

# Molecular Dynamics: from basic techniques to applications

## (A Molecular Dynamics Primer)<sup>1</sup>

E. R. Hernández

*Institut de Ciència de Materials de Barcelona, ICMA B - CSIC,  
Campus de Bellaterra, 08193 Bellaterra, Barcelona, Spain*

**Abstract.** It is now 50 years since the first papers describing the use of Molecular Dynamics (MD) were published by Alder and Wainright, and since then, together with Monte Carlo (MC) techniques, MD has become an essential tool in the theoretical study of materials properties at finite temperatures. In its early days, MD was used in combination with simple yet general models, such as hard spheres or Lennard-Jones models of liquids, systems which, though simple, were nevertheless not amenable to an analytical statistical mechanical treatment. Nowadays, however, MD is most frequently used in combination with rather sophisticated models, ranging all the way between empirical force fields to first-principles methods, with the aim of describing as accurately as possible any given material. From a computational aid in statistical mechanics and many-body physics, MD has evolved to become a widely used tool in physical chemistry, condensed matter physics, biology, geology and materials science. The aim of this course is to describe the basic algorithms of MD, and to provide attendees with the necessary theoretical background in order to enable them to use MD simulations in their research work. Also, examples of the use of MD in different scientific disciplines will be provided, with the aim of illustrating the many possibilities and the wide spread use of MD simulation techniques in scientific research today.

**Keywords:** Atomistic Simulation, Condensed Matter, Thermodynamics.

**PACS:** 02.70.-c, 02.70.Ns

## 1. INTRODUCTION

As a scientific tool for the study of condensed matter, Computer Simulation really started in the 1950s with the development of the first sufficiently powerful computers. The two main families of simulation techniques, Monte Carlo (MC) methods [1] and Molecular Dynamics (MD) methods [2], were described for the first time in that decade. During the second half of the XX century, Computer Simulation methods have established themselves as a mature and powerful research tool in condensed matter and molecular physics and chemistry, and are at present slowly but steadily extending their usefulness and applicability to other, more challenging areas such as biological systems and soft condensed matter.

The aim of this series of lectures is to provide an introduction to MD techniques, illustrating the power of these simulation tools. The outline of the material to be covered

---

<sup>1</sup> Dedicated to my teachers and mentors C. R. A. Catlow and M. J. Gillan, in recognition of an unpayable debt.

will be the following. First I will discuss, in very general terms, the aims and usefulness of Computer Simulation, and in particular MD methods, at least as I see it. A self-contained and to-the-point description of MD will be provided. Next, some historical background of MD will be reviewed, although inevitably we will not do justice to the extensive literature on this topic that has been published over the last five decades! In actual fact, I will just mention a few landmark papers, which in my own (admittedly personal) view are of key importance in the history of MD. Then we will move into the practical aspects of performing MD simulations, how to integrate the equations of motion, how to simulate bulk systems, etc. I will also make some general comments on the different approaches used to model the interactions between atoms and molecules, discussing some examples of potentials, and we will also see alternative approaches involving electronic structure calculations. After all that introductory material, we will get slightly more technical, and discuss how standard MD can be extended to simulate systems in contact with a thermal bath, emulating the conditions of constant temperature, or to simulate systems in which the volume and/or cell-shape fluctuate in such a way as to reproduce conditions of constant pressure. Illustrative examples of these different techniques will be provided along the way. We will wrap up with some discussions on the possible shortcomings and limitations of MD, and attempt to guesstimate some of the developments we are likely to see in the future, which inevitably will attempt to ameliorate some of those limitations.

Much of the material presented here (though not all) has been taken from standard references about computer simulation, and can be found there in more detail, together with lots of useful references to the literature. Standard text books on atomistic computer simulation are those of Allen and Tildesley [3], Frenkel and Smit [4] and Thijssen [5], which I recommend for more details on the topics discussed here.

## 2. MOLECULAR DYNAMICS IN A NUTSHELL

In a nutshell, MD consists of numerically solving the classical equations of motion for a collection of atoms<sup>2</sup>. For doing this, three basic ingredients are necessary; firstly, we must have some law describing the mutual interactions between the atoms in the system, from which we can calculate, given the atomic positions, the associated potential energy, the forces on the atoms, and if necessary the stress on the container walls. This law is in general unknown, but it can be approximated with different degrees of accuracy (and realism) by a force field, or it can be modelled by means of electronic structure calculations, which can also be done at different levels of theory. Secondly, we need an algorithm to integrate numerically the equations of motion for the atoms in the system. Over the years many different schemes have been put forward for doing this. Thirdly and finally, in order to solve the equations of motion, the integration scheme needs to be fed with some valid initial conditions, i.e. initial positions and velocities for all atoms in the system. With these three basic ingredients, one is set for performing MD simulations.

Before going into describing these different ingredients in somewhat more detail

---

<sup>2</sup> Here the word *atom* is used in a loose sense to refer indistinctly to atoms, ions or entire molecules

below, it is worthwhile to pause for a moment and consider why it may be useful to perform an MD simulation, and what can be extracted from it. With such a simulation, we are emulating, i.e. simulating in an approximate way, the real dynamics of the system under study, and in so doing we can keep track of the doings of individual atoms in an incredibly detailed way; so much so that one can easily feel like Big Brother among the atoms. In this way MD simulations can help us to gain new insight into important processes taking place at the atomic and molecular level, an insight which is often impossible to obtain purely from experiments, as these rarely have sufficient resolution. Furthermore, when performing simulations, one can easily prepare the conditions (temperature, pressure, atomic configuration, etc.) at will, and has a level of control over them that is much greater than is usually possible in experiments.

Aside from the numerical approximations involved in the integration of the equations of motion, there are two basic approximations inherent in MD simulations. The first one is that we assume that atoms behave like classical entities, i.e. they obey Newton's equations of motion. How much an approximation this is depends on the particular system under study, and on the actual conditions in which it is simulated. One can expect this approximation to be crude for light atoms at low temperatures, but in general it is not a bad approximation. In this respect it is fortunate that normally quantum effects on the atomic dynamics are relatively small, except in a few notable examples such as liquid He, and other light atoms. For those cases where quantum effects cannot be neglected, one should use the Path Integral approach [6] or some similar method.

The other key approximation is the model used to describe the interactions between the atoms in the system. It is clear that only through a sufficiently realistic description of those interactions one has any chance of getting useful and reliable information on the atomic processes taking place in the system. On the other hand, if one wants to address generic questions about a particular class of systems, such as low density gases, or liquid metals, say, one probably does not need to describe a particular example of such systems with a very accurate potential; it will be sufficient to use a generic model that captures the essential features, the defining physics, of that particular class of systems. To be too specific in this case can actually be counterproductive and obscure the general picture. It is therefore important to find the right level of description for the particular problem at hand.

This nutshell description may give you the idea that MD is simply solving Newton's equations for atoms and molecules. But in reality MD is much more than this: one can design rather artificial-looking forms of MD, which nevertheless serve a useful purpose, such as simulating a system under conditions of constant temperature and/or constant pressure (see section 7), something that is not possible to do by a straightforward solution of the standard equations of motion, or one can combine the physical dynamics of ions with a fictitious dynamics of electronic wave functions, which makes possible the effective realisation of atomic dynamics from first principles (the so-called Car-Parrinello method, see section 6). In essence, MD is extremely powerful and flexible, and far from being a simple numerical recipe for integrating the equations of motion for atoms and molecules.

### 3. SOME HISTORY

It is no surprise that the two most fundamental methodologies for simulating condensed matter systems, namely Monte Carlo (MC) and Molecular Dynamics (MD) made their first appearance in the 1950s. At this time the first computers, originally available only for classified military research, were made available to scientists in the US, and the possibility of performing fast automated calculations was immediately seen to have great potential for problems in statistical mechanics, for example. Ever since the first publication describing the MD technique, by Alder and Wainright [2] in 1957, applications of the technique have been growing in number, and nowadays MD is an extensively used research tool in disciplines which include physics, chemistry, materials science, biology and geology.

In the early days of MD, covering mostly the 1960s and 70s, the technique was mostly used as an aid in statistical mechanics. For the largest part, there was no attempt to model realistic systems, but rather the focus was on simple, generic model systems such as hard spheres or the Lennard-Jones fluid. The aim was not so much to address questions concerning specific systems, but rather to learn about entire families of systems, e.g. simple liquids. In time, models grew in complexity and in their degree of specificity. Empirical (i.e. derived from experimental information) models began to be developed for specific classes of systems, such as the CHARMM [7] or AMBER [8] force fields for organic and biological molecules, the ionic potentials for oxide materials [9], the embedded atom potentials for metals [10], or the bond-order inspired potentials for covalent materials [11].

At the same time, new methodological developments were being carried out. Since MD consists basically of integrating the classical equations of motion for the atoms or molecules of a system, it was implicitly accepted that MD could only be used to simulate systems in microcanonical conditions, i.e. conditions of constant number of particles,  $N$ , constant volume,  $V$ , and constant energy,  $E$ . This was somewhat limiting, as experiments are most often conducted on samples which are not isolated, but in thermal and/or mechanical contact with their surroundings. However, in an influential paper, Andersen [12] demonstrated that new, more general forms of MD could be devised. Andersen introduced two new tools, known as the Andersen thermostat and the Andersen barostat, which, as their name indicates, serve the purpose of controlling the temperature and the pressure during the simulation, respectively. In section 7 we will discuss the details of Andersen's thermostat and barostat; for now let us just remark that particularly the idea of the barostat has proved to be very influential in the subsequent history of MD. In essence, Andersen introduced a new variable into the dynamics of the system, namely the system's volume, with an associated velocity, a fictitious mass, and a potential energy term depending on the external pressure. The coupled dynamics of atoms and volume proposed by Andersen ensured that the system samples the isoenthalpic (constant enthalpy) ensemble, which is useful for analysing how the system may react to an externally imposed pressure. Andersen showed that, by introducing a small number of additional fictitious degrees of freedom (the volume) it was possible construct a new dynamics which effectively achieved the same effect as coupling the system to the infinitely many degrees of freedom of a reservoir. As pointed out above, this idea was to prove extremely influential.

Shortly after Andersen's paper was published, Nosé [13] showed that the introduction of an additional fictitious variable coupled to the atomic dynamics could be done in such a way as to obtain sampling in the canonical (constant temperature) ensemble. Contrary to the thermostat already introduced by Andersen [12], which affects the atomic dynamics in a stochastic way, Nosé's thermostat is fully deterministic. Nosé's approach, as later modified by Hoover [14], has now become perhaps the most commonly used scheme for performing MD simulations in the canonical ensemble.

The constant-pressure scheme of Andersen, originally conceived for the simulation of bulk fluids, was not generally applicable to crystalline solids, because only volume, and not shape fluctuations were considered. Parrinello and Rahman [15] generalised the method of Andersen by incorporating the components of the lattice vectors of the simulation cell as new fictitious dynamical variables, thus making possible the observation of solid-solid phase transitions in MD simulations. This scheme also made possible the study of systems under non-hydrostatic stress conditions.

Andersen's barostat and Nosé's thermostat proved that MD was potentially much more than simply a scheme for solving the equations of motion for a collection of atoms isolated from the rest of the universe. By adequately incorporating appropriately designed fictitious variables, these developments showed that more general and experimentally relevant statistical ensembles could be sampled. But the introduction of fictitious variables was soon to be found to have even wider possibilities: in 1985, i.e. only 5 years after Andersen's barostat had been introduced, Car and Parrinello [16] demonstrated a new use of fictitious dynamical variables. In their seminal paper, Car and Parrinello showed for the first time that it was possible to perform *ab initio* MD, i.e. MD in which the forces on the atoms are not extracted from an empirical force field, but rather from a full blown first principles electronic structure calculation. This combination of methods has been given the name first principles molecular dynamics (FPMD), also known as *ab initio* molecular dynamics (AIMD).

Before Car and Parrinello's paper, FPMD had been regarded as essentially impossible mostly due to the computational cost involved in performing a time-consuming electronic structure calculation for each time step of an MD simulation, i.e. thousands or even tens of thousands of times. Computers were simply not fast enough for the task in 1985. However, Car and Parrinello showed that with a clever introduction of new fictitious variables, the cost of FPMD could be brought down significantly, so much so as to make it a realistic undertaking, even with the computers of the day. Briefly, Car and Parrinello's idea consisted of bringing electrons and ions simultaneously into the picture, but in a very unusual and imaginative way. Just as in conventional force-field MD, ions moved subject to the forces acting on them, but these forces came not from an empirical potential, but from their mutual (coulombic) interaction, and from their interaction with the valence electron density around them. Car and Parrinello formulated their FPMD in the context of density functional theory (DFT) [17, 18] formalism. Within this formalism, the electron density is obtained from a series of so-called Kohn-Sham orbitals, which are the solutions of a Schrödinger-like equation, the Kohn-Sham equation. These orbitals must be obtained for the given ionic configuration before the total energy of the system and the forces of the ions can be calculated, and this process is considerably more costly than any calculation based on force fields. Typically, Kohn-Sham orbitals are represented by means of some basis set of appropriately chosen functions, such as

atomic-like orbitals, or plane-waves, the latter being particularly convenient in the case of periodic systems. Then, solving the electronic structure problem consists of finding the appropriate expansion coefficients for the relevant Kohn-Sham orbitals in terms of the basis set functions. The break-through of Car and Parrinello was to incorporate the expansion coefficients of the Kohn-Sham orbitals in terms of the basis functions as fictitious dynamical variables, with associated fictitious masses. By choosing these masses appropriately (set to values much smaller than those of the ions), the Kohn-Sham orbitals evolve much more rapidly than the ions, and as a result a regime is established in which the orbitals adapt quasi-instantaneously to the comparatively slow change in the ionic positions. This is the Born-Oppenheimer approximation again, but in a new, imaginative setting.

The achievement of Car and Parrinello served the purpose of waking up the scientific community to the fact that FPMD was indeed viable, and soon many groups worldwide began to perform FPMD simulations, either directly employing the Car-Parrinello scheme, or alternative ones. Throughout the 1990s and this century, FPMD has now become a relatively standard tool, with an impressive showcase of applications. This is not to say that FPMD has completely supplanted the simpler, more approximate force-field based MD; far from it. There are many problems that remain too challenging to tackle via FPMD, either because the system is too large, too complex, or because it cannot be modelled accurately enough with DFT. In such cases empirical force fields continue to be the only viable option, and this is likely to remain the case for some time to come.

Up to here what is now the history of MD, according to an admittedly personal view. As for future developments, well, as the saying goes, making predictions is extremely difficult (especially about the future!), but at the end of this chapter I will try to summarise what we can already begin to see, or guess, for the relatively short term future of MD.

## 4. MD: BASIC TECHNIQUES

In this section we are going to review some practical aspects of MD simulations, such as how to integrate numerically the classical equations of motion, how to deal with infinite systems, how to start and run a simulation, and how to analyse the results.

### 4.1. Integrating the equations of motion

Much has been written about how to integrate the equations of motion of a dynamical system most effectively and accurately. This is more an issue of applied mathematics [19] than of physics (although some methods have a very physical inspiration), therefore we are not going to go in depth here. I just want to provide a simple recipe, which will be useful in most cases that we are likely to encounter.

The classical equations of motion have the general form

$$\dot{q} = G(p, q), \quad \dot{p} = F(p, q), \quad (1)$$

where  $G(p, q) = \partial H / \partial p$  and  $F(p, q) = -\partial H / \partial q$ , and  $H$  is the Hamiltonian, which in the standard case is given by

$$\mathcal{H} = \sum_i \frac{\mathbf{p}_i^2}{2m_i} + U(\mathbf{q}), \quad (2)$$

where  $\mathbf{q}_i$  represents the coordinates of atom  $i$ , and  $\mathbf{p}_i$  is its conjugate momentum. The recipe which we will use is known as the *generalised leapfrog*, and is summarised as follows. First, we advance the momenta in time half a time step, then, with the momenta at half time step we move the coordinates forward in time by a full time step, recalculate the forces at the new positions, and with these new forces, advance the momenta to full time step. The algorithm is symbolically written down as:

$$\begin{aligned} p(t + \Delta t/2) &= p(t) + \Delta t F[p(t + \Delta t/2), q(t)]/2, \\ q(t + \Delta t) &= q(t) + \Delta t \{G[p(t + \Delta t/2), q(t)] + \\ &\quad G[p(t + \Delta t/2), q(t + \Delta t)]\}/2, \\ p(t + \Delta t) &= p(t + \Delta t/2) + \Delta t F[p(t + \Delta t/2), q(t + \Delta t)], \end{aligned} \quad (3)$$

where  $\Delta t$  is the time step. For the simple case of a separable Hamiltonian such as that of Eq. (2), the generalised leapfrog algorithm reduces to:

$$\begin{aligned} p(t + \Delta t/2) &= p(t) - \frac{\Delta t}{2} \frac{\partial U}{\partial q}(t), \\ q(t + \Delta t) &= q(t) + \Delta t p(t + \Delta t/2), \\ p(t + \Delta t) &= p(t + \Delta t/2) - \frac{\Delta t}{2} \frac{\partial U}{\partial q}(t + \Delta t), \end{aligned} \quad (4)$$

which is known simply as the leapfrog algorithm. Using Eqs. (4) repeatedly, one can map out a trajectory from specified initial conditions (coordinates and momenta of all atoms in the system). Provided the time step  $\Delta t$  is sufficiently small, this scheme conserves the energy reasonably accurately, and is time reversible, as the equations of motion are. The generalised leapfrog (and therefore the leapfrog) is accurate to second order in  $\Delta t$ .

The classical equations of motion have many properties, of which the most obvious one is time reversibility, but there are other ones. A particularly important symmetry is that known as *symplecticity* or *symplecticness*. Consider the following sum of infinitesimal areas:

$$\sum_i \delta \mathbf{r}_i \times \delta \mathbf{p}_i, \quad (5)$$

where the sum extends over all degrees of freedom of the system, and the deltas imply infinitesimally short vectors centred at each position and momentum. It can be easily shown that this infinitesimal area is a constant of motion of classical mechanics. It is important that any numerical scheme for integrating the equations of motion respects as many as possible of the intrinsic properties of the equations of motion; the more of such properties that are respected, the greater the guarantee that we will have that the

numerical solution found will resemble a physically correct trajectory. The generalised leapfrog scheme described above is time reversible and symplectic, and due to this it is particularly stable. A more in-depth discussion of these issues can be found in the book by Sanz-Serna and Calvo [19].

There are many other schemes for integrating numerically the equations of motion, and excellent discussions can be found in the literature [3, 4, 5], but the generalised leapfrog will suffice as an example for us, and turns out to be one of the best schemes in the market anyway.

## 4.2. Periodic boundary conditions

Frequently one is confronted with the need to study a system, periodic or not, which contains large numbers of atoms or molecules, where large means of the order of  $N_A$ , Avogadro's number. Naturally, we cannot deal with such large numbers, so we must resort to some computational tricks in order to emulate a system in these conditions. The trick used in this case is referred to as *periodic boundary conditions* (PBC), and consists of assuming that the simulation box (i.e. the box containing the atoms in the simulation) is surrounded by identical copies of itself in all directions. In the simulation of periodic systems the simulation box is typically a (super)cell with the periodicity of the system. If the simulation does not involve studying the dynamics of the system (as when we do a structural relaxation) then there is no approximation involved in the use of PBC, unless we are doing an electronic structure calculation, in which case care has to be taken in order to sample the electronic states in regions of the Brillouin zone beyond the  $\Gamma$  point to ensure the convergence of the calculation<sup>3</sup>. If, however, we are interested in the dynamics of the system, the use of PBC involves an approximation, even if the system is periodic. Adopting PBC implies assuming that all periodic images of atoms in the central simulation box move in exactly the same way.

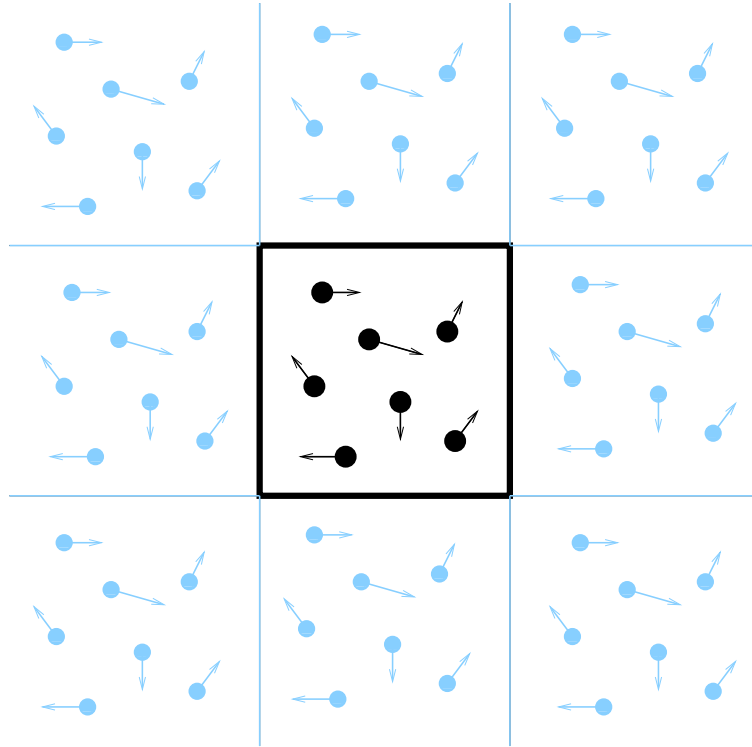
In a liquid, or in an amorphous solid, the use of PBC imposes an artificial symmetry, the consequences of which can be subtle. Further limitations arise in the study of defects and impurities in solids: the use of PBC generally implies that one is considering a large concentration of the defect or impurity under consideration, since we cannot always make the simulation box as large as we would like due to the computational cost involved in doing so. Furthermore, there are certain kinds of defects that cannot be easily accommodated in a periodic cell. This happens with dislocations. In such cases one has to include two dislocations of opposite sign so that they can both be included within a periodic cell, or renounce to the use of PBC altogether.

Another consideration when using PBC is that long-range interactions have to be dealt with appropriately. Electrostatic forces have such a long-range that it is necessary to include the contribution of far-away periodic images of the simulation box on the atoms contained in the central simulation box in order to get meaningful results. Simply truncating electrostatic interactions beyond a certain cutoff is crude, and generally frowned

---

<sup>3</sup> The discussion of this very important point is however specific to electronic structure calculations, and we will not discuss it here; see e.g. [20].





**FIGURE 1.** Periodic Boundary Conditions; illustration in two dimension. The simulation box is highlighted at the centre, and is surrounded by periodic images of itself.

upon. Several ways have been described in the literature to deal with long-range interactions. Perhaps the most popular procedure is that due to Ewald, known as Ewald Summation. I will not go into details here, but I would like to describe the basic idea. The problem with electrostatic interactions is that terms of order  $1/r$  decay very slowly with  $r$ , the interatomic distance. The Ewald summation method consists of splitting this term as follows:

$$\frac{1}{r} = \frac{\text{erf}(\alpha r)}{r} + \frac{\text{erfc}(\alpha r)}{r}, \quad (6)$$

where erf is the error function, erfc is the complementary error function, and  $\alpha$  is a parameter. The idea is that the first term on the rhs of Eq. (6) can be shown to be short-ranged in reciprocal space, where it can be easily evaluated, and likewise, the second term is short-ranged in real space. The parameter  $\alpha$  is chosen so that an optimal split between the real-space sum and the reciprocal-space sum is obtained.

A final consideration concerning PBC is the following. PBC are a useful device for calculating the energy and its derivatives (forces and stresses, see below) as if the atoms of the system were indeed in an infinite system, or at least sufficiently far away from any surface to notice its presence. However, in general it is not required, nor is it desirable, to modify the positions of the atoms as they move so that they all lie in the central simulation box, particularly if one is interested in monitoring the diffusion or other dynamical properties of the system. So, when using PBC, one calculates the total energy

and the forces on the atoms as if all atoms were relocated in the central simulation box, regardless of their actual position in space, but they are not moved back to the simulation box if during the simulation they drift out of it.

### 4.3. Derivatives of the total energy

In order to integrate the equations of motion for the atoms constituting the system, we must be able to obtain the forces, i.e. the derivatives of the total energy with respect to the atomic positions:

$$\mathbf{f}_i = -\nabla_{\mathbf{r}_i} E_{tot}, \quad (7)$$

where  $E_{tot}$  is the total energy of the system. In the standard case, the kinetic energy does not depend on the atomic positions, and so only the derivative of the potential energy has to be considered.

Another useful derivative of the total energy, needed if one wishes to calculate the pressure of the system, or conduct a constant pressure simulation (constant pressure simulations will be discussed in section 7, is the stress. The stress is also useful because it can be related with the elastic constants of a crystal. The stress is defined as the derivative of the total energy with respect to the components of the strain tensor. The strain,  $\varepsilon$ , defines infinitesimal distortions of the simulation box. For example, consider that the simulation box (not necessarily cubic or orthorhombic) is defined by the three vectors  $\mathbf{a}_\alpha$ , with  $\alpha = 1, 2, 3$ . Then, a distortion of the simulation box defined by the strain tensor  $\varepsilon$  will lead to new cell vectors  $\mathbf{a}'_\alpha$  given by

$$a'_\alpha = a_\alpha + \sum_{\beta} \varepsilon_{\alpha\beta} a_\beta. \quad (8)$$

Because a distortion of the cell will cause the distances between atoms and the angles between bonds in the cell to change, such a distortion changes the total energy, and that change is given to first order by the stress, defined as

$$\sigma_{\alpha\beta} = \frac{\partial E_{tot}}{\partial \varepsilon_{\alpha\beta}}. \quad (9)$$

Note that in this definition we have dropped a minus sign, so that the stress is not minus the derivative of the total energy with respect to the strain. With this definition, a negative applied stress is tensile, while a positive applied stress is compressive, resulting in an intuitive convention.

From Eq. (8) we see that the strain is a tensor with dimensionless components. Furthermore, we are only interested in symmetric strain tensors, because any asymmetric strain tensor involves rotation as well as distortion of the cell. Rigid rotations of the system are not interesting, however (at least in the absence of external fields), and complicate the dynamics, so, in practise, we will always be concerned with symmetric strain tensors, and, for the same reasons, the stress tensor will also be symmetric.

## 4.4. Start-up of an MD simulation

Imagine that you have written an MD code, which integrates the equations of motion for a given model, and everything works correctly. You are now in a position to perform a simulation. How does one start? Typically, one needs a starting configuration of the system. For a crystal, the perfect lattice will serve the purpose, although one can use also a slightly distorted version of the lattice (in fact this may have some advantages in achieving the thermalisation of the system; see below). For a liquid the initial configuration may be less obvious. In this case one can start with a lattice configuration known to be unstable at the temperature of the simulation, and hope that this will evolve rapidly toward configurations typical of the liquid phase.

As well as coordinates, one needs to generate initial velocities for the atoms. The almost universally adopted choice is to generate random velocities sampled from the Maxwell-Boltzmann distribution defined for the desired temperature of the simulation. These may need to be corrected, so that the centre of mass of the system has zero velocity (this avoids the drift of the system as the simulation proceeds). In the case of finite systems it is also useful to avoid the rotation of the system, by ensuring that its total angular momentum is zero.

Generally, one wishes to conduct a simulation at a given temperature. However, even if one generates initial velocities sampling from the Maxwell-Boltzmann distribution corresponding to the desired temperature, the evolution of the system will drive the temperature to other values, such that its average over the simulation run will not in general coincide with (or even be near) the desired temperature. This happens essentially because the starting atomic positions are not necessarily consistent with the desired temperature. It is therefore necessary to drive the system from its starting conditions to other conditions, compatible with the desired temperature. This is generally done by scaling the velocities during an initial period of the simulation, usually referred to as *thermalisation* or *equilibration*. Each atomic velocity is scaled by a factor  $\sqrt{T_{ext}/T_{inst}}$ , where  $T_{inst}$  is the *instantaneous* temperature (see below), and  $T_{ext}$  is the desired equilibrium temperature. This scaling will slowly drive the system toward the desired conditions; it can be done every time step of the equilibration period, or every few time steps. Obviously, this tampering with the velocities results in a lack of energy conservation. The dynamics is thus artificial, and only serves the purpose of preparing the system in conditions from which the real simulation can start. Therefore, no information obtained during this period is useful, and should not be included in the subsequent analysis of results.

How long should the equilibration period be really depends on the nature of the simulated system, but also on practical considerations, such as the cost involved. Ideally, one should run the equilibration period for long enough so that the system has lost any “memory” of its initial conditions, and is fully at equilibrium at the desired temperature. Once this is achieved, the average temperature should be close to (usually not more than a few degrees away from) the desired temperature. If this does not happen, then obviously the equilibration period was not sufficiently long.

## 5. ANALYSING THE RESULTS

MD simulations can produce a wealth of information, ranging from the time evolution of the coordinates and velocities of individual atoms to other so called “collective” properties such as the temperature, pressure, and so on. In this section we review the standard magnitudes that are monitored during an MD simulation.

### 5.1. Temperature

The temperature in a simulation can be calculated directly from the standard expression from statistical mechanics relating it to the kinetic energy of the atoms. This expression is

$$T_{inst} = \frac{2}{gk_B} E_{kin}, \quad (10)$$

where  $E_{kin}$  is the kinetic energy at the present time,  $g$  is the number of degrees of freedom of the system and  $k_B$  is Boltzmann’s constant. This expression gives the *instantaneous* temperature of the simulation. This value will be different at different time steps; really only its average value over the length of the simulation gives a meaningful value to the temperature:

$$\langle T \rangle = \frac{1}{N} \sum_n T_{inst}(n), \quad (11)$$

where the sum extends over all time steps (or a subset) of the simulation,  $N$ . Only if an appropriate equilibration period has been undertaken before the actual simulation (see 4.4 above) will one have that the temperature of the simulation will be close to the desired target temperature, i.e.  $\langle T \rangle \approx T_{target}$ .

### 5.2. Pressure

The pressure is another useful magnitude to monitor through a simulation. Its average value will provide information on the mechanical state of the system, i.e. if the system is compressed or expanded with respect to its equilibrium volume at the temperature of the simulation. The instantaneous pressure of the system is given by

$$P_{inst} = \rho k_B T + \frac{1}{3V} \left\langle \sum_{i<j} \mathbf{f}_{ij} \cdot \mathbf{r}_{ij} \right\rangle, \quad (12)$$

where  $\rho$  is the number density,  $V$  is the volume,  $\mathbf{f}_{ij}$  is the pair force between atoms  $i$  and  $j$ , and  $\mathbf{r}_{ij}$  is the distance vector. This expression is actually only valid for the case of pair interactions, and must be generalised in more complicated models. Like in the case of the temperature, only the average value,  $\langle P \rangle$ , makes sense, as the value of  $P_{inst}$  will fluctuate strongly in time.

If it is desired to perform the simulation at a pre-specified pressure, one has to adjust the volume of the simulation cell in such a way that the average pressure has the desired value. This usually requires performing several short simulations. The average pressure is a smooth function of the simulation volume, so it is usually sufficient to find two volumes which give average pressures which bracket the desired pressure, and then use linear interpolation to obtain the volume which would correspond to the desired average pressure. This procedure is simple if the shape of the simulation box is fixed, as is the case in a liquid, or in a solid where no phase transition is expected. However, if the crystal shape is complicated or unknown a priori, it is probably more desirable to conduct a variable-shape MD simulation at constant pressure, which allows the system to dynamically adopt volume and shape to the most favourable values at the conditions of the simulation. These are called constant-pressure MD simulations, and will be discussed in section 7.

### 5.3. Structure


Even though at finite temperatures the atoms of the system will never be at rest, the system will have a definite structure. Let us consider the atomic density of the system at each point in space,  $\rho(\mathbf{r})$ . We can write such a density as the following thermal average:

$$\rho(\mathbf{r}) = \langle \sum_{i=1}^N \delta(\mathbf{r} - \mathbf{r}_i) \rangle . \quad (13)$$

If the system is crystalline or amorphous at the conditions of the simulation,  $\rho(\mathbf{r})$  will peak at the average positions of the atoms, and will fall to low values close to zero at interstitial regions, which will be visited infrequently by the atoms. If, on the other hand, the system is fluid, then  $\rho(\mathbf{r})$  should be constant everywhere, and equal to the bulk density. In the crystalline case  $\rho(\mathbf{r})$  will have the periodicity of the lattice and the system will possess so-called long-range order. But even if the system is amorphous or fluid, it will possess short-range order, and it will be interesting to characterise it. Two quantities are frequently used to achieve this, known as the *radial distribution function* and the *bond-angle distribution function*.

#### 5.3.1. Radial distribution function

The radial distribution function (RDF) is constructed as a histogram of the distances between an atom and its neighbours during the simulation. Assuming that we are dealing with a one-component system, all atoms are equivalent, and the RDF is then averaged over all atoms. Suppose that we want to calculate the RDF, typically called  $g(r)$ , in the range  $r_0$  to  $r_{max}$ . To do so, we divide this range into a number of equally spaced segments of length  $\delta r$ , and add a 2 in the appropriate segment of the  $g(r)$  histogram for each pair of atoms separated by a distance  $r$  corresponding to that segment. Such a histogram will diverge for large distances, since the probability of finding two atoms separated by  $r$




**FIGURE 2.** Radial distribution function calculated for Pb using the empirical potential due to Cleri and Rosato [21]. Results are shown for the solid fcc phase at 500 K and for the liquid at 700 K, both at zero pressure.

when  $r$  is large grows very rapidly, so it is customary calculate  $g(r)$  as the probability of finding two atoms separated by distance  $r$  relative to the probability of finding two atoms at the same distance in the ideal gas of the same density. According to this definition,  $g(r)$  should tend to 1 as  $r$  becomes large.

In Fig. (2) a typical RDF function obtained from a simulation is shown for the particular case of solid and liquid Pb. As can be seen, the RDF is zero at short distances, reflecting the exclusion volume around a given atom; it then grows rapidly to reach a maximum at the nearest neighbour distance, falling down again to a first minimum, which can be followed by other (smaller) maxima and minima, or smoothly evolve to its large-distance limit of 1. The volume integral of the RDF from zero to the first minimum after then nearest neighbour peak gives the number of first neighbours of each atom, i.e. the average coordination of atoms through the simulation:

$$n_c = 4\pi\rho_0 \int_0^{r_{min}} r^2 g(r) dr, \quad (14)$$

where  $\rho_0$  is the bulk density.




**FIGURE 3.** Bond-angle distribution function for liquid Na, calculated using ab initio simulations [22].

### 5.3.2. Bond-angle distribution function

Since the RDF gives a distribution of distances between atoms, it does not have any angular resolution. Therefore, in order to complete the picture of the short-range environment of atoms in the simulation, it is frequent to calculate what is called the Bond-angle distribution, or BAD. The BAD is exactly what its name implies, i.e. a distribution of the bond-angles, or actually, the cosine of the bond angles, found between an atom and its first shell of neighbours, taken two by two, with the atom in question forming the apex of the bond angle. A typical example of the bond-angle distribution found from a simulation is shown in Fig. (3), calculated for liquid Na.

## 5.4. Dynamics

The power of MD simulation as compared to MC is that it also provides information on the dynamics of the system, not just the structure. The dynamics of a system can be regarded from many angles: firstly, there is atomic motion, which can be vibrational around the equilibrium sites of atoms in a solid, or hopping from one site to another, or even diffusive. But we also have the so-called collective dynamics, such as density fluctuations, sound, or viscous flow. Much could be written (and has indeed been written, see for example [23]) about all these forms of dynamics and how to study them. But here we are going to limit ourselves to the most frequent kind of dynamical analysis, that related to diffusive motion.



**FIGURE 4.** Mean-squared displacements of Pb at different temperatures, just before (500 K, left panel) and after melting (700 K, right panel) has taken place. The non-diffusive behaviour of the solid phase, and the diffusive behaviour of the liquid phase can be clearly appreciated. Taken from the simulations reported in [24].

#### 5.4.1. Mean-squared displacements

In order to answer the question *how much does an atom move from its initial position in a given time  $t$ ?* one can calculate the averaged mean-squared displacements (MSD), given by

$$\langle |\mathbf{r}(t) - \mathbf{r}(0)|^2 \rangle = \sum_{t_0} \sum_{i=1}^N \frac{|\mathbf{r}_i(t+t_0) - \mathbf{r}_i(t_0)|^2}{N_{t_0} N}, \quad (15)$$

where we have calculated the average taking advantage of the fact that different time origins  $t_0$  can be taken,  $N_{t_0}$  being the number of such time origins, and since all atoms of the same species are equivalent, we can also average over them. An example of MSDs as obtained from a typical simulation is illustrated in Fig. (4).

MSDs are important because they provide information on how fluid-like a system is. In Fig. (4) we show examples of how the MSD look in a system (in this particular case Pb) before and after the system has melted. It can be seen that while the system is in the solid phase, the MSD are flat, having zero slope at all times but very short ones. This is because in the solid atoms don't travel large distances, but rather oscillate around their equilibrium positions. The amplitude of those oscillations is related to the saturation value of the MSDs at long times. On the other hand, if the system is in the liquid phase, atoms can move and diffuse through the system. In this case the MSDs grow with time having a well-defined slope, which is related with the diffusion coefficient:

$$6Dt + b = \langle |\mathbf{r}(t) - \mathbf{r}(0)|^2 \rangle, \quad (16)$$

where  $b$  is a constant, and  $D$  is the diffusion constant of the system. If there is more than one species of atoms in the system, one can define a diffusion constant for each species by generalising the above expression. For the particular example illustrated in Fig. (4),



one obtains a value of  $D$  equal to  $2.0 \times 10^{-5} \text{cm}^2/\text{s}$ , very close to the experimental value of  $2.0 \times 10^{-5} \text{cm}^2/\text{s}$  found just after melting has taken place.

#### 5.4.2. Velocity auto-correlation function

The last magnitude we will discuss is that of the velocity autocorrelation function (VAF). The VAF is defined as follows:

$$\text{VAF}(t) = \langle \mathbf{v}(t) \cdot \mathbf{v}(0) \rangle, \quad (17)$$

and is calculated from a simulation following the same scheme as for the MSD given in Eq. (15) above. It is also frequent to work with the normalised VAF, defined by

$$\hat{\text{VAF}}(t) = \frac{\langle \mathbf{v}(t) \cdot \mathbf{v}(0) \rangle}{\langle \mathbf{v}(0) \cdot \mathbf{v}(0) \rangle}. \quad (18)$$

In actual fact, the VAF also provides information on the diffusivity of the different species in the system, as it is linked with the diffusion constant through the following relation:

$$D = \frac{1}{3} \int_0^\infty \langle \mathbf{v}(t) \cdot \mathbf{v}(0) \rangle dt. \quad (19)$$

The VAF is really telling us how much time it takes for an atom in the system to “forget” its original velocity at time zero, through collisions with other atoms in the system. It starts at a positive large value, equal to 1 in the case of the  $\hat{\text{VAF}}$ , and has an oscillating behaviour, falling more or less exponentially to zero as time increases. The Fourier transform of the VAF is given by

$$\text{VAF}(\omega) = \frac{1}{\pi} \int_0^\infty \cos(\omega t) \text{VAF}(t) dt, \quad (20)$$

and is related to the phonon density of states through

$$\text{VAF}(\omega) \propto \text{VDOS}(\omega) e^{-\beta \hbar \omega}, \quad (21)$$

where  $\text{VDOS}(\omega)$  is the vibrational density of states at frequency  $\omega$ , and  $e^{-\beta \hbar \omega}$  is the corresponding Boltzmann factor for that frequency ( $\beta = (k_B T)^{-1}$ ).

## 5.5. Summary

In this section we have discussed some of the practical aspects of performing an MD simulation, ranging from the numerical integration of the equations of motion to how to deal with infinite systems, and how to start up the simulation. We have also discussed some of the typical properties that one analyses during or after an MD simulation has

been performed, and how this analysis provides information on the properties of the system. Really, we have only skimmed the surface. There is much more to say about how to perform simulations and how to analyse the results than we can discuss in these notes, but at least I hope that you have got the basics, and can build from here if needed.

## 6. MODELLING INTERACTIONS IN ATOMIC SYSTEMS

A key issue in any form of modelling atomistic systems, be it by means of MD, MC, or any other technique, is the representation of the interactions between the atoms or molecules that make up the system. Ultimately, these interactions are the result of the subtle interplay of electrons and nuclei. This interplay can give rise to a wide variety of behaviours; some systems display a covalent type of bonding, while others would be better described as ionic, although more generally the situation is intermediate between these two extremes. Yet in other systems neither of these patterns takes place; rather, one has a metallic behaviour, where a significant portion of the electrons are free to move through the entire system, without being associated to a particular atom or bond. The opposite extreme to this is that of noble gases, where electrons are tightly attached to individual atoms, adopting a closed-shell electronic structure. To complicate matters further, several of these widely different behaviours can be displayed by one and the same system at different temperature and/or pressure conditions. For example, at ambient conditions, silicon is known to adopt the diamond structure, which is a covalent semi-conductor. But upon raising the temperature beyond  $\approx 1670$  K silicon melts, and in so doing it becomes metallic. Likewise, upon applying pressure at fixed temperature, the diamond structure eventually undergoes a phase transition to the so-called  $\beta$ -Sn structure, which is also metallic.

In general, we can find two different approaches to describing the interactions between atoms and molecules in a system. The first one is to employ some form of potential, i.e. a (in general complicated) function which depends on the relative interatomic positions (distances and angles) and on a series of parameters which must be fitted in order to reproduce as closely as possible some relevant properties of the system, such as a crystal structure, elastic or vibrational properties, etc. In this approach the electronic structure is obviated; rather one attempts to account for its effects with the potential function. The second approach involves retaining the picture of the system as composed by electrons and nuclei, and to obtain the energetics of the system as well as the forces on the atoms from a quantum mechanical treatment of the electronic structure, either at the semi-empirical level or through a fully first-principles treatment. This approach is theoretically more sound, but obviously more expensive.

The first approach is usually termed the *empirical potential* approach. The name makes reference to the fact that in general the form of the potential is *ad hoc*, i.e. there is no underlying guiding principle as to what the mathematical expression of the potential should be, beyond the fact that it should be repulsive at short distances, attractive at intermediate ones and decay to zero at infinite separation. It also refers to the fact that the potential has a series of disposable parameters that must be fitted, traditionally to empirical information on the system, though lately it is very common to parametrise potentials to results obtained with more accurate theoretical calculations, usually based

on electronic structure calculations. To give an exhaustive review of the different types of potentials used in the literature would be a daunting task, far beyond the scope of what only aims to be an introduction to MD. Rather, I will just name a few common examples, and refer the interested reader to the appropriate literature.

It is customary to assume that the total potential energy of an atomic system in the absence of external fields can be written as a series of the form

$$U = \sum_{i,j} V_2(\mathbf{r}_i, \mathbf{r}_j) + \sum_{i,j,k} V_3(\mathbf{r}_i, \mathbf{r}_j, \mathbf{r}_k) + \sum_{i,j,k,l} V_4(\mathbf{r}_i, \mathbf{r}_j, \mathbf{r}_k, \mathbf{r}_l) + \dots, \quad (22)$$

where  $V_2$  represents the energy of interactions of pairs of atoms,  $V_3$  that of triads, and so on. In practice this series is rarely taken beyond the sum containing the  $V_4$  terms, and frequently it is truncated after the first or second sums.

Although much work has been done with discontinuous potentials such as hard spheres, here we will focus on continuous potentials. Of these, perhaps the simplest, though still extensively used, is the Lennard-Jones potential, which takes the form

$$V_{LJ} = 4\varepsilon \left[ \left( \frac{\sigma}{r} \right)^{12} - \left( \frac{\sigma}{r} \right)^6 \right]. \quad (23)$$

The potential is characterised by two parameters, namely  $\varepsilon$ , which has dimensions of energy, and which determines the minimum value of the potential, and  $\sigma$ , which has dimensions of length, and is related to the position of the minimum. The first term in the squared brackets of Eq. (23) causes the potential to be strongly repulsive at short distances, while the second term has the typical form expected for dispersion-type interactions, which decay as  $r^{-6}$ . This potential has been frequently used to model e.g. noble gases, and in such a case it is an example of a model in which Eq. (22) is truncated after the first sum. Eq. (23) is frequently used also to describe non-bonded type interactions (i.e. interactions between atoms that are not linked by a chemical bond) in more complex molecular systems.

Covalent systems generally have a more complex bonding, where interactions are not only distance dependent but also directional. Several potentials have been put forward for such systems, out of which perhaps the best well known are those of Stillinger and Weber [25] for silicon, and that of Tersoff [11], also for silicon, but which has been parametrised also for carbon. Both potentials have been extensively used for modelling covalent materials, and have inspired the formulation of more sophisticated models. For example, the Tersoff [11] potential has the following expression:

$$U = \frac{1}{2} \sum_i \sum_{i \neq j} f_C(r_{ij}) [f_R(r_{ij}) + b_{ij} f_A(r_{ij})], \quad (24)$$

where  $f_R(r_{ij}) = A_{ij} e^{-\lambda_{ij} r_{ij}}$  and  $f_A(r_{ij}) = -B_{ij} e^{-\mu_{ij} r_{ij}}$  are repulsive and attractive pair potentials, respectively, and the parameters  $A_{ij}$ ,  $B_{ij}$ ,  $\lambda_{ij}$  and  $\mu_{ij}$  depend on the chemical species of atoms  $i$  and  $j$ . It would appear from Eq. (24) that this model is a pair-wise potential, but this is not so. The third-body dependence of the potential is contained in the  $b_{ij}$  term, which is a function of  $\theta_{ijk}$ , the angle defined by the vectors connecting one

atom with every possible pair of its neighbours. The parametrisation for silicon is such that a tetragonal arrangement of each atom's neighbours, at appropriate first-neighbour distances, minimises the energy. In the case of carbon, a second minimum at  $120^\circ$  allows the obtain also the single-layer graphite structure (graphene).

Potentials have also been developed for metallic systems, such as the Embedded Atom Model (EAM) and its derivations, or the Cleri and Rosato [21] potential. In these models there is also an environment dependence of the potential, but it is not so strongly directional as in the case of covalent materials. This model, like many of its kind, is based on the observation that the energetics of d-band metals are largely dictated by the width and centre of gravity of the d-band density of states (d-DOS), but are fairly insensitive to its detailed shape. Since the width of the d-DOS is proportional to  $\sqrt{\mu_2}$ , where  $\mu_2$  is the second moment of the d-DOS, and  $\mu_2$  can be related to the hopping integrals of the Hamiltonian, the idea is then to write down the binding energy of the system as an expression reminiscent of this. Cleri and Rosato [21] proposed to use

$$E_b^i = - \left[ \sum_j \eta^2 e^{-2q(r_{ij}/r_0-1)} \right]^{1/2}, \quad (25)$$

where  $\eta$  plays the role of a hopping integral, and  $r_0$  is the nearest-neighbour distance. The binding energy is complemented by a pairwise repulsive energy of the form

$$E_r^i = \sum_j A e^{-p(r_{ij}/r_0-1)}. \quad (26)$$

The total potential energy of the system,  $\mathcal{U}$ , is then given as the sum over all atoms of Eqs. (25) and (26) above. This model has been parametrised for a series of metals, including Ni, Cu, Rh, Pd, Ag, Ir, Pt, Au, Al and Pb [21].

Organic and biological molecules are frequently simulated with potentials of the form

$$\begin{aligned} U &= \sum_{\text{bonds}} k_{\text{bond}}(d - d_0)^2 + \sum_{\text{angles}} k_{\text{angle}}(\theta - \theta_0)^2 + \\ &+ \sum_{\text{torsions}} k_{\text{torsion}}[1 + \cos(n\chi - \delta)] + \\ &+ \sum_{ij} \left\{ 4\epsilon_{ij} \left[ \left( \frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left( \frac{\sigma_{ij}}{r_{ij}} \right)^6 \right] + \frac{q_i q_j}{r_{ij}} \right\}, \end{aligned} \quad (27)$$

where the first term is a sum over the bonds, the energy of which is modelled by a harmonic spring or some similar potential (e.g. a Morse potential). The second term accounts for bond-angle vibrations, also modelled by a harmonic spring on the deviation from the equilibrium angle. The third term describes dihedral angles, and involves sequences of four atoms linked by three adjacent bonds. The last term encompasses the energy of interaction between pairs of atoms that are not directly bonded, and it includes a Lennard-Jones type potential (see above) and a Coulomb term to account for the electrostatic interaction between charged ions. Potentials similar to that of Eq. (27) form the core of such programs as CHARMM [7] and AMBER [8].

Tight-Binding (TB) models occupy the middle ground in the spectrum of models for materials simulation, between the extremes of empirical potentials and first-principles methods. TB models, unlike empirical potentials, do incorporate a description of the electronic structure, although they do so at a much more simplistic and approximate level than first-principles methods. In TB models the matrix elements of the electronic Hamiltonian are not evaluated rigorously from the Hamiltonian operator and a chosen basis set, but rather are assumed to have a certain parametrised dependence on the interatomic distances. This makes the cost of constructing the matrix representation of the TB Hamiltonian rather small, while it is a significant part of the calculation in first-principles methods. However, this is at the cost of assuming a given form of the matrix elements, which may be physically sound, but is ultimately *ad hoc*, just as the form of an empirical potential is. In spite of this, TB models have become extremely popular in materials modelling [26], due to their combination of methodological simplicity and comparative accuracy. We will not go into details on the different TB models; interested readers may find details of these techniques in several review papers [26, 27] and books [20, 28].

As discussed in Section 3, one of the landmarks in computational condensed matter physics was the development of first-principles MD by Car and Parrinello [16] (CP). These authors obviated the need to employ a potential for modelling the atomic interactions by means of an empirical potential. Rather, the potential energy and its derivatives were directly obtained from first-principles electronic structure calculation. Specifically, CP formulated FPMD within the context of Density Functional Theory (DFT). DFT was formulated in the 1960s by Kohn and collaborators. Hohenberg and Kohn [17] demonstrated that the energy of an ensemble of electrons moving in an external field, and in particular the field generated by the nuclei or ions, is a unique functional of the electron density, and furthermore, that this functional adopts a minimum value when the electron density is that corresponding to the ground state. Subsequently, Kohn and Sham [18] showed that the electronic structure problem could be cast into an independent particle problem in which the wave functions of each particle obey a Schrödinger-like equation of the form

$$\left[ -\frac{1}{2}\nabla^2 + V_{KS}(\mathbf{r}) \right] \psi_i = \varepsilon_i \psi_i, \quad (28)$$

where  $\psi_i$  and  $\varepsilon_i$  are the particle wave functions and energies, respectively, and  $V_{KS}$  is the Kohn-Sham potential, given by

$$V_{KS}(\mathbf{r}) = V_{ext}(\mathbf{r}) + \int d\mathbf{r}' \frac{n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} + \frac{\delta E_{xc}[n(\mathbf{r})]}{\delta n[\mathbf{r}]}. \quad (29)$$

Here the first term on the rhs is the potential of interaction with the ions or nuclei, the second term is the potential due to the electrostatic interaction with the electron density, and the last term is the exchange-correlation potential. To cut a long story short, the Kohn-Sham orbitals must be obtained by self-consistently solving Eqs. (28) (note that  $V_{KS}$  depends on the  $\psi_i$  orbitals through the density  $n(\mathbf{r}) = 2 \sum_i |\psi_i(\mathbf{r})|^2$ ). Once the Kohn-Sham equations have been solved, the total energy, forces and stress can be obtained, and used in a conventional MD simulation. In order to solve Eqs. (28), it is

customary to expand the orbitals  $\psi_i$  as a linear combination of basis functions, like so:

$$\psi_i = \sum_k c_{i,k} \phi_k, \quad (30)$$

where different choices of basis functions  $\phi$  exist. The problem then is reformulated into finding out the coefficients of the expansion,  $c_{i,k}$ . This can be done by any of a number of techniques [20]. Let's imagine starting an MD simulation from an atomic configuration for which the  $c_{i,k}$  coefficients in Eq. (30) had been previously obtained. For such a configuration the total energy and forces on the ions are available once the electronic structure problem is solved, so one can use these forces to perform a conventional MD step on the ions. Once the ions move, however, in principle one would have to go back and solve the electronic problem all over again. But CP proposed to do something different: they showed that it is possible to incorporate the  $c_{i,k}$  expansion coefficients as fictitious classical variables in the dynamics, with adequately chosen fictitious masses. Thus one ends up with a combined dynamics of ions and wave function coefficients. This looks very strange indeed, but it is in fact a very clever trick: with a suitable choice of fictitious masses for the  $c_{i,k}$  coefficients and a small enough time step, it is possible to arrange things in such a way that the dynamics of the  $c_{i,k}$  follows closely the Born-Oppenheimer surface, or in other words, the  $c_{i,k}$  automatically adapt to the slowly varying ionic configuration, giving wave functions that are very close to the Kohn-Sham ground state for the current ionic configuration. The fictitious masses of the  $c_{i,k}$  need to be small enough so that their dynamics is faster than that of the ions; this in turn imposes a smaller time step than would be required for a stable dynamics of the ions with a conventional force field, but the gain is that one has done away with the force field altogether!

Another consideration to take into account is that the dynamics of the  $c_{i,k}$  must be subject to the constraints of orthonormality of the Kohn-Sham orbitals, i.e.

$$\sum_k c_{i,k}^* c_{j,k} = \delta_{ij}, \quad (31)$$

where  $\delta_{ij}$  is the Kronecker delta. There are standard techniques for performing MD subject to constraints<sup>4</sup> which can be applied to impose Eq. (31) at each time step. Imposing such constraints ( $N^2$  of them, where  $N$  is the number of Kohn-Sham orbitals) is a significant bottle-neck of FPMD, as this carries a computational cost that grows as  $\mathcal{O}(N^3)$ .

There are many intricacies in CPMD and DFT calculations, which we cannot cover here in any detail, but fortunately all this is extensively documented in the literature (see e.g. [20, 29, 30, 31]). Suffice it to say that FPMD in the CP flavour and in others is now a fairly standard and frequently used simulation technique. It is still computationally very demanding compared to MD based on force fields, but the cost is affordable in many cases, thanks in part to the continuing improvement of algorithms and numerical techniques, and to the ever increasing tendency of computer power.

---

<sup>4</sup> Two well-known algorithms for imposing constraints are the so-called *rattle* and *shake* methods. We will not discuss them here, but interested readers will find full accounts in refs. [3, 4].

## 7. BEYOND THE MICROCANONICAL (NVE) ENSEMBLE

Let us now briefly discuss how MD has been extended beyond microcanonical conditions, so as to simulate systems in mechanical and thermal contact with their surroundings. As mentioned in Sec. 3, the first work to consider the possibility of performing MD simulations in conditions of constant pressure was that of Andersen [12]. Andersen proposed to couple the dynamics of the atoms with that of the volume,  $\Omega$ , of the system, in such a way that they would be both described by the following Lagrangian:

$$\mathcal{L}_{\text{Andersen}} = \frac{1}{2} \sum_i m_i \Omega^{2/3} \dot{\mathbf{q}}_i^2 - \mathcal{U}(\Omega^{1/3}, \{\mathbf{q}\}) + \frac{1}{2} m_\Omega \dot{\Omega}^2 - P_{\text{ext}} \Omega. \quad (32)$$

The first two terms here are the usual kinetic and potential energy of the atoms, but now rewritten in such a way as to make their dependence on the volume of the system explicit. Note that instead of the usual Cartesian positions for the atoms,  $\mathbf{r}_i$ , we have now used the scaled, or lattice, coordinates,  $\mathbf{q}_i = \Omega^{-1/3} \mathbf{r}_i$ , which are more convenient in this case. The third and fourth terms in Eq. (32) above correspond to the kinetic and potential energies for the volume, which is now itself a dynamical variable.  $m_\Omega$  is the thermostat fictitious mass, and  $P_{\text{ext}}$  is the external pressure which is exerted on the system. If the volume was to stay fixed, its kinetic energy would be zero (no volume momentum), and the volume potential energy would be constant. In this case  $\mathcal{L}_{\text{Andersen}}$  reduces to the conventional microcanonical Lagrangian for the atoms. However, when the volume is free to move, it will react to the external pressure, increasing or decreasing as dictated by the pressure and the combined dynamics of atoms and volume. The volume will eventually settle and oscillate around an average value, which will be the mean volume for the imposed external pressure. It is easy to make the transition from the Lagrangian formulation of Eq. (32) to the Hamiltonian form, by simple application of the usual rules of classical mechanics [32], with momenta defined as

$$\mathbf{p}_i = \frac{\partial \mathcal{L}}{\partial \dot{\mathbf{q}}_i} \quad (33)$$

and defining the Hamiltonian function

$$\mathcal{H} = \sum_i \dot{\mathbf{q}}_i \cdot \mathbf{p}_i - \mathcal{L}. \quad (34)$$

It is a useful exercise for the reader to transform Eq. (32) to Hamiltonian form and then use the generalised leapfrog scheme discussed in 4.1 to obtain a constant pressure integration algorithm.

Andersen's approach only considers volume fluctuations, i.e. the size of the simulation box is allowed to change, but its shape is constrained to remain cubic. This is ok for liquids, but for crystalline solids it is actually restrictive. If a undergoes a phase transition to another solid phase, it will in general change not only its cell volume, but also its shape. In order to account for such situations, Parrinello and Rahman [15, 33] generalised the method of Andersen with the following Lagrangian:

$$\mathcal{L}_{\text{PR}} = \frac{1}{2} \sum_i m_i \dot{\mathbf{q}}_i \mathbf{H}^t \cdot \mathbf{H} \dot{\mathbf{q}}_i - \mathcal{U}(\mathbf{q}, \mathbf{H}) + \frac{1}{2} M_H \text{Tr}(\dot{\mathbf{H}}^t \dot{\mathbf{H}}) - P_{\text{ext}} V. \quad (35)$$

This can be seen to be somewhat similar to Andersen's Lagrangian, Eq. (32), though there are some differences. The key difference is that now, instead of the volume, the cell degrees of freedom are the components of the vectors defining the shape of the simulation box. These vectors can be arranged into a matrix which is labelled as  $\mathbf{H}$ ; the Cartesian coordinates of the atoms are  $\mathbf{r}_i = \mathbf{H}\mathbf{q}_i$ . In Eq. (35) a fictitious kinetic energy term which now includes the kinetic energy of each of the cell vector components. There are several subtleties here, not least the fact that the fictitious dynamics of the cell in the formulation of Parrinello and Rahman allows not only for the deformation of the cell, but also for changes in its orientation in space. This is certainly rather inconvenient in simulations, as one would have to somehow distinguish between the motion of the atoms which is intrinsically due to the atomic dynamics, and that which results from cell rotation. There are cures for this, but we do not need to concern ourselves with such technicalities here [34, 35]. Again, it is possible to transform Eq. (35) to Hamiltonian form, to obtain the equations of motion from the Hamiltonian, and with them derive an algorithm for their numerical integration using the generalised leapfrog scheme. This would be a recommended exercise for anyone who wants to become familiar with MD techniques.

In his famous paper, Andersen [12] proposed not only a way to perform simulations in conditions of constant pressure (i.e. the isobaric-isoenthalpic, or NPH ensemble), but also in conditions of constant temperature (canonical or NVT ensemble). His strategy to attain canonical sampling consisted of selecting an atom at random, and changing its velocity to a new value, selected from the Maxwell-Boltzmann distribution corresponding to the desired temperature of the simulation. This process, repeated at regular intervals, was shown to sample the canonical ensemble, and is now known as Andersen's thermostat. However, there is one key difference between Andersen's procedure for sampling at constant pressure and that for constant temperature, and it is that while the first one is deterministic (i.e. the barostat obeys a certain equation of motion), Andersen's thermostat is stochastic. In this sense it is a bit like introducing ingredients of MC into MD. This is not a bad thing in itself, provided one is not interested in the dynamical properties of the system, such as transport properties. If this is, however, the case, one must be aware of the fact that the stochastic nature of Andersen's thermostat affects the dynamics of the atoms, and will cause an artificially rapid loss of coherence in their dynamics. In other words, the VAF (see 5.4.2) will decay faster than it would otherwise do, and this clearly affects the value of e.g. the diffusion coefficient [see Eq. (19)]. In general, if you are interested in obtaining dynamical information of the studied system, my recommendation would be to avoid the use of thermostats and other artifacts, which may affect the dynamics of the atoms in the system, and to stick to NVE sampling. This is not to say that thermostats and barostats are not useful; far from it! there are many situations in which one is not really interested in the atomic dynamics, and in which one needs to simulate systems in conditions of constant pressure and/or temperature. In such cases the use of these artful devices is highly recommended.

Andersen's paper was the starting point of a plethora of different methods to allow MD to sample other ensembles. Among these, a key development was that of Nosé [13], who proposed a new thermostat to achieve canonical sampling. Contrary to that of Andersen, though, Nosé's thermostat has the peculiarity of being deterministic, instead



of stochastic. Nosé proposed the following Lagrangian,

$$\mathcal{L}_N = \frac{1}{2} \sum_i m_i s^2 \dot{\mathbf{r}}_i^2 + \frac{1}{2} m_s \dot{s}^2 - U(\{\mathbf{r}\}) - g k_B T_{\text{ext}} \ln s, \quad (36)$$

where  $g$  is the number of degrees of freedom of the system,  $k_B$  is Boltzmann's constant,  $s$  is the thermostat *position* variable,  $m_s$  is its associated fictitious mass, and  $T_{\text{ext}}$  is the temperature at which we desire the system to be. For once, we will miss the opportunity of proposing an exercise for the reader, and give directly the Hamiltonian function for the Nosé thermostat; this is

$$\mathcal{H}_N = \sum_i \frac{\mathbf{p}_i^2}{2m_i s^2} + U(\{\mathbf{r}\}) + \frac{p_s^2}{2m_s} + g k_B T_{\text{ext}} \ln s, \quad (37)$$

where the atomic momenta are  $\mathbf{p}_i = m_i s^2 \dot{\mathbf{r}}_i$ , and similarly,  $p_s = m_s \dot{s}$ . I will, however, suggest to the reader to derive equations of motion from Nosé's Hamiltonian Eq. (37). Then, it will be seen that the force acting on the thermostat is proportional to the difference between the kinetic energy of the atoms and the  $g k_B T_{\text{ext}}/2$ ; in words, what this means is that the thermostat variable  $s$  is going to increase when the kinetic energy grows to values above that corresponding to  $T_{\text{ext}}$ , and the opposite will happen when the kinetic energy is below the target value. In this way, one ensures that the average kinetic energy of the system, during a sufficiently long run, will correspond to the correct value at  $T_{\text{ext}}$ . Not only this, but it can actually be shown that because of the chosen form of the potential energy for the thermostat in Eqs. (36) and (37), the dynamics of the atoms samples the canonical ensemble, under the usual assumption of ergodicity.

The thermostat variable  $s$  in Eq. (36) is actually a time-scaling factor. The *real* time of the simulation is actually given by

$$t_{\text{real}} = \int_0^t \frac{dt}{s(t)}, \quad (38)$$

which means that the actual length of the time step varies during the simulation. This is somewhat inconvenient, particularly if one wishes to calculate time dependent properties of the system. Motivated by this, Hoover [14] modified Nosé's original formulation through a change of variables which resulted in a scheme in which the length of the time step is constant. The resulting scheme is frequently referred to in the literature as the Nosé-Hoover thermostat. There is one small caveat with this, though, which is perhaps of little practical significance, but it is worth commenting. Hoover's modifications amounted to a non-canonical (non-Hamiltonian) transformation of Eq. (37), and as a result of this the method loses its Hamiltonian structure. This means that the Nosé-Hoover equations of motion cannot be derived from a Hamiltonian, and as a consequence, one cannot use the generalised leapfrog scheme to obtain an integration algorithm. This is not such a serious problem, because alternative integration schemes can be used which do not rely on having a Hamiltonian structure (see e.g. [4]).

In more recent times Bond *et al.* [36] have shown that it is indeed possible to reformulate the Nosé thermostat by means of a canonical transformation (thus respecting the

underlying Hamiltonian structure), a transformation which is designed to counteract the troublesome time scaling implicit in Nosé’s original formulation. They did this by using a so-called Poincaré transformation, resulting in a new Hamiltonian given by:

$$\mathcal{H}_{\text{NP}} = s(\mathcal{H}_{\text{N}} - H_0), \quad (39)$$

where  $H_0$  is a constant, and  $H_{\text{N}}$  is given by Eq. (37). Because this is a Hamiltonian, one can use the generalised leapfrog scheme, and this would be my recommended option for canonical sampling MD.

The use of new variables such as the barostats and thermostats discussed above has been called the *extended system approach*. These extended variables are designed in such a way that they emulate the effect of having the system placed in contact with its surroundings, i.e. with an essentially infinite number of degrees of freedom. It is quite remarkable that one can do this with just a few additional degrees of freedom. But in introducing these artificial variables one must assign values to their corresponding fictitious masses [ $m_{\Omega}$  and  $m_s$  in Eqs. (32) and (36), respectively]; the dynamics of the extended variables and to some extent that of the atoms of the system do depend on the values chosen for these fictitious masses, and some tuning of their values may be necessary in order to achieve sensible results. The values of average properties are in general not very sensitive to the values of  $m_{\Omega}$  and  $m_s$ ; their choice should be guided by an efficiency of sampling, while at the same time trying to affect minimally the dynamics of the atoms.

Finally, before leaving this section, let us remark that both Andersen [12] and Nosé [13] considered sampling the isothermal-isobaric or NPT ensemble by simultaneously coupling the system to a thermostat and a barostat. This combination is also considered by Sturgeon and Laird [37] and Hernández [35].

## 8. PROBLEMS, CHALLENGES, ..., AND ALL THAT!

It would not be fair to conclude these introductory notes on MD without making some comments about the limitations of MD, which indeed exist and are not few. The most obvious one is the issue of time scales. Depending on the level at which you model your system (first-principles or empirical force field) MD may be limited to time scales ranging from a few pico-seconds to up to a few nano-seconds at most. Yet many processes of chemical and physical interest happen over time scales which can be many orders of magnitude larger than this (slow diffusion problems in solids, dynamics of glassy or polymer systems, or protein folding, to name but a few), and straightforward MD simply can get you nowhere in such cases. In recent years, Voter and others have developed several techniques to try to address this problem, such as Hyperdynamics [38, 39], Temperature-accelerated dynamics [40] or the Parallel-replica method [41].

One of the reasons why the time scale that can be covered is limited is related to the computational cost involved in calculating the energy and forces necessary to perform MD simulations. It has been recently suggested [42] that specially designed neural networks may be trained to predict with sufficient accuracy the energetics and forces of a given system after being fed with a sufficient training data set obtained from simulations. This could potentially reduce significantly the cost of performing accurate simulations,

and thus also extend the length of time scales accessible with such simulations. However, this methodology is still rather new, and its full potential is yet to be demonstrated.

MD is ultimately a sampling technique, like MC, with the added bonus of providing dynamical information, at the extra cost of calculating the forces. Systems with complicated energy landscapes are inherently more problematic to sample adequately, so special care has to be taken in such cases. In such systems one may have to wait for a long time for the dynamics to explore the configuration space. To ameliorate this problem, Parrinello and co-workers have proposed the technique known as metadynamics [43, 44]. In this technique, a dynamical trajectory is followed which is discouraged from visiting regions of configuration space that have already been visited by adding a Gaussian potential of a pre-specified height and width to each visited point. In this way potential energy minima are gradually filled up, facilitating the escape of the system from such trapping funnels, and improving the configuration space sampling.

Yet another challenge for MD is the issue of varying length-scales. In many systems, the phenomena under observation cover many different length scales. A typical example of this is crack propagation, where a material is loaded (stressed) until a crack tip forms and starts to propagate in the material. Close to the crack tip, chemical bonds are being broken, and atoms strongly rearrange. A bit further away from the tip the material may be severely deformed, but without bonds being actually broken, and yet further away the atomic positions may hardly deviate from those in the perfect crystal. To model such systems directly at the atomistic level requires extremely large simulation cells, soon growing into six orders of magnitude figures and beyond. For tackling such problems effectively it is necessary to treat different length scales at different levels of theory, effectively embedding a quantum mechanical description of the tip crack into a force field description for the atoms a certain distance away from the crack. This in turn must be matched at some point with a continuum mechanics description, valid for large length scales. A similar situation is encountered e.g. in enzymatic reactions, where the active site of the protein and the reactants must be described at a quantum mechanical level, while the rest of the protein and perhaps the solvent (typically water) may be accounted for at a lower level of theory.

To summarise, both challenges and exciting times lie ahead; MD in particular, and simulation in general, are very open fields, in constant evolution, and responding to the new issues which are continuously being raised by experimental progress in the physics and chemistry of materials and by nanotechnology. I have no doubt that very exciting times lie ahead in this field, a field full of opportunities for unfolding a productive and fulfilling career in science.

## **ACKNOWLEDGEMENTS**

Many of the results of MD simulations used here as illustrations of the different aspects of the technique have been conducted in collaboration with or directly by my former students O. Natalia Bedoya-Martínez, Marcin Kaczmarek and Riccardo Rurali. I also wish to thank the organisers of this summer-school. My work is funded by the Spanish Ministry of Science and Innovation through grants FIS2006-12117-C04 and TEC2006-12731-C02-01 and by the AGAUR through grant 2005SGR683.

## REFERENCES

1. N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (2005).
2. B. J. Alder, and T. E. Wainright, *J. Chem. Phys.* **27**, 1207–1208 (1957).
3. M. P. Allen, and D. J. Tildesley, *Computer Simulation of Liquids*, Clarendon Press, Oxford, 1987.
4. D. Frenkel, and B. Smit, *Understanding Molecular Simulation*, Academic Press, San Diego, 1996.
5. J. M. Thijssen, *Computational Physics*, Cambridge University Press, Cambridge, 1999.
6. D. M. Ceperley, *Rev. Mod. Phys.* **67**, 279–355 (1995).
7. B. R. Brooks, R. E. Bruccoleri, B. D. Olafson, D. J. States, S. Swaminathan, and M. Karplus, *J. Comput. Chem.* **4**, 187–217 (1983).
8. D. A. Case, T. E. Cheatham, T. Darden, H. Gohlke, R. Luo, K. M. Merz, A. Onufriev, C. Simmerling, B. Wang, and R. Woods, *J. Computat. Chem.* **26**, 1668–1688 (2005).
9. G. V. Lewis, and C. R. A. Catlow, *J. Phys. C: Solid State Physics* **18**, 1149–1161 (1985).
10. B. J. Lee, J. H. Shim, and M. I. Baskes, *Phys. Rev. B* **68**, 144112–1,11 (2003).
11. J. Tersoff, *Phys. Rev. B* **37**, 6991–7000 (1988).
12. H. C. Andersen, *J. Chem. Phys.* **72**, 2384–2393 (1980).
13. S. Nosé, *J. Comput. Phys.* **81**, 511–519 (1984).
14. W. G. Hoover, *Phys. Rev. A* **31**, 1695–1697 (1985).
15. M. Parrinello, and A. Rahman, *Phys. Rev. Lett.* **45**, 1196–1199 (1980).
16. R. Car, and M. Parrinello, *Phys. Rev. Lett.* **55**, 2471–2474 (1985).
17. P. Hohenberg, and W. Kohn, *Phys. Rev.* **136**, 864 (1964).
18. W. Kohn, and J. J. Sham, *Phys. Rev.* **140**, 1133 (1965).
19. J. M. Sanz-Serna, and M. P. Calvo, editors, *Numerical Hamiltonian Problems*, Chapman and Hall, New York, 1995.
20. R. M. Martin, *Electronic structure: basic theory and practical methods*, Cambridge University Press, Cambridge, 2004.
21. F. Cleri, and V. Rosato, *Phys. Rev. B* **48**, 22–33 (1993).
22. E. R. Hernández, and J. Íñiguez, *Phys. Rev. Lett.* **98**, 055501–(1,4) (2007).
23. J. P. Hansen, and I. R. McDonald, *Theory of simple liquids*, Elsevier, London, 1986, 2nd edn.
24. O. N. Bedoya-Martínez, M. Kaczmariski, and E. R. Hernández, *J. Phys.: Condens. Matter* **18**, 8049–8062 (2006).
25. F. H. Stillinger, and T. A. Weber, *Phys. Rev. B* **22**, 3790–3794 (1980).
26. C. M. Goringe, D. R. Bowler, and E. R. Hernández, *Rep. Prog. Phys.* **60**, 1447–1512 (1997).
27. L. Colombo, *Nuovo Cimento* **28**, 1–59 (2005).
28. A. P. Sutton, *Electronic Structure of Materials*, Clarendon Press, Oxford, 1993.
29. G. Kresse, and J. Furthmüller, *Phys. Rev. B* **54**, 11169–11186 (1996).
30. M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, *Rev. Mod. Phys.* **64**, 1045–1097 (1992).
31. D. K. Remler, and P. A. Madden, *Mol. Phys.* **70**, 921–966 (1990).
32. H. Goldstein, *Classical Mechanics*, Addison-Wesley, Reading, Massachusetts, 1980.
33. M. Parrinello, and A. Rahman, *J. Appl. Phys.* **52**, 7182–7190 (1981).
34. I. Souza, and J. L. Martins, *Phys. Rev. B* **55**, 8733–8742 (1997).
35. E. Hernández, *J. Chem. Phys.* **115**, 10282–10290 (2003).
36. S. D. Bond, B. J. Leimkuhler, and B. B. Laird, *J. Comput. Phys.* **151**, 114–134 (1999).
37. J. B. Sturgeon, and B. B. Laird, *J. Chem. Phys.* **112**, 3474–3482 (2000).
38. A. F. Voter, *Phys. Rev. Lett.* **78**, 3908–3911 (1997).
39. A. F. Voter, *J. Chem. Phys.* **106**, 4665–4677 (1997).
40. M. R. Sorensen, and A. F. Voter, *J. Chem. Phys.* **112**, 9599–9606 (2000).
41. A. F. Voter, *Phys. Rev. B* **57**, 13985–13988 (1998).
42. J. Behler, and M. Parrinello, *Phys. Rev. Lett.* **98**, 146401–(1,4) (2007).
43. A. Laio, S. Bernard, G. L. Chiarotti, S. Scandolo, and E. Tosatti, *Science* **287**, 1027–1030 (2000).
44. R. Martoňák, A. Laio, and M. Parrinello, *Phys. Rev. Lett.* **90**, 075503–1,4 (2004).