <nz>
```

The command-line arguments `<rho>`, `<nx>`, `<ny>` and `<nz>` constitute the free parameters for the `lattices_in` file to be created. The first argument `<rho>` is the density (i.e., the number of particles per unit volume) of the reference states to construct. (Note that both states necessarily have the same density). The second, third and fourth arguments `<nx>`, `<ny>` and `<nz>` are integers which correspond to the number of unit cells (described in a moment) which will be tiled in the x-, y- and z-directions respectively to construct the supercell for each phase. The programs output the desired `lattices_in` file to stdout. Hence one must redirect the output to create the required `lattices_in` file, e.g., `$ ./lattices_in_hcp_fcc 0.5 2 3 5 > lattices_in`

What follows is a description of the unit cells for each pair of phases for each program.

#### 6.1.1. *lattices\_in\_hcp\_fcc*

The unit cell here contains 12 particles. The particles are spread over 6 planes in the z-direction; each plane contains two particles. The positions of the particles corresponds to a stacking sequence for the planes of ABCABC for the fcc unit cell, and ABABAB for the hcp unit cell.

#### 6.1.2. *lattices\_in\_bcc\_fcc*

The unit cell here is the conventional 2-particle body-centred tetragonal (bct) unit cell; for the bcc(fcc) lattice the relative dimensions of the bct unit cell in each Cartesian direction correspond to the bct representation of the bcc(fcc) lattice.

#### 6.1.3. *lattices\_in\_bcc\_hcp*

The unit cell here contains 4 particles. The bcc unit cell is the 4-particle face-centred tetragonal (fct) corresponding to the fct representation of the bcc lattice. The hcp unit cell is the 'fct-like' representation of the hcp lattice.

### 6.2. Post-processing *state* files

`monteswitch_post` is a tool for post-processing the file `state` generated by `monteswitch` or `monteswitch_mpi`. It can be used to extract useful information from that file. The `state` file which the program operates on is that in the current directory. The command-line arguments determine the task performed by the program. Usage of `monteswitch_post` is as follows, where the function of each command-line argument is described below:

```
monteswitch_post -extract_wf
monteswitch_post -extract_M_counts
monteswitch_post -extract_pos [<species>]
monteswitch_post -extract_R_1 [<species>]
monteswitch_post -extract_R_2 [<species>]
```

```
monteswitch_post -extract_u [<species> <phase>]
monteswitch_post -calc_rad_dist <bins>
monteswitch_post -merge_trans <state_in_1> <state_in_2> <state_out>
monteswitch_post -extract_lattices_in <vectors_in_1> <vectors_in_2>
monteswitch_post -extract_pos_xyz
monteswitch_post -set_wf <wf_file>
monteswitch_post -taper_wf
```

#### 6.2.1. *-extract\_wf*

Extract the weight function from `state` and output it to `stdout`. In the output the first token on each line is the order parameter, and the second is the corresponding value of the weight function.

#### 6.2.2. *-extract\_M\_counts*

Extract order parameter histograms from `state` and output them to `stdout`. In the output the first token on each line is the order parameter, the second is the corresponding number of counts for phase 1, and the third is the corresponding number of counts for phase 2.

#### 6.2.3. *-extract\_pos [<species>]*

Extract the current positions of the particles, and output them to `stdout`. In the output the first, second and third tokens on each line are the x-, y- and z-coordinates respectively for a particle. If the optional argument `<species>` is present then only the positions for particles belonging to species `<species>` are output.

#### 6.2.4. *-extract\_R\_1 [<species>]*

Extract the current positions of the lattice sites for phase 1, and output them to `stdout`. In the output the first, second and third tokens on each line are the x-, y- and z-coordinates respectively for a particle. If the optional argument `<species>` is present then only the sites for particles belonging to species `<species>` in phase 1 are output.

#### 6.2.5. *-extract\_R\_2 [<species>]*

Extract the current positions of the lattice sites for phase 2, and output them to `stdout`. In the output the first, second and third tokens on each line are the x-, y- and z-coordinates respectively for a particle. If the optional argument `<species>` is present then only the sites for particles belonging to species `<species>` in phase 2 are output.

#### 6.2.6. *-extract\_u [<species> <phase>]*

Extract the displacements of the particles, and output them to `stdout`. In the output the first, second and third tokens on each line are the x-, y- and z-displacements respectively for a particle. If the optional arguments `<species>` and `<phase>` are present then only the displacements for particles belonging to species `<species>` in phase `<phase>` are output.

#### 6.2.7. *-calc\_rad\_dist <bins>*

Calculate the radial distribution function, based on the current state of the system, and output it to `stdout`. `<bins>` is the number of bins to be used in the function.

#### 6.2.8. `-merge_trans <state_in_1> <state_in_2> <state_out>`

Combine the **trans** matrices from the files `<state_in_1>` and `<state_in_2>`, and store the result in the file `<state_out>`, where all variables in `<state_out>` other than the matrix **trans** are inherited from `<state_in_1>`. This argument can be used for pooling the results of multiple simulations which utilise the same underlying ‘order parameter grid’ **M\_grid\_size**.

#### 6.2.9. `-extract_lattices_in <vectors_in_1> <vectors_in_2>`

Output the geometrical properties of the system in the format of a `lattices_in` file. The arguments `<vectors_in_1>` and `<vectors_in_2>` can be either `pos` or `R`. If `<vectors_in_1>` is `pos`, then the positions of the particles in phase 1 of the forthcoming `lattices_in` file will be the positions of the particles in phase 1 in the `state` file; if `<vectors_in_1>` is `R`, then the positions of the particles in phase 1 of the `lattices_in` file will be the current lattice vectors (i.e., **R\_1**) corresponding to phase 1 in the `state` file. Similar applies for `<vectors_in_2>` with phase 2.

#### 6.2.10. `-extract_pos_xyz`

Extract the positions of the particles, and output them to stdout in ‘.xyz’ format. In the output the first line contains the number of particles, the second line is a comment line, and the subsequent lines contain the particle positions and species: the first token is the ‘element’ (set to ‘A’ for species ‘1’, ‘B’ for species ‘2’, ..., ‘Z’ for species 26, and ‘?’ otherwise), and the second, third, fourth and fifth tokens are the x-, y- and z-coordinates respectively.

#### 6.2.11. `-set_wf <wf_file>`

Alters the weight function in `state` to correspond to that specified in the file `<wf_file>`. The format of the file `<wf_file>` must be analogous to the format of the weight function output by this program via the `-extract_wf` argument described above: the file must contain **M\_grid\_size** lines, each containing two tokens (extra lines and tokens are ignored), which both should be of type REAL. The first token on each line is ignored, while the second tokens are treated as the new weight function: the value of the weight function for macrostate  $i$  is set to the value of the second token on line  $i$  in `<wf_file>`.

#### 6.2.12. `-taper_wf`

‘Tapers’ the weight function in the `state` file. To elaborate, it is assumed that the weight function has a single local minimum in the  $M < 0$  region (the region associated with phase 1) and a single local minimum in the  $M > 0$  region (associated with phase 2), with a maximum at  $M \approx 0$  separating the two minima. Let  $\eta_{\min,1}$  denote the value of the weight function at the minimum in the  $M < 0$  region, and let  $M_{\min,1}$  denote the location of this minimum. In the  $M < 0$  region the weight function is tapered by setting  $\eta_M = \eta_{\min,1}$  for all  $M < M_{\min,1}$ . Similarly the weight function is tapered in the  $M > 0$  region by setting  $\eta_M = \eta_{\min,2}$  for all  $M > M_{\min,2}$ . Thus the weight function is ‘flattened’ for the regions in  $M$ -space ‘outwith’ the two minima. This prevents oversampling of these regions, which is unnecessary and inefficient.

## 7. Examples

We now present results obtained using *monteswitch* for two systems in order to elucidate how *monteswitch* could be used in practice.

### 7.1. Hard-sphere solid: hcp vs. fcc

The first system we consider is the hard-sphere solid. This system has been studied extensively with LSMC [2, 5, 3, 13], and accordingly makes an excellent testing ground for *monteswitch*. LSMC studies of this system have focused on calculating the free energy difference between the hcp and fcc phases. We have done the same using *monteswitch* for a 216-particle system, with the aim of validating *monteswitch* against the results of other studies. Specifically we consider the NVT ensemble at reduced density  $\tilde{\rho} = 0.7778$ , where  $\tilde{\rho}$  is related to the number density of the system  $\rho$  by the equation  $\tilde{\rho} = \rho/\sqrt{2}$ ; and the NPT ensemble at  $P\beta\sigma^3 = 14.58$ , where  $\sigma$  denotes the hard sphere diameter. For the NPT ensemble we consider both *anisotropic* volume moves which allow the Cartesian dimensions of the supercell to vary independently, and *isotropic* volume moves in which the supercell’s shape, but not size, is constrained. Note that for each of the three ensembles we considered, i.e., the NVT ensemble, the isotropic NPT ensemble and the anisotropic NPT ensemble, we used the utility program `lattices_in_hcp_fcc` (see Section 6) to generate the relevant `lattices_in` input file required by the main programs `monteswitch` and `monteswitch_mpi` (Section 5).

For each of the three ensembles we used the following procedure to obtain the hcp–fcc free energy difference, and, in the case of the NPT ensembles, the densities of each phase  $\tilde{\rho}_{\text{hcp}}$  and  $\tilde{\rho}_{\text{fcc}}$ . We first performed a number of short conventional Monte Carlo simulations in each phase. The aim of these was to determine quantities to be used in the forthcoming LSMC simulations, such as the maximum particle and volume move sizes, and the appropriate range of order-parameter space to consider. Then we performed an LSMC simulation in order to generate the weight function. In the weight-function-generation simulations we used artificial dynamics to quickly generate an accurate weight function. Moreover we used `monteswitch_mpi` to parallelise the simulation: we ran 16 replicas of the system in parallel. Our weight-function-generation simulations each consisted of 18,000,000 Monte Carlo sweeps (1,125,000 sweeps per replica). Fig. 2 shows the weight function obtained from a weight-function-generation simulation for the NVT ensemble, along with the evolution of the order parameter for one of the replicas during that simulation. Note that order-parameter space is explored systematically, concordant with artificial dynamics as described in Section 2.5.3. Using this weight function we then we performed ‘production’ LSMC simulations to determine the hcp–fcc free energy difference – and also  $\tilde{\rho}_{\text{hcp}}$  and  $\tilde{\rho}_{\text{fcc}}$  for the NPT systems. For each ensemble we performed two production simulations, where each production simulation consisted of 125,000,000 sweeps, again using 16 replicas in parallel via `monteswitch_mpi` (7,812,500 sweeps per replica). Thus the total number of production sweeps for each ensemble was 250,000,000. Finally, the results from all production replicas for each ensemble were pooled to obtain the final results.

The final results are presented in Table 7.1 along with the results of other studies. As can be seen from the table *monteswitch* is in agreement with the other studies. Note that slightly different NPT results are obtained using isotropic and anisotropic volume moves. This is expected since isotropic moves constrain the underlying lattice for each phase to be ‘strictly’ hcp or fcc, while anisotropic moves allow the system to ‘stretch’ along any Cartesian dimension. Interestingly the densities of both phases are higher when isotropic moves are used.

### 7.2. Embedded atom model for Zr: bcc vs. hcp

As our second example we consider the hcp–bcc transition in Zr, modeled using the EAM potential ‘#2’ developed in Ref. [37]. In Ref. [37] the authors determined that the zero-pressure

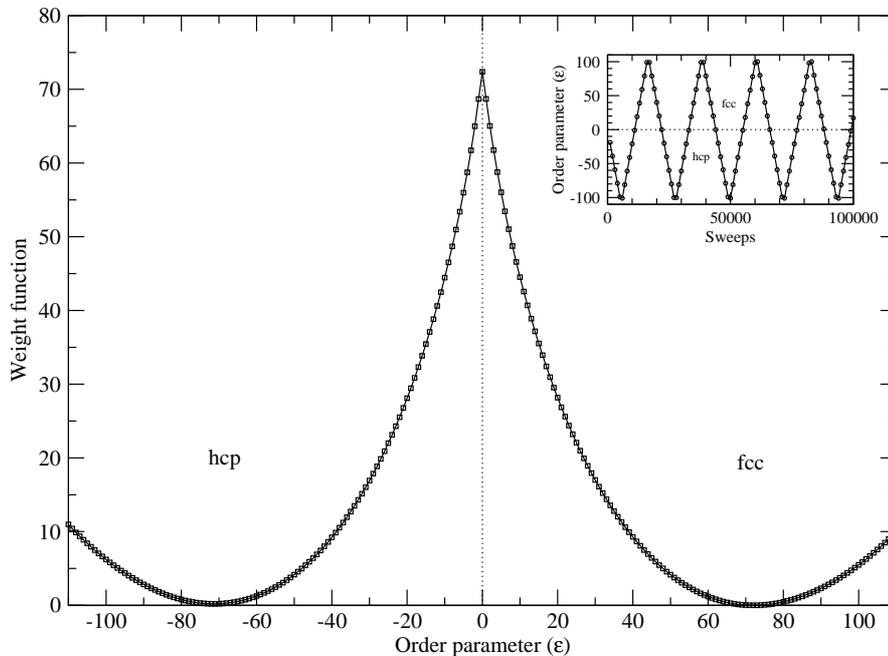


Figure 2: Weight function for an LSMC simulation of a 216-particle hard-sphere system in the NVT ensemble at  $\tilde{\rho} = 0.7778$ ; and the evolution of the order parameter vs. simulation time for one of the replicas during the corresponding weight-function-generation simulation (inset; only the first 100,000 sweeps are shown). The regions of order-parameter space corresponding to hcp and fcc are indicated. The order parameter is in units of the sphere overlap energy  $\epsilon$ , which is set sufficiently high during the simulation to realise hard spheres (see Section 4.2.4).

System	Study	$\beta\Delta\mathcal{F}_{\text{fcc}\rightarrow\text{hcp}}$	$\tilde{\rho}_{\text{hcp}}$	$\tilde{\rho}_{\text{fcc}}$
NVT, $\tilde{\rho} = 0.7778$	<i>This work</i>	0.00135(5)	-	-
	Ref. [5]	0.00132(4)	-	-
	Ref. [3]	0.00133(4)	-	-
	Ref. [13]	0.00133(3)	-	-
NPT, $P\beta\sigma^3 = 14.58$ , iso	<i>This work</i>	0.00123(6)	0.77820(6)	0.77820(6)
NPT, $P\beta\sigma^3 = 14.58$ , aniso	<i>This work</i>	0.00117(9)	0.77759(6)	0.77768(6)
	Ref. [3]	0.00113(4)	0.7776(1)	0.7775(1)

Table 4: LSMC results for various 216-particle hard-sphere systems.  $\Delta\mathcal{F}_{\text{fcc}\rightarrow\text{hcp}}$  is the *intensive* Helmholtz(Gibbs) free energy difference in the NVT(NPT) ensemble. ‘iso’ refers to isotropic volume moves, while ‘aniso’ refers to anisotropic volume moves in which the shape of the supercell is allowed to vary independently in each Cartesian dimension. Uncertainties in the LSMC quantities of this work are standard errors of the mean.

transition temperature  $T_{\text{hcp}\rightarrow\text{bcc}}$  for this potential was 1233K. They also determined that the enthalpy change  $\Delta H_{\text{hcp}\rightarrow\text{bcc}}$  and fractional volume change  $\Delta V_{\text{hcp}\rightarrow\text{bcc}}/V_{\text{hcp}}$  associated with the transition were 0.039 eV/atom and -0.8% respectively. However these quantities were determined *indirectly* from the results of multiple molecular dynamics simulations (see Ref. [37] for details). By contrast we have used *monteswitch* to determine  $T_{\text{hcp}\rightarrow\text{bcc}}$ ,  $\Delta H_{\text{hcp}\rightarrow\text{bcc}}$  and  $\Delta V_{\text{hcp}\rightarrow\text{bcc}}/V_{\text{hcp}}$  *directly* via LSMC. Note that this is a more ‘realistic’ example than the previous one: usually one is interested in, e.g., the temperature and pressure at which a phase transition between two phases occurs, as opposed to the value of the free energy difference between the phases at one state point (which was the nature of the previous example).

While LSMC provides a means for calculating the Gibbs free energy difference  $\Delta G$  between two phases at a given temperature and pressure, there is the question of how to exploit LSMC to determine the transition temperature, which by definition is the temperature at which  $\Delta G = 0$ . We used the following iterative procedure, which in Ref. [4] was shown, for the Lennard-Jones solid, to outperform other existing methods. We first considered an initial estimate for the transition temperature  $T^{(1)}$ . We then used LSMC to obtain the Gibbs free energy difference between the phases  $\Delta G \equiv G_{\text{bcc}} - G_{\text{hcp}}$ , the enthalpies and volumes for each phase ( $H_{\text{hcp}}$ ,  $H_{\text{bcc}}$ ,  $V_{\text{hcp}}$  and  $V_{\text{bcc}}$ ), as well as uncertainties in these quantities, at  $T^{(1)}$ . We then substituted these quantities into the following equation to generate a more accurate estimate for the transition temperature  $T^{(2)}$ :

$$T^{(n+1)} = T^{(n)} \left[ 1 - \left( 1 - \frac{\Delta H^{(n)}}{\Delta G^{(n)}} \right) \right], \quad (18)$$

where  $\Delta H \equiv H_{\text{bcc}} - H_{\text{hcp}}$ , and the superscript ‘ $(n)$ ’ signifies a quantity obtained from an LSMC simulation at temperature  $T^{(n)}$ . (The above equation is, in fact, the Newton-Raphson estimate of the temperature  $T^{(n+1)}$  at which  $\Delta G^{(n+1)} = 0$ ; see Ref. [33] for details). We then repeated all of the above for  $T^{(2)}$ ,  $T^{(3)}$ , etc. until the procedure converged. The procedure was deemed to have converged if the change in temperature for the next iteration,  $T^{(n+1)} - T^{(n)}$ , was less than its corresponding uncertainty – calculated from propagating the uncertainties in  $\Delta H^{(n)}$  and  $\Delta G^{(n)}$  through Eqn. (18). At this point the next iteration was not performed, and the enthalpies and volumes obtained from iteration  $n$  were used to form our quoted values for  $\Delta H_{\text{hcp}\rightarrow\text{bcc}}$  and  $\Delta V_{\text{hcp}\rightarrow\text{bcc}}/V_{\text{hcp}}$ ; and  $T^{(n+1)}$  was used as our quoted value for  $T_{\text{hcp}\rightarrow\text{bcc}}$ .

We considered two system sizes: a smaller system containing 384 atoms and a larger system containing 1296 atoms. Note that both these systems are significantly smaller than those used in Ref. [37]. For both system sizes we used the NPT ensemble with anisotropic volume moves (in the sense described in the previous example), and bootstrapped the iterative procedure with  $T^{(1)}=900\text{K}$ . Furthermore we used the utility program `lattices_in_bcc_hcp` to generate the relevant `lattices_in` input files for all simulation. Our procedure for calculating  $\Delta G$ ,  $H_{\text{hcp}}$ ,  $H_{\text{bcc}}$ ,  $V_{\text{hcp}}$  and  $V_{\text{bcc}}$  at a given temperature is similar to that described in the previous example: we first performed short conventional Monte Carlo simulations in each phase, then a weight-function-generation simulation, and finally a production simulation. Noteworthy aspects of the procedure are as follows. Firstly, our weight-function-generation simulations consisted of 160,000 Monte Carlo sweeps using 4 replicas in parallel via `monteswitch_mpi` (40,000 sweeps per replica). Secondly, we ‘tapered’ the weight function obtained from a weight-function-generation simulation using the utility program `monteswitch_post` before using the weight function in a production simulation. Fig. 3 shows the weight function obtained for the 384-atom system at  $T = 1275.86\text{K}$  (the final iteration,  $n = 3$ , as is discussed below) before and after tapering. Tapering was done in order to improve the efficiency of the production simulation (see Section

	$n$	$T^{(n)}$	$\Delta G_{\text{hcp}\rightarrow\text{bcc}}$	$\Delta H_{\text{hcp}\rightarrow\text{bcc}}$	$\Delta V_{\text{hcp}\rightarrow\text{bcc}}/V_{\text{hcp}}$	$T_{\text{hcp}\rightarrow\text{bcc}}$
$N = 384$	1	900.00	0.0114(11)	0.0459(4)	-0.8(2)	1200(40)
	2	1199.91	0.00238(3)	0.0401(2)	-0.810(6)	1275.9(11)
	3	1275.86	-0.00003(5)	<b>0.0379(2)</b>	<b>-0.787(7)</b>	<b>1275(2)</b>
$N = 1296$	1	900.00	0.01265(2)	0.04589(7)	-0.996(3)	1242.6(11)
	2	1242.61	0.001031(10)	0.03930(6)	-0.791(7)	1276.1(3)
	3	1276.10	-0.00003(2)	0.03835(13)	-0.772(5)	1275.1(6)
	4	1275.13	-0.000000(10)	<b>0.0383(2)</b>	<b>-0.772(7)</b>	<b>1275.1(3)</b>
Ref. [37]	-	-	-	0.039	-0.8	1233

Table 5: Results of LSMC simulations for Zr potential ‘#2’ of Ref. [37]. Results are given for each iteration  $n$  of the iterative procedure described in the main text, for each system size considered ( $N$  denotes the number of atoms in the system). The results of the final iteration, which are to be compared to the results of Ref. [37] (bottom row), are in bold text. Note that for the LSMC results  $T_{\text{hcp}\rightarrow\text{bcc}}$  is the transition temperature as predicted from the results of the corresponding LSMC simulation, whereas  $T^{(n)}$  is the temperature at which that simulation was performed. All temperatures are in K,  $\Delta G_{\text{hcp}\rightarrow\text{bcc}}$  and  $\Delta H_{\text{hcp}\rightarrow\text{bcc}}$  are in eV/atom, and  $\Delta V_{\text{hcp}\rightarrow\text{bcc}}/V_{\text{hcp}}$  is a percentage. Uncertainties in the LSMC quantities are standard errors of the mean.

6.2.12 for details). Finally, at each temperature we performed a single production simulation consisting of 700,000 sweeps, again using 4 replicas in parallel (175,000 sweeps per replica).

Our results are presented in Table 7.2, where they are compared against those of Ref. [37]. Note that our 384-atom and 1296-atom results are in agreement with each other, which means that finite-size effects are insignificant in the 384-atom system for the quantities we consider here. Note also that the iterative procedure converged quickly for both system sizes: only a few iterations were needed to pinpoint the transition temperature to an accuracy of  $\lesssim 2\text{K}$ . We emphasise that only relatively modest computational effort was required to obtain this accuracy: for the 384-atom system each iteration (i.e., one 160,000-sweep weight-function-generation simulation and one 700,000-sweep production simulation) took a wall-clock time of  $\approx 1.9$  hours on a desktop machine (exploiting 4 cores),<sup>9</sup> while for the 1296-atom system each iteration took  $\approx 17.7$  hours. Our results are in agreement with those of Ref. [37] for  $\Delta H_{\text{hcp}\rightarrow\text{bcc}}$  and  $\Delta V_{\text{hcp}\rightarrow\text{bcc}}/V_{\text{hcp}}$ . Regarding  $T_{\text{hcp}\rightarrow\text{bcc}}$ , we find  $T_{\text{hcp}\rightarrow\text{bcc}}$  to be 42K higher than the 1233K quoted in Ref. [37]. Whether this is sufficiently close to 1233K to constitute ‘agreement’ depends on the uncertainty in the Ref. [37] value, which unfortunately is not provided in that study. However in Ref. [38] the same method as in Ref. [37] was used to determine  $T_{\text{hcp}\rightarrow\text{bcc}}$  for a Mg EAM potential; moreover the uncertainty in the value was stated to be 40K. (Specifically,  $T_{\text{hcp}\rightarrow\text{bcc}}$  was found to be  $645\pm 40\text{K}$  for the Mg potential). If this uncertainty carries over to the Zr study of Ref. [37], then it follows that our  $T_{\text{hcp}\rightarrow\text{bcc}}$  is in agreement with that of Ref. [37].

## 8. Conclusions and outlook

We have described *monteswitch*, a package for performing lattice-switch Monte Carlo (LSMC) simulations for atomic systems, and have presented results demonstrating its efficacy. While here we have only presented results for single-component crystalline phases, we emphasise that *monteswitch* can be applied more generally. *monteswitch* could be used to calculate

<sup>9</sup>Specifically, the machine we used was an iMac14,2 with a 3.2GHz Intel Core i5-4570 processor.

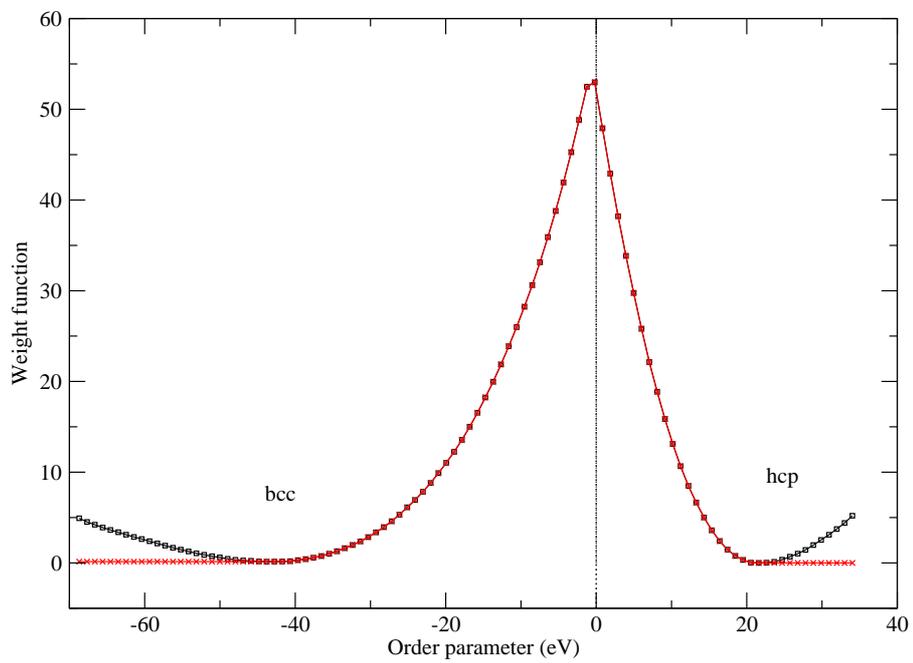


Figure 3: Weight function for Zr potential '#2' of Ref. [37] applied to a 384-atom system at  $T = 1275.86\text{K}$ . The black curve with squares corresponds to the 'raw' weight function obtained from the weight-function-generation simulation, while the red curve with crosses corresponds to the weight function after 'tapering'. The regions of order-parameter space corresponding to bcc and hcp are indicated.

the free energy difference between two multicomponent phases. Moreover *monteswitch* could be used to evaluate the free energy of an interface between two solids, or the free energy of a crystallographic defect. The free energy of an interface between two phases  $A$  and  $B$  can be calculated via LSMC by choosing the two supercells used in the LSMC simulation to have the same amounts of  $A$  and  $B$ , but different amounts of  $A$ – $B$  interface [5]. For example if the first supercell were comprised of superimposed  $A$  and  $B$  slabs with a stacking sequence  $A|B|A|B$ , and similarly for the second supercell but with a stacking sequence  $A|A|B|B$ , then, while both supercells contain two  $A$  and  $B$  slabs, the first supercell contains four  $A$ – $B$  interfaces (taking into account periodic boundary conditions) but the second contains only two interfaces. Denoting the area associated with one interface in the supercell as  $\mathcal{A}$ , it follows that calculating the free energy difference between these supercells via LSMC would yield the free energy difference associated with an area  $2\mathcal{A}$  of an  $A$ – $B$  interface – assuming that the slabs are sufficiently large that interfaces do not ‘interact’ with one another. A similar approach could conceivably be applied to calculate the free energy of planar defects such as twin boundaries. In principle point defects are also accessible to LSMC. For example the free energy of a Frenkel defect could be evaluated by having one supercell be the ‘defect-free’ crystal, while the other contains a single Frenkel defect. The prospect of using LSMC in this way requires further exploration.

The main strength of *monteswitch* is its versatility regarding the interatomic potentials it can implement, which we have demonstrated here by using *monteswitch* in conjunction with an embedded atom model (which is a many-body potential), and the hard-sphere model. This versatility is achieved by having the source code for the potentials housed within a Fortran module which is amenable to customisation. As well as using modules included with the package which implement some commonly-used potentials, it is anticipated that users will wish to write their own versions of this module to implement their own potentials. Templates and guidance are provided with the package to facilitate this. An especially interesting prospect is to develop modules which interface *monteswitch* with quantum chemistry programs, in order to calculate the energy using *ab initio* methods. Such ‘*ab initio* LSMC’ would be a valuable tool in examining the phase stability of systems in which classical models are inappropriate. Needless to say any new modules we develop will be made available to the wider community on the home page for the package [39].

## Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council [a Doctoral Prize Fellowship; and grant EP/M011291/1]. Valuable discussions with Mikhail Mendeleev, Nigel Wilding, Andrey Brukhno and Kevin Stratford are gratefully acknowledged. This work made use of the Balena High Performance Computing Service at the University of Bath.

- [1] D. Frenkel, B. Smit, *Understanding Molecular Simulation*, 2nd Edition, Academic Press, Inc., 2001.
- [2] A. D. Bruce, N. B. Wilding, G. J. Ackland, Free energy of crystalline solids: A lattice-switch monte carlo method, *Phys. Rev. Lett.* 79 (1997) 3002–3005.
- [3] A. D. Bruce, A. N. Jackson, G. J. Ackland, N. B. Wilding, Lattice-switch monte carlo method, *Phys. Rev. E* 61 (2000) 906–919.
- [4] A. N. Jackson, A. D. Bruce, G. J. Ackland, Lattice-switch monte carlo method: Application to soft potentials, *Phys. Rev. E* 65 (2002) 036710.
- [5] S. Pronk, D. Frenkel, Can stacking faults in hard-sphere crystals anneal out spontaneously?, *J. Chem. Phys.* 110 (9) (1999) 4589–4592.
- [6] S.-C. Mau, D. A. Huse, Stacking entropy of hard-sphere crystals, *Phys. Rev. E* 59 (1999) 4396–4401.
- [7] A. N. Jackson, G. J. Ackland, Lattice-switch monte carlo simulation for binary hard-sphere crystals, *Phys. Rev. E* 76 (2007) 066703.

- [8] M. Yang, H. Ma, Effect of polydispersity on the relative stability of hard-sphere crystals, *J. Chem. Phys.* 128 (13).
- [9] M. Marechal, M. Dijkstra, Stability of orientationally disordered crystal structures of colloidal hard dumbbells, *Phys. Rev. E* 77 (2008) 061405.
- [10] P. Raiteri, J. D. Gale, D. Quigley, P. M. Rodger, Derivation of an accurate force-field for simulating the growth of calcium carbonate from aqueous solution: A new model for the calcite-water interface, *J. Phys. Chem. C* 114 (13) (2010) 5997–6010.
- [11] D. Wilms, N. B. Wilding, K. Binder, Transitions between imperfectly ordered crystalline structures: A phase switch monte carlo study, *Phys. Rev. E* 85 (2012) 056703.
- [12] D. Quigley, Communication: Thermodynamics of stacking disorder in ice nuclei, *J. Chem. Phys.* 141 (12) (2014) 121101.
- [13] S. Bridgwater, D. Quigley, Lattice-switching monte carlo method for crystals of flexible molecules, *Phys. Rev. E* 90 (2014) 063313.
- [14] T. L. Underwood, G. J. Ackland, Lattice-switch monte carlo: the fcc-bcc problem, *JPCS* 640 (1) (2015) 012030.
- [15] N. B. Wilding, A. D. Bruce, Freezing by monte carlo phase switch, *Phys. Rev. Lett.* 85 (2000) 5138–5141.
- [16] J. R. Errington, Solid-liquid phase coexistence of the lennard-jones system through phase-switch monte carlo simulation, *J. Chem. Phys.* 120 (7) (2004) 3130–3141.
- [17] G. C. McNeil-Watson, N. B. Wilding, Freezing line of the lennard-jones fluid: A phase switch monte carlo study, *J. Chem. Phys.* 124 (6) (2006) 064504.
- [18] N. B. Wilding, Freezing parameters of soft spheres, *Mol. Phys.* 107 (4-6) (2009) 295–299.
- [19] N. B. Wilding, Solid-liquid coexistence of polydisperse fluids via simulation, *J. Chem. Phys.* 130 (10) (2009) 104103.
- [20] P. Sollich, N. B. Wilding, Crystalline phases of polydisperse spheres, *Phys. Rev. Lett.* 104 (2010) 118302.
- [21] N. B. Wilding, P. Sollich, Phase behavior of polydisperse spheres: Simulation strategies and an application to the freezing transition, *J. Chem. Phys.* 133 (22) (2010) 224102.
- [22] B. Fultz, Vibrational thermodynamics of materials, *Prog. Mater. Sci.* 55 (4) (2010) 247–352.
- [23] A. van de Walle, G. Ceder, The effect of lattice vibrations on substitutional alloy thermodynamics, *Rev. Mod. Phys.* 74 (2002) 11–45.
- [24] C. Vega, E. Sanz, J. L. F. Abascal, E. G. Noya, Determination of phase diagrams via computer simulation: methodology and applications to water, electrolytes and proteins, *J. Phys. Condens. Mat.* 20 (15) (2008) 153101.
- [25] M. B. Sweatman, Comparison of absolute free energy calculation methods for fluids and solids, *Mol. Phys.* 113 (9-10) (2015) 1206–1216.
- [26] J. Purton, J. Crabtree, S. Parker, DL\_monte: a general purpose program for parallel monte carlo simulation, *Mol. Simulat.* 39 (14-15) (2013) 1240–1252.
- [27] G. Ackland, K. DMellow, S. Daraszewicz, D. Hepburn, M. Uhrin, K. Stratford, The moldy short-range molecular dynamics package, *Comput. Phys. Comm.* 182 (12) (2011) 2587 – 2604.
- [28] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, E. Teller, Equation of state calculations by fast computing machines, *J. Chem. Phys.* 21 (6) (1953) 1087–1092.
- [29] B. A. Berg, T. Neuhaus, Multicanonical algorithms for first order phase transitions, *Phys. Lett. B* 267 (2) (1991) 249–253.
- [30] B. Berg, T. Neuhaus, Multicanonical ensemble: A new approach to simulate first-order phase transitions, *Phys. Rev. Lett.* 68 (1992) 9–12.
- [31] G. R. Smith, A. D. Bruce, A study of the multi-canonical monte carlo method, *J. Phys. A: Math. Gen.* 28 (23) (1995) 6623.
- [32] M. Fitzgerald, R. R. Picard, R. N. Silver, Canonical transition probabilities for adaptive metropolis simulation, *Europhys. Lett.* 46 (3) (1999) 282.
- [33] A. N. Jackson, Structural phase behaviour via monte carlo techniques, Ph.D. thesis, University of Edinburgh, Edinburgh, UK (2001).
- [34] M. S. Daw, M. I. Baskes, Embedded-atom method: Derivation and application to impurities, surfaces, and other defects in metals, *Phys. Rev. B* 29 (1984) 6443–6453.
- [35] [http://lammps.sandia.gov/doc/pair\\_eam.html](http://lammps.sandia.gov/doc/pair_eam.html) (accessed 3rd April 2016).
- [36] M. Matsumoto, T. Nishimura, Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator, *ACM Trans. Model. Comput. Simul.* 8 (1) (1998) 3–30.
- [37] M. I. Mendeleev, G. J. Ackland, Development of an interatomic potential for the simulation of phase transformations in zirconium, *Phil. Mag. Lett.* 87 (5) (2007) 349–359.
- [38] D. Y. Sun, M. I. Mendeleev, C. A. Becker, K. Kudin, T. Haxhimali, M. Asta, J. J. Hoyt, A. Karma, D. J. Srolovitz, Crystal-melt interfacial free energies in hcp metals: A molecular dynamics study of mg, *Phys. Rev. B* 73 (2006) 024116.
- [39] <https://github.com/tomlunderwood/monteswitch>.