

# An introduction to atomistic simulation methods

***Martin T. Dove***

Department of Earth Sciences, University of Cambridge, Downing Street, Cambridge  
CB2 3EQ, UK

*Various methods for simulating materials and minerals at an atomic level are reviewed, including lattice energy relaxation, lattice dynamics, molecular dynamics and Monte Carlo methods, and including also the use of empirical interatomic potentials and ab-initio quantum mechanical methods. A small number of diverse applications are described. Approaches to job and data management are also discussed.*

## **1. Introduction**

Over the past two decades or more all sciences have seen an explosion of the use of computer simulations to the point where computational methods are now stand alongside theoretical and experimental methods in value. The birth of the use of computer simulations was actually around five decades ago, but their impact in modern science has exactly mirrored the exponential growth in the power of computers and the use of computers across the whole of science (e.g. for instrument control). In turn, the growing power of computers has spurred the development of methods and code interfaces, widening the potential of simulations to tackle a wide range of scientific issues and placing tools in the hands of a wider group of scientists. Let me illustrate this with a simple example. 15 years ago using a computer system that cost the equivalent of 20 desktop computers in modern equivalent terms I could calculate the energy-relaxed structure of cordierite ( $\text{Mg}_2\text{Al}_4\text{Si}_5\text{O}_{18}$ ) in around 3 hours. The same calculation on my desktop has just taken me 1.5 s. A study of energy modelling of leucite ( $\text{KAlSi}_2\text{O}_6$ ) (Dove et al., 1993) from 15 years ago that involved computing the energy a number of different configurations required the use of a powerful vector computer, but today the same calculations would be trivial to perform. 15 years ago setup time would be considered to be trivial compare to the job run time, and thus no-one worried about program interfaces. Today I expect to be able to download a file containing an atomic configuration and have it running in a simulation within a very short period of time.

In this introductory chapter I aim to set the scene for the range of applications discussed in the following chapters by outlining some of the key technologies now being used in computational materials and mineral sciences. I will begin by reviewing and comparing the relative benefits of the two most common approaches to modelling the interactions between atoms, namely using empirical representations and ab initio quantum mechanical methods. Both approaches can then be used in each of the four methods that are reviewed in subsequent sections, namely lattice energy modelling, calculations of lattice dynamics, the molecular dynamics simulation method, and the Monte Carlo simulation method (focussing on the Metropolis method). Throughout this review I cite a number of representative applications, and I follow the review of the methods with a small number of brief case studies. I conclude by enlarging the discussion to consider practical issues of running simulations on modern computing resources, and of the critical role of data management.

## 2. Empirical models for forces between atoms

### 2.1. Functions

The use of empirical models for the interactions between atoms depends critically on there being a set of simple equations that give a reasonable description of the dependence on interatomic separation or other details of local atomic coordination, and there being the possibility to tune the parameters in the equation. One model that has worked well for ionic materials is to use long-range Coulomb interactions with short-range repulsive interactions of the form:

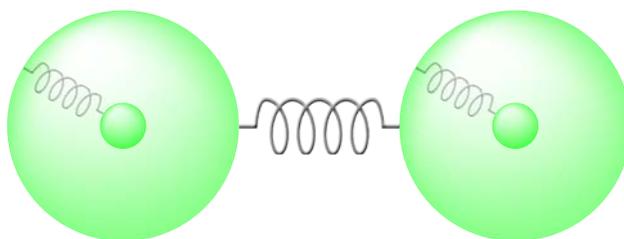
$$E(r) = A \exp(-r / \rho) \tag{1}$$

where  $r$  is the interatomic spacing, and  $A$  and  $\rho$  are parameters whose values are to be tuned against some data. This specific function was introduced by Born and Mayer as long ago as 1932 (Born and Mayer, 1932). This model has worked well for many minerals, as reviewed by Burnham (1990), and appears to have some justification when compared with the results of quantum mechanical calculations. However, there are cases where this is an over-simplification to an extent that renders it inappropriate. One example is where there are important angular forces such as might be expected for ions that favour certain atomic coordination. One example of the use of bond-bending terms is to model the tetrahedral  $\text{SiO}_4$  coordination found in many silicates, and the use of a simple harmonic function of bond angle (Sanders et al., 1984) often suffices:

$$E(\theta) = \frac{1}{2} K (\theta - \theta_0)^2 \tag{2}$$

where  $\theta_0$  is the nominal equilibrium angle ( $90^\circ$  for an octahedral coordination,  $109.47^\circ$  for a tetrahedral coordination). Other functions, such as more complicated distance-dependent pair terms, or multi-atom terms, can be incorporated into a model, provided that there might be some physical or chemical justification. Terms that merely represent the first few terms in a Taylor expansion, such as the harmonic bond-bending term, are clearly justified because they are little more than a numerical approximation. Terms of the form  $-Cr^{-6}$  can be justified as representing the known dispersive interactions. A wide range of models has been developed, and are incorporated within most standard modelling codes: e.g. GULP (Gale, 1997; Gale and Role, 2003) and DL\_POLY (Smith, 2006; Todorov and Smith, 2004).

The models outlined above assume that ions are rigid in the sense that they cannot be deformed (although clearly the models allow some degree of spatial overlap of the spatial extent of individual ions and do not treat ions as perfectly hard spheres). This means that the models do not allow for ionic polarisation in response to local electrical fields, whether static or dynamic. In turn, this means that such models are unable to be used to calculate the refractive index, and it follows that it will not be possible to obtain the correct differences between the frequencies of longitudinal and transverse optic phonons. The simplest solution to this problem is to introduce the shell model, in which the ion is represented by an inner charged core with the full mass of the ion, and an outer mass-less shell. The simplest version of the shell model has a harmonic force between the centres of the core and shell that opposes their separation, and does not allow for any deformation of either core or shell; this is illustrated in Figure 1.



**Figure 1.** Representation of the shell model, showing the inner cores containing the whole of the ionic mass and the outer mass-less shells. The core and shell on any ion interact via a potential energy function that assume an equilibration separation of zero. In simplest models this is a harmonic interaction, but it can be extended to higher powers of separation. Moreover, the model can be extended to allow the shell to be deformable either in shape or radius. In most applications neighbouring atoms interact through short-range potentials that act on shells only (in addition to the Coulomb interactions). From Dove (2003).

## 2.2. Tuning empirical functions

The existence of usable and appropriate functions is only part of the issue; a greater challenge is whether the parameters in the function can be tuned to appropriate values (Gale, 1996). The traditional approach, anticipating the use of methods to be described later, is to tune the parameters such that they lead to the prediction of the crystal structure that is in closest agreement with the experimental structure. Formally this procedure is carried out by adjusting the parameters in the model until all computed stresses on the crystal and forces on the atoms are as close to zero as is possible. This approach can suffer from the fact that frequently there are more parameters in the model than the number of independent structure variables; in the case of a simple cubic binary material such as MgO there is only one structural parameter. The paucity of data may be circumvented by keeping some of model parameters at fixed values, such as using formal charges and by only allowing one of the two parameters in the Born–Mayer term to be tuned.

A more refined approach is to incorporate experimental data. For example, and again anticipating later discussion, it is possible to calculate physical properties such as dielectric and elastic constants, and vibrational frequencies. Thus not only can the model be tuned to minimise the stresses and forces, it can also be tuned to simultaneously give closest agreement with the experimental data. In fact this approach is essential when using shell and bond-bending models, because invariably these components of the model are often not probed in the details of the structure.

In the approach I have outlined, the result might be a well-tuned model that is capable of reproducing a range of experimental data. However, there is one caveat that needs to be appreciated. By fitting against crystal structure and physical properties we have tuned the model so that it best represents the first and second differentials of the interatomic potential energy function at the positions of the interatomic distances. There is no guarantee that it gives a good estimate of real bond energies, nor that it will work well for different interatomic distances.

Let us consider two examples. First, we might have a well-tuned model that reproduces the structure and properties of  $\text{Al}_2\text{MgO}_4$  spinel. However, if we then calculate the energy change when swapping some of the Al and Mg cations, it is highly unlikely that the model will properly reproduce the energy associated with the change in coordination number of the two types of cation (Al is octahedrally coordinated in the spinel structure, and Mg is tetrahedrally coordinated) if there is any degree of covalent bonding in the crystal. Second, take the same example and perform a calculation under pressure. If the model correctly reproduces the elastic constants, there is a chance that

the low-pressure results will be reasonable. But once the bond lengths have changed appreciably, there is no guarantee at all that the model will match experiment because, quite simply, the model has not been tuned to reproduce the energy function for these new interatomic distances. The take-home message here is to not trust simulations using empirical models at pressures for which the models were not specifically tuned.

An emerging approach to tackle this problem is to tune models against the predictions of quantum mechanical calculations (again, see below). Early approaches in this direction were to compute the interactions between pairs of ions as a function of separation and fit the resultant energy curve (Post and Burnham, 1986). A refinement on this was to compute the energy of a small cluster of atoms for different configurations, and fit the resultant energy surface (i.e. the energy as a function of many variables) (Tsuenyuki et al., 1988). The trend now is to obtain the quantum mechanical energy surface for many configurations of the crystal and use this as the primary source of data; this approach enables the use of many independent data values and thus facilitates the tuning of more complicated models for a wider range of interatomic separations. Furthermore, this approach also enables the models to incorporate bond energies directly, and would enable a reliable calculation of the energy changes due to swapping ions with different coordination number.

### 2.3. *Transferability*

One of the big areas of success for empirical modelling is organic materials, particularly with applications for pharmaceuticals. The success relies on the observation that models for standard non-bonded atom pairs, such as C...C and C...H, and also for internal molecular distortions (such as flexing of C–C–C bonds) can be modelled using a set of standard functions without the need to retuned the model for each new application. Such a model is said to be “transferable”, and this feature underpinned the original growth of energy modelling.

For modelling of aluminosilicates we are fortunate that we can also use transferable models (Dove, 1989; Winkler et al., 1991; Price et al., 1987), although it has to be admitted that less effort has gone into tuning transferable models for these systems than was given to the models for organic crystals.

One model that is commonly used, at least in the UK materials/modelling community, has its roots in an initial model for quartz, SiO<sub>2</sub> (Sanders et al., 1984). This model used formal ionic charges, a shell model for oxygen anions, Born–Mayer repulsive interactions, dispersive Si...O and O...O interactions (although the term used for Si...O may not be doing much than accommodating residual errors in the overall

Si...O interaction), and an O–Si–O bond-bending term. It has been extended by including interactions for  $\text{Al}^{3+}$  in both tetrahedral and octahedral coordination (the same parameters appear to work well in both cases), and other cations can be incorporated using Born-Mayer short-range potentials (parameters can be taken from the Modified Electron Gas calculations of Post and Burnham, 1986). Furthermore, it has proven to be straightforward to incorporate hydrogen into these models (Price et al., 1987).

The transferable model described here has been successfully used for a wide range of different systems (Dove, 1989; Winkler et al., 1991; Price et al., 1987), including rock-forming minerals (Patel et al., 1991; Dove and Redfern, 1997), clays and other layer silicates (Palin et al., 2001; Sainz-Díaz et al., 2001), and zeolites (Ohsuna et al., 2004), with applications in mineralogy and materials chemistry.

### 3. Quantum mechanical methods

#### 3.1. General notes

Quantum mechanical methods hold out the possibility of performing simulations that can give reliable results without the prior need for tuning. In recent years the rapid growth of computer power coupled with recent methodological developments have transformed this area of simulation science from the domain of the specialist into the hands of a wider range of researchers.

The challenge in developing quantum mechanical methods is that we are dealing with equations for which there is no solution. The wave function properly describes all electrons within a system, which for crystalline materials really means the primitive unit cell because of translational symmetry. However, the electrons interact with one another, and we cannot easily separate out equations for each electron. In fact it is possible to write the overall wave function as a product of single-electron wave functions (the Slater determinant) and then recast the underlying Schrödinger equation in terms of the component single-particle wave functions. In so doing, the theory generates new terms from the electron–electron interactions that are divided into the ‘exchange term’ (the term generated within the Hartree–Fock approach) and the ‘correlation term’ (everything else that is generated by this approach and not included within the exchange term).

Different levels of theory handle the electron–electron terms in different ways and to different degrees of sophistication, but there are some common approaches regardless. One feature is the need to use trial functions for the wave functions – known as the “basis set” – rather than being able to derive analytical wave functions from scratch. For

methods that are focussed on the atoms, including quantum-chemical (molecular) methods, the basis sets are often designed to model atomic electron orbitals, but for methods designed to model crystalline materials the existence of translational symmetry is exploited through the use of plane waves as the basis set. The other common approach is to treat the nuclei as classical objects whose role is to generate electric fields that have the same status as an external field; the electrons move within the electric field of the nuclei, but no wave function is assigned to the nuclei. This is known as the “Born–Oppenheimer approximation”.

### 3.2. Density Functional Theory and modern applications

Most solid-state applications are now making heavy use of Density Functional Theory (DFT) (Payne et al., 1992; Segall et al., 2002). In this approach, the primary equations are cast in terms of the electron density rather than on the wave functions. However, although the theory is well-developed, it remains the case that for certain components (particularly the exchange and correlation terms) exact functionals are not available. Thus approximations are required. One common approach is to use numerical approximations to the exchange–correlation energy taken from simulations of a gas of electrons. The “Local Density Approximation” (LDA) assumes that the exchange–correlation energy for a point in the crystal can be replaced by the corresponding energy for a gas of electrons with the same electron density as at that point. A more sophisticated approach is to take account of the local gradient in the electron density – the so-called “Generalised Gradient Approximation” – but this approach tends to over-compensate for errors in the LDA.

DFT has been used for a wide range of atomic basis sets. The early major DFT codes for periodic systems, including CASTEP (Segall et al., 2002), VASP<sup>1</sup> (Vienna ab-initio simulation Package, 2007), ABINIT (Gonze et al., 2002) and WIEN2k (Schwarz and Blaha, 2003), use plane wave representations of the electrons (see below for comments on the approach taken by codes such as WIEN in this regard). Plane waves are actually very good for basis sets because they have a high degree of flexibility, provided that sufficient numbers of waves are used to model fine details in the variation of the electron density. However, plane-wave methods do not scale very well with system size, typically scaling between  $N^2$  and  $N^3$ , where  $N$  is the number of atoms. In the quest for a method that scales linearly with  $N$ , the emphasis is on using point-centred atomic orbitals. The SIESTA code, for example, uses numerical atomic orbitals as the basis sets (Soler, et al., 2002). This approach enables simulations of larger sample size than is

---

<sup>1</sup> Vienna ab initio simulation package (2007). <http://cms.mpi.univie.ac.at/vasp>.

possible using plane waves, but the cost is that results will be constrained by the flexibility of the basis set.

One challenge is how to model the electrons near the centres of the atoms. In this region of space, the wave functions change rapidly with distance from the atomic centre, and such a variation can be very hard to model. The so-called all-electron codes such as WIEN2k (Schwarz and Blaha, 2003) switch from plane waves to the use of atomic orbitals within a defined volume encompassing the atom centres. But many DFT codes make the approximation of replacing the inner electrons by an effective potential seen by the valence electrons, the so-called “pseudopotential approximation”. The key advantage of this approach is computational speed.

## 4. Lattice energy calculations

### 4.1. Fundamental idea

We now turn to discussing the tools that are used to provide the interface between models (whether empirical or quantum mechanical) and scientific applications. The first of these is lattice energy modelling. In this tool, the lattice parameters and atomic coordinates are systematically adjusted until the energy (computed either from an empirical or quantum mechanical model) is minimised (i.e. to that there are no residual stresses or forces). Typically the process involves a non-linear regression approach. The approach can easily be extended to incorporate pressure; in this case the quantity to be minimised is the lattice enthalpy.

The primary outputs of a lattice energy minimisation are the crystal structure and associated energy. One example where this approach can provide new information is when the energies of different phases are to be compared; although neither empirical nor quantum mechanical models will give useful absolute values of energy, calculations on related structures will give reliable energy differences. Thus these models might be used to predict phase stability at different pressures. Moreover, the evolution of the structure under pressure can be investigated, and the results might provide some rationalisation of the predicted phase stability.

We remark here that a good model should be capable of calculating lattice parameters and interatomic distances to within 2% of the experimental values, with the anticipation that one should usually achieve better than this. In empirical modelling the residual errors typically come from using over-simplified models, and the fact that any attempt to improve the accuracy in the computed structure will inevitably come at the expense of the accuracy of computing other properties. In quantum mechanical

modelling, residual errors reflect the limitations of approximations used in the model. For example, it is generally found that the computed interatomic distance within the GGA are larger than experimental values, but the LDA tends to give shorted computed distances.

Energies per se may not be immediately useful in many cases, but one example of an application for which this approach has been successfully used is in the study the energetics of cation ordering (Palin et al., 2001; Bosenick et al., 2001). Such studies involve calculating the relaxed energies of many configurations with different cation distributions, typically resulting in the tuning of a model Hamiltonian for use in Monte Carlo models (as described below).

#### 4.2. Long-range electrostatic interactions

The short-range contributions to the lattice energy are easy to evaluate. Starting from each atom in the unit cell, the interaction between that atom and all other atoms lying within a sphere of a pre-defined radius are calculated. With the speed of modern computers, this radius can be selected such that all short-range interactions have effectively fallen to zero at this value. Here we note that one of the technical issues is the fact that the sum of the electrostatic interactions, and also dispersive interactions, cannot be accurately summed this way. Instead it is necessary to sum to infinite atomic separation, which is clearly impossible to achieve. The common approach is to use a summation technique developed by Ewald, in which the summation is separated into two parts, one in real space and one in reciprocal space, which are both converge reasonably quickly. The Ewald sum was originally developed for the Coulomb interaction, but it can be generalised to include any function that varies as  $r^{-n}$ , with  $n = 6$  as the most common example. Although most simulation scientists know this, it is one of the major technical issues that the novice needs to quickly appreciate. The Ewald sum was originally developed for three-dimensional crystals, but it has been generalised for one (Bródka, 2002) and two (Harris, 1998) dimensions (e.g. for studying line defects or surfaces).

#### 4.3. Computation of properties

One common application of lattice energy modelling is to calculate physical properties. In simulations using empirical models, the energy expression can be extended to include the contribution due to external elastic stresses and electric fields, taking account of the coupling to relaxation of internal structural coordinates and strains.

From this expression it is possible to derive equations for dielectric constants (both static and high frequency, with the latter yielding the refractive index tensor), piezoelectric tensors, and elastic constant tensor.

Typically it is not common to achieve the same level of accuracy as is attainable for structures. Errors of order 10–20% are not uncommon for the best models without expending a lot of effort, and not surprisingly accuracy can be worse for small off-diagonal tensor components.

It is anticipated that quantum mechanical methods might give better results for physical properties. However, this is offset by the fact that the level of theory to obtain physical properties from quantum mechanical calculations is much harder. Properties such as elastic constants (Baroni et al., 1987a) and dielectric constants (Baroni et al., 1987b) can be computed using a perturbation theory approach (“linear response” method), but actually elastic constants are more commonly computed by applying a set of strains and fitting the elastic constant values to the computed energies.

#### *4.4. Calculations of defect energies*

Using empirical models it is possible to calculate the energies of defects using the Mott-Littleton method (Catlow, 2000). In this method, the structure contained within a sphere surrounding the defect is relaxed, and the structure on the surface of the sphere is connected with the infinite crystal through the use of an interface region. There have been many studies of different types of defects in minerals and ceramics using this approach (e.g. Wright et al., 1994; Walker et al., 2003).

#### *4.5. Studies of surfaces*

Lattice energy models have found widespread use for simulations of surfaces, particularly for the study of adsorption of ions and molecules onto surfaces. Some codes – e.g. GULP (Gale and Rohl, 2003) – use a two-dimensional version of the Mott-Littleton algorithm, but with quantum mechanical approaches it is more common to use two-dimensional slabs of the structure separated by a vacuum which allows the calculation to be performed using a standard three-dimensional modelling code. These calculations are invariably considerably more demanding than standard crystalline calculations, and carry the risk that any minimum energy state is not the global minimum.

## 5. Lattice dynamics calculations

### 5.1. Introduction to the theory

The theory of lattice dynamics (Dove, 1993) relies on the ability to expand the crystal energy in terms of small atomic displacements to quadratic order. Restricting ourselves here to using one-dimensional atomic displacements in order to make nomenclature more transparent, the energy for a crystal is written as:

$$E = \frac{1}{2} \sum_{ij,\ell\ell'} u_i(\ell) \frac{\partial^2 E}{\partial u_i(\ell) \partial u_j(\ell')} u_j(\ell') = \frac{1}{2} \sum_{ij,\ell\ell'} u_i(\ell) \Phi_{ij}(\ell\ell') u_j(\ell') \quad (3)$$

where  $i,j$  label the atoms within the unit cell, and  $\ell,\ell'$  label different unit cells. This equation is easily generalised to three-dimensions, albeit at the cost of additional book-keeping, to three dimensions.

The vibrations themselves are described in terms of orthogonal linear combinations of atomic motions, namely the normal modes. When the above equation is recast in terms of the normal modes, it is found that the equations yield a matrix – the “dynamical matrix” – that can be constructed in terms of interatomic force constants and phase factors and which can be diagonalised to give the squares of the angular frequencies of all vibrations for any given wave vector.

We write the atomic displacement in terms of normal mode coordinates  $Q_k$  and mode eigenvectors  $\mathbf{e}$  written as vector allowing for three-dimensional atomic displacements within this one-dimensional system:

$$\mathbf{u}(j\ell,t) = \frac{1}{\sqrt{Nm_j}} \sum_k \mathbf{e}(j,k) \exp(i\mathbf{k} \cdot \mathbf{r}_{j\ell}) Q_k \quad (4)$$

where we use the shorthand  $k$  to denote both the wave vector  $\mathbf{k}$  and the phonon branch number (and  $-k$  denotes the complex conjugate, namely the same mode at wave vector  $-\mathbf{k}$ ). Substitution yields the following equation for the energy:

$$E = \frac{1}{2} \sum_k \mathbf{e}^T(k) \cdot \mathbf{D}(k) \cdot \mathbf{e}(k) Q_k Q_{-k} \quad (5)$$

where  $\mathbf{e}(k)$  is now a column matrix that incorporates the motions of all atoms in the unit cell the dynamical matrix, and  $\mathbf{D}(\mathbf{k})$ , is given as:

$$D_{ij}(\mathbf{k}) = \frac{1}{\sqrt{m_i m_j}} \sum_{\ell} \Phi_{ij}(0\ell) \exp(i\mathbf{k} \cdot [\mathbf{r}_i(0) - \mathbf{r}_j(\ell)]) \quad (6)$$

This is an eigenvalue/eigenvector problem, in which the dynamical matrix can be diagonalised to yield:

$$E = \frac{1}{2} \sum_k \omega_k^2 Q_k Q_{-k} \quad (7)$$

where  $\omega_k$  is the angular frequency of the normal mode. The task is to compute the elements of the dynamical matrix from the calculation of the dynamical matrix. Expansion to three dimensions is straightforward, but with messy notation.

## 5.2. Practical calculations

The dynamical matrix lends itself naturally to simulations using empirical calculations, where the second derivatives are easily computed. Practically it is not costly to perform a set of calculations for a range of wave vectors. In the early days of lattice dynamics modelling – in the 1960's–80's, following the development of inelastic neutron scattering methods for the measurement of phonon dispersion curves – the dynamical matrix was set up using individual and independent force constants. Nowadays it is more typical to calculate the values of the force constants from the empirical energy functions. In this approach, it is essential that the calculation is performed on the structure after it has been relaxed (see previous section), and as a result codes such as GULP (Gale, 1997; Gale and Rohl, 2003) integrate lattice energy and lattice dynamics calculations. We remark that for models that are not specifically tuned against vibrational data it is possible to achieve accuracy of around 20% or better on calculated frequencies.

The situation is not so straightforward for quantum mechanical calculations, because the derivatives within the dynamical matrix are not easy to form. These calculations do not contain the equivalent of an interaction between two atoms that makes setting up the dynamical matrix so easy for empirical models. There have been a number of attempts to compute the derivatives numerically by using large samples and computing the energy for finite displacements of pairs of atoms (Ackland et al., 1997). However, the emerging technique is to use linear response methods (Baroni et al., 1987b), recognising that these are both theoretically complex and computationally demanding. On the other hand, this approach is capable of high accuracy for a wide range of systems, certainly

better than possible for most empirical models. It is possible to achieve accuracy on computed frequency of a few percent.

### 5.3. Lattice dynamics and thermodynamic functions

One important application for lattice dynamics models is the computation of thermodynamic properties. For example, the free energy is given as:

$$\begin{aligned}
 F &= \frac{1}{2} \sum_k \hbar \omega_k + k_B T \sum_j \ln[1 - \exp(-\hbar \omega_k / k_B T)] \\
 &= k_B T \sum_k \ln[2 \sinh(\hbar \omega_k / 2k_B T)]
 \end{aligned}
 \tag{8}$$

where in practice the sum is over a grid of wave vectors.

Lattice dynamics modelling may be used to compute thermodynamic properties, such as the specific heat, for a range of temperatures and pressures. In fact some of these functions are actually rather insensitive to the accuracy of the calculated frequencies, and it is not hard to obtain agreement with experiment to within 2%. A further application is to use calculations of free energies to compute phase stability and phase diagrams. The pioneering study in this regard was the computation of the  $\text{Mg}_2\text{SiO}_4$  phase diagram using empirical models (Price et al., 1987), with the discrepancies with experiment being due to small errors in the relative computed enthalpies of different phases (see comments about the performance of empirical models under pressure above). On the other hand, the slopes of phase boundaries were computed reasonably accurately, leading us to expect that when using this approach with quantum mechanical methods we might expect to be able to make useful predictions.

### 5.4. The quasiharmonic approximation and thermal expansion

The theory discussed so far uses the harmonic approximation, one limitation of which is that it does not allow for thermal expansion. To overcome this, it is common to use what is known as the “quasiharmonic approximation”, in which the harmonic approximation is retained but with force constants that will vary with volume. If we start from the equilibrium structure at zero pressure, and increase the volume, the energy will increase but the majority of vibrational frequencies will decrease. If we then compute the free energy, we will find that the increase in energy is offset by an increase in the entropy with a subsequent contribution of  $-TS$  to the free energy. Thus there will be a

temperature at which this new volume corresponds to the minimum of the overall free energy. It is straightforward to see how this approach can be used to simulate thermal expansion (Allan et al., 2001).

## 6. Molecular dynamics methods

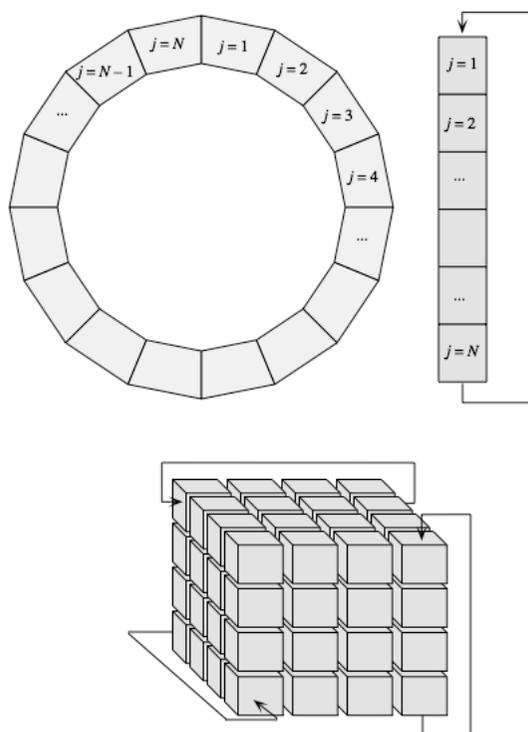
### 6.1. Essential details of the method

The previous discussion has pointed out the need to incorporate temperature and anharmonic terms into the simulation of materials, and the best approach that takes this further is the molecular dynamics (MD) method (Allen and Tildesley, 1988; Rapaport, 2004). The MD method is simple in concept: the idea is to exploit the fact that Newton's equation of motion links forces to acceleration. We start with a large configuration of atoms, and use either empirical or quantum mechanical methods to compute the force on each atom. This force is converted to acceleration, and using a numerical time step algorithm the acceleration can be combined with information on the current and previous atomic positions, velocities and/or accelerations to predict the position of each atom a small time interval (time step) later. There are actually several algorithms in common use, giving the level of stability and accuracy required by the specific application. The result of a simulation is the evolution of the positions and velocities of an ensemble of atoms through time, and using the approach discussed here this ensemble will have a constant energy and volume (i.e. the microcanonical ensemble). It is possible to extend the set of equations to include additional dynamical variables that enable the simulations to correspond to constant temperature (variable energy) and constant pressure/stress (variable sample volume and shape).

The idea outlined above is very simple, but is remarkably powerful. Typically time steps correspond to around 1/20th of the smallest natural period of vibration (whether of a high-frequency bond-stretching vibration or the oscillation of the shell model), and it is common to run simulations for time of 5 ps (in the case of large quantum mechanical simulations) up to several 100 ps (with empirical models). In fact there are good reasons why to simulate much longer leads to diminishing returns (see below). With quantum mechanical simulations sample sizes of around a few 100 atoms are as much as is feasible for most researchers; the limiting factor is that computational requirements often scale as between the second and third power of the number of atoms. Using modern algorithms that scale linearly with the number of atoms, and by using large multi-processor computers with low-latency inter-processor communications, it is possible to run simulations containing a few million atoms with empirical models.

## 6.2. Some technical details

The key technical point that now needs to be discussed concerns how one handles the boundaries of the sample. In lattice energy and lattice dynamics models, we embed the unit cell within an infinite crystal of identical unit cells without giving it much thought. When we extend the method to simulate larger-scale processes, such as the adsorption of molecules on mineral surfaces, we simply define a large unit cell. In the MD method we commonly use the same approach, which is now said to be using “periodic boundary conditions”. The periodic boundary construction is illustrated in Figure 2. It means that the simulation does not have any free surfaces and thus more accurately resembles a true bulk material.



**Figure 2.** Representation of periodic boundary conditions, showing both one and three-dimensional representations. The building blocks are typically unit cells.

Although the use of periodic boundary conditions solves the problems that would be created if our models had free surfaces, they do introduce some specific artefacts. One is that a travelling disturbance will propagate through one surface and reappear through the opposite surface. After a time related to the ratio of the sample length and the velocity of

sound the disturbance will reappear at the same point, giving the possibility of a feedback mechanism. As a result, any phenomena that lasts for time longer than this “recurrence time” will be hard to quantify in a way that is free from the effects of this artefact. Another consequence of using finite-sized samples is that the number of wave vectors contributing to the overall behaviour is limited. This course-grained sampling of wave vector space can cause some of the amplitudes of fluctuations to be dependent on sample size, typically giving lower values for smaller samples.

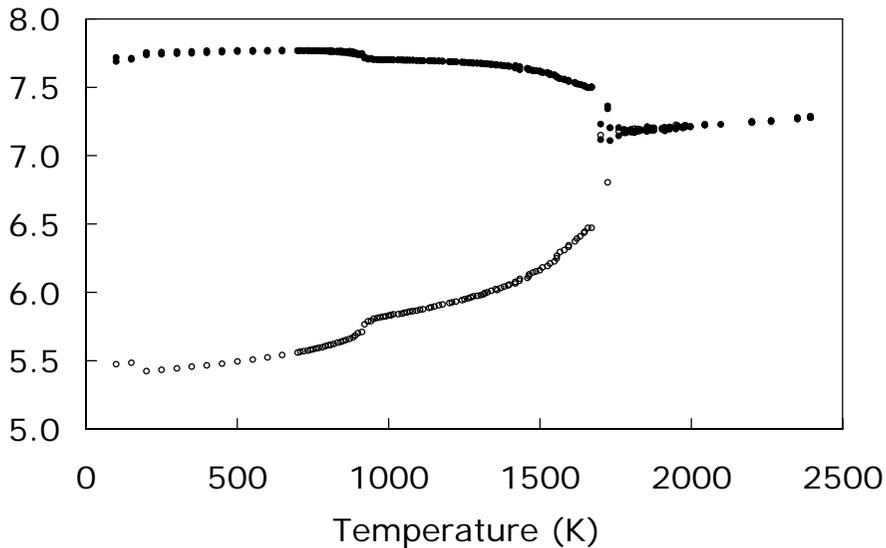
### 6.3. What can you calculate with the MD method?

In short, barring constraints of statistical accuracy and sensitivity, and of finite sample size and running time, one can compute many things that can be measured experimentally and more beside. To discuss this we consider a number of different types of standard analysis.

*First* we have the evolution of the system with temperature and pressure in the case of equilibrium phenomena, and the evolution of the system with time in the case where the sample is reacting to a specific disturbance. For example, it is straightforward to compute the equation of state for a range of temperatures and pressure; an example is our recent work on understanding the compressibility of amorphous silica (section 8.3). Similarly it is possible to study phase transitions using this approach; Figure 3 gives an example, showing the lattice parameters of rhombohedral  $\text{BaCO}_3$  on heating through one disordering phase transition and then through a phase transition into a cubic disordered rocksalt structure. In the case of the time evolution of a sample, we cite work on studying the response of minerals and ceramics to the high-energy impact of an atom that has undergone radioactive decay (see section 8.2).

*Second* it is possible to compute distribution functions, such as the distribution of orientations of molecules or of the positions of atoms. A common example is the pair distribution function, which is particularly useful to understand atomic coordination in fluids.

*Third* is the analysis of dynamics, frequently through the use of time-correlation functions. In simulations of fluids or highly-dynamic systems, it is possible to calculate diffusion constants from analysis of single atom motions. In simulations of solids, studies of fluctuations through time-correlation functions will give detailed information on the time scales for dynamic processes. In this regard we note that the Fourier transform of the single-particle velocity autocorrelation function for a crystalline material yields the phonon density of states. Recent developments have led to methods to compute properties such as the thermal conductivity (Archer et al., 2003).



**Figure 3.** Temperature dependence of the lattice parameters of  $\text{BaCO}_3$  obtained by molecular dynamics simulations using a rigid-ion empirical model with the interatomic potentials of reference 41. The starting structure at low-temperature is the cubic  $R\bar{3}c$  structure of calcite, and the final structure at high-temperature is the cubic rocksalt structure. The lower set of data (open circles) correspond to the  $c$  lattice parameter in the trigonal setting scaled to the value in the cubic phase, and the upper set of data (closed circles) correspond to the  $a$  and  $b$  lattice parameters in the trigonal setting similarly scaled to their values in the cubic (data are unpublished at the time of writing).

#### 6.4. Thermodynamic integration

There has been a lot of recent interest in using molecular dynamics methods to provide quantitative information on phase stability under extreme conditions, one key example being recent work on the solid–fluid phase diagram of iron (Guillan, 2006). The task revolves around the need to be able to compute the free energy at any temperature or pressure. One important approach is “thermodynamic integration”. In this method, a direct comparison is made with a similar model for which the free energy is known exactly. We define the Hamiltonian of the system we are interested in as  $H$ , and the Hamiltonian of the reference state as  $H_0$ , and we define a mixed Hamiltonian as  $H_\lambda = \lambda H - (1 - \lambda)H_0$ . The procedure is to run a series of simulations for a range of Hamiltonians  $H_\lambda$ . The free energy of the system of interest is then obtained as:

$$F = F_0 + \int_0^1 \frac{\partial F_\lambda}{\partial \lambda} d\lambda \quad (9)$$

where  $F_0$  is the known free energy of the reference state defined by  $H_0$ , and  $F_\lambda$  is the free energy associated with the Hamiltonian  $H_\lambda$ . It can be shown that the differential can be obtained from the relation:

$$\frac{\partial F_\lambda}{\partial \lambda} = \langle \Delta H \rangle_\lambda \quad (10)$$

where the right hand side equals the expectation value of  $H - H_0$  evaluated for configurations generated by the Hamiltonian  $H_\lambda$ .

Clearly the route to obtaining the free energy at a single state point involves a large number of separate simulations, but computing resources to match this task are now available.

### 6.5. Methods for non-equilibrium and rare events

We conclude this survey of the MD method by noting that there is growing interest in using variants of the MD method for simulations of events such as transition states that involve thermally-activated crossing of high potential energy barriers. In order to study such events, it is essential to vastly accelerate their rate. A number of methods exist for this, including “temperature accelerated dynamics” (Voter et al., 2002) and “metadynamics” (Laio and Parrinello, 2002) approaches.

## 7. Metropolis Monte Carlo methods

### 7.1. Basics of the Monte Carlo method

The Monte Carlo (MC) method is a technique to sample the infinite number of available configurations of a material (Binder and Heermann, 2002). In principle the MD method is one route to obtaining such a sample, but there are cases where the MD method is not sufficiently efficient, particularly when sampling configurations that can only evolve dynamically at an extremely slow rate. An example is the study of cation site ordering.

The Metropolis implementation of the MC method, the one that is common use in computational science, requires an algorithm for generating a new configuration by changing a previous configuration. Any change in the configuration or atomic ensemble will give a change in energy. If the energy change  $\Delta E$  is negative, leading to a lowering of the energy, the change is automatically accepted. On the other hand, if the energy

change is positive, the configuration is only accepted with probability  $\exp(-\Delta E/k_B T)$ . This procedure is repeated for a large number of steps, leading to an evolution of the ensemble through the multi-dimensional phase space. This approach ensures that the sampling procedure is consistent with thermodynamics; for example, states with any energy  $E$  occur with the relative probability  $\exp(E/k_B T)$ . The complete set of configurations, including duplicated configurations, can be analysed to give averaged quantities that have the correct thermodynamic weighting automatically ensured. For example, it is possible to calculate the average energy as a function of temperature, and the variance of the energy will yield the heat capacity using the standard formulation based on fluctuations.

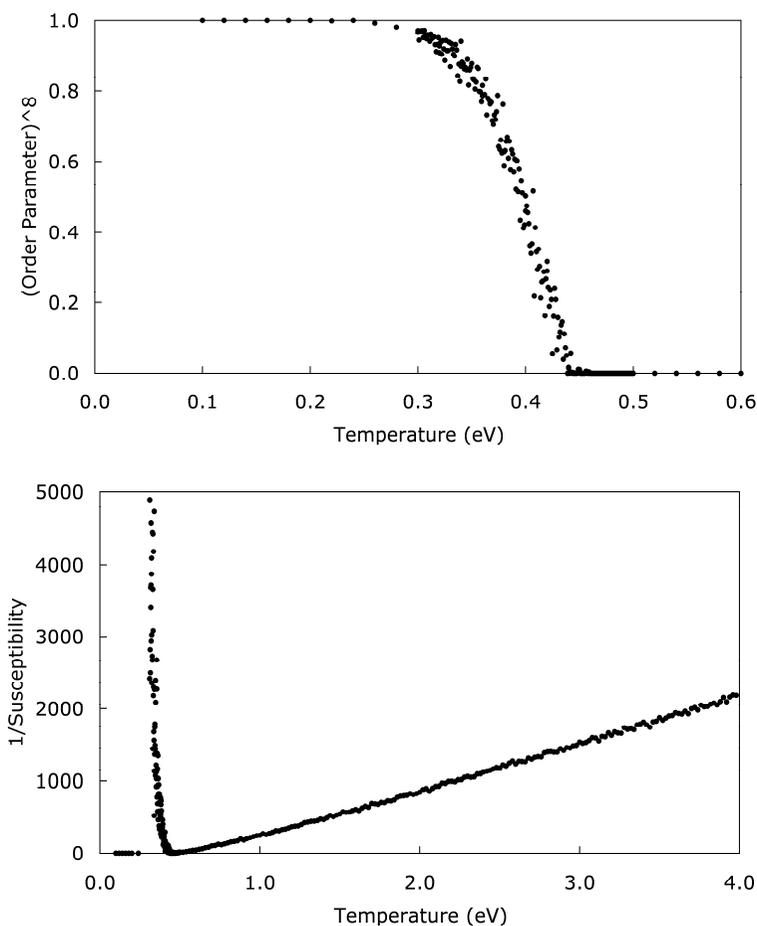
Briefly we mention two types of study that might be conducted using MC methods. The first is where each MC step consists of a small displacement of one of the atoms within the simulation, and the energy change is computed using either empirical or quantum mechanical methods. This is similar to the sort of approach used in the MD method, but is providing an alternative way to sample the phase space defined by atomic positions, but will not provide information on dynamics. It can be more efficient than MD in obtaining a good sample. The second type of study is when using different types of variable and a model Hamiltonian. The common example is a spin model, where each site has a variable that can be mapped onto a vector with the interactions between neighbours being described by a model Hamiltonian. The problem of atomic site ordering, where each MC step consists of swapping the sites occupied by two atoms chosen at random, is an example of a system that can be mapped onto a spin model.

## 7.2. Example of cation ordering

The basic ideas are most easily amplified with reference to an example, and we consider here the case study of the ordering of two types of cation in a clay silicate layer (Warren et al., 2001; Sainz-Díaz et al., 2003). We define a model Hamiltonian in which pairs of sites, labelled by  $n$ , have energy  $J_n$  if they are occupied by the same cation and  $-J_n$  if they are occupied by different cations. Figure 4 shows the temperature-dependence of various quantities that have been calculated using the MC model; the two points to make here are that these are representative data of the sort of task for which the MC method is good at obtaining, and that the MC method lends itself to comprehensive sweeps through parameter space (in this case temperature) in order to give a detailed picture of the behaviour being studied.

### 7.3. Some technical points for discussion

One of the points highlighted by Figure 4 is that the results from an MC run are subject to some degree of statistical error. Indeed, the MC method is a statistical sampling through a multi-dimensional phase space, and thus it is essential that MC simulations run for enough steps to minimise the effects of statistical errors on the detailed analysis. A second point highlighted by Figure 4 is the rounding of the results around the transition temperature; this is due to the finite size of the sample, and with the type of detail that MC methods can give it is important to minimise these effects using samples that are as large as possible.



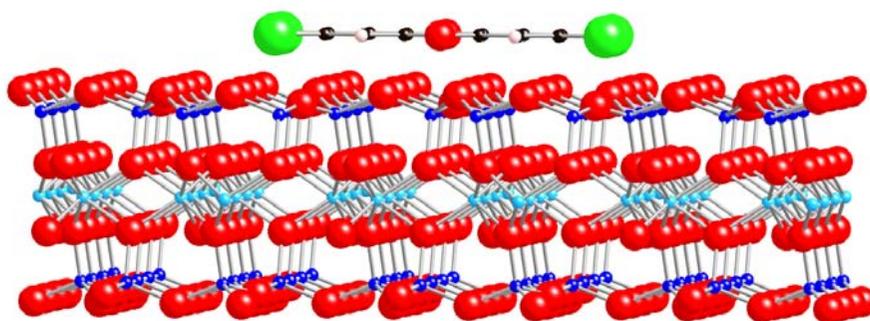
**Figure 4.** Temperature dependence of the order parameter (top) and inverse susceptibility (bottom) computed for a set of Mg and Al cations within an octahedral layer in a model mica/clay structure (Sainz-Díaz et al., 2003). The simulation method was the Monte Carlo method, using a model spin Hamiltonian with interactions taken to fifth neighbour.

## 8. Case studies

Here I cite a few illustrative examples from my own work to highlight a number of issues; representative examples from other studies have been cited above, and are given in the other contributions to this volume.

### 8.1. Study of pollutant molecules on mineral surfaces

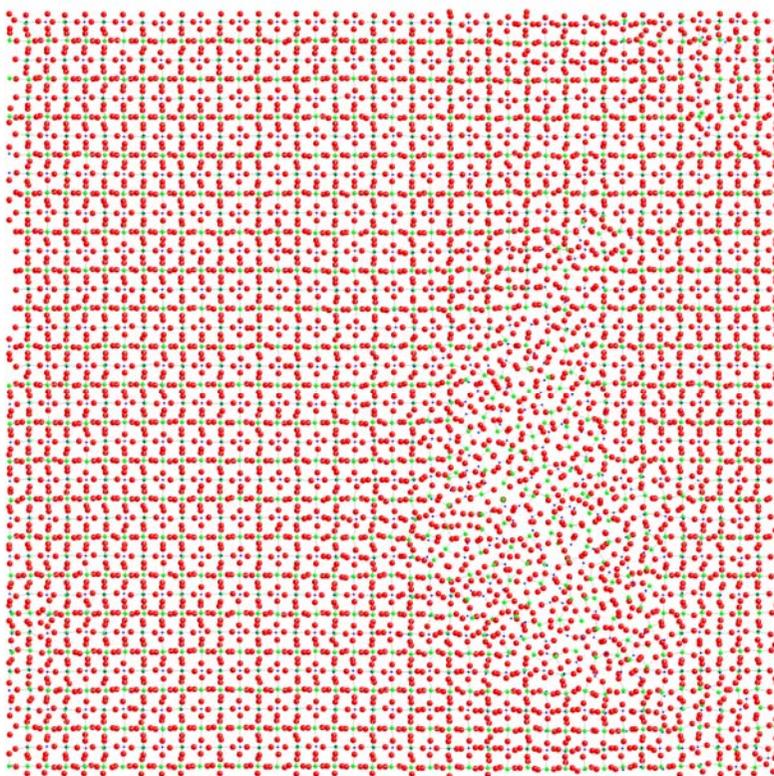
We are engaged in a detailed study of the energetics of adsorption of pollutant molecules on mineral surfaces, an example being dioxin ( $C_{10}O_2Cl_xH_{8-x}$ ) molecules on clay surfaces (Austen et al., 2008), as illustrated in Figure 5. Our tools for this work include a combination of DFT and empirical models, using both lattice energy and MD methods. The scientific challenge is to understand the factors associated with adsorption, and one key finding from a sweep across all molecules (congeners) within this family of molecules is that the most important factor is the number of Cl atoms. Less important are the factors of where they are in the molecule and the specific location or orientation of the molecule. Adsorption energies of less than 1 eV per molecule have been identified, and the similar results have been found for adsorption on silica and carbonate minerals. The simulation challenge for the empirical modelling is that it is not easy to parameterise all the interactions in this study, and for the DFT work is the fact that the DFT approach does not explicitly take account of dispersive interactions (which may be important for physisorption processes) and the simulations require a large number of atoms. Linear-scaling DFT codes such as SIESTA are required for such large systems.



**Figure 5.** Dioxin molecule on a clay surface, an example of the type of environmental simulation now amenable to molecular simulation using both empirical and *ab initio* models (Austen et al., 2008).

## 8.2. Radiation damage in ceramics

Long-term storage of high-level nuclear waste products, such as plutonium, may require encapsulation within ceramic matrices. We have been running large simulations of high-energy (some tens of keV) recoils of heavy atoms with crystalline ceramics such as zircon (Geisler et al., 2003) and perovskite (Trachenko et al., 2005). For this work we require MD studies using millions of atoms. The MD code DL\_POLY\_3 (Todorov and Sminth, 2004) has been specifically optimised for this work. Figure 6 shows an example of the damage caused by a single recoil event in a simulation. It is clear that large samples are required to contain the extent of the damage, and even larger samples for the study of the aggregation of recoil damage due to several events. Animations show the time evolution of the damage process. These simulations require several hours running over 512 or 1024 processors on a national supercomputing facility.



**Figure 6.** Part of an atomic configuration of zircon,  $ZrSiO_4$ , that has been damaged following a high-energy recoil of a heavy atom (Geisler et al., 2003). This structure is now frozen into the structure following a quasi-melting of the structure around the trajectory of the recoil atom.

### 8.3. Compressibility of amorphous silica

Amorphous silica has the strange property that its compressibility has a maximum at a pressure of 2 GPa. The headline restatement of this fact is that it initially becomes softer under increasing pressure before beginning to harden again. We have run a set of MD studies using empirical models parameterised by quantum mechanical calculations, running these as a set of parameter sweeps through pressure for different configurations and two potential models (Walker et al., 2007). Results for the pressure-dependence of volume are shown in Figure 7, which show a consistent maximum in the gradient –  $dV/dP$  at a pressure above zero. To understand the behaviour, the results were analysed in terms of the inherent flexibility of the glass network, showing a strong correlation between compressibility and flexibility.

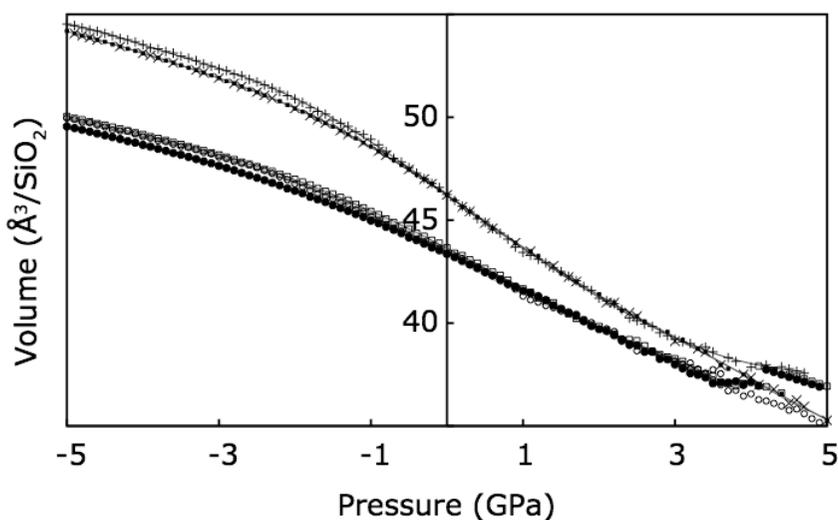


Figure 7. Simulated volume of amorphous silica,  $\text{SiO}_2$ , obtained using the molecular dynamics simulation method using two interatomic potential models and 3 different atomic configurations (Walker et al., 2007). The key point to note is that there is a maximum in the slopes of the curves at around 2 GPa.

## 9. Computing and data issues

### 9.1. New computing paradigms

The discussion presented above has focussed on the algorithms and scientific motivation. In this section I address wider practical computing issues. As noted in the introduction, the past 1–2 decades have seen enormous changes in the capabilities of

computers. Even the simplest of calculations would once require some waiting time on a standard workstation, and MD and MC methods required high-performance computing. Both raw computing power and processor memory would not be available on the desktop. *Ab initio* methods would necessarily be the preserve of specialised computing resources. This situation has changed dramatically. For example, the capability of my laptop computer far exceeds what once was used by the US Department of Commerce<sup>2</sup> as the definition of a supercomputer, so defined to control exports to unfriendly nations. It is now quite possible to run large-scale computations on low-cost computing resources.

Over the past few years we have seen the emergence of two complementary type of computing resource that can give large amounts of computing power. One is the traditional “high-performance” resource, which typically consists of thousands of processors with high-speed and low-latency communication links between individual processors. Programs that exploit this type of resource need to be able to farm out parts of the overall task to separate processors to be carried out in parallel with other parts, and procedures are required to facilitate data transfer and orchestration of the various tasks. Some types of simulation, particularly large-scale MC methods, work particularly well with this type of architecture, facilitating the sort of large scale simulation illustrated by the work on radiation damage. Unfortunately not all methods are amenable to algorithms that scale well across hundreds or thousands of processors, and many *ab initio* codes fall into this class. Such cases tend to work well up to a few tens of processors, exploiting both the increased number of processors and available memory, but beyond a certain limit adding more processors does not yield a significant reduction in the time required to complete a simulation. The other type of computing resource is what is known as “high-throughput” computing. In this case large numbers of processors are used, with each processor being assigned an independent task. There is no need for fast communications between processors because individual processors do not share data. Whereas in high-performance applications one task is shared over many processors – this is called parallel processing – in high-throughput applications it is the actual study that is parallelised, not the individual tasks. Thus in the MC example shown in Figure 4, each processor in the high-throughput resource ran one or more separate simulations at given temperatures. There are two types of high-throughput resources in common use. One is a purpose-built cluster of processors – typically a rack of processors – and the other is a grid of desktop computers that is managed using middleware tools such as Condor (Thain et al., 2003). It should be noted that the

---

<sup>2</sup> US Department Of Commerce, Bureau of Export Administration. US federal document 94-4156. <http://www.fas.org/spp/starwars/offdocs/940224.htm>, 1994 (URL active November 2007)

boundaries between the two types of resource can be blurred. The emerging computing paradigm is for multi-core processors, so that we will soon all have parallel computers on our desks as the norm, and people will run clusters or grids of small multi-processors machines. Whilst these can be run as separate processors, we will also run small parallel computations in high-throughput mode. This will be a particular boon to many *ab initio* studies.

## 9.2. Data and information management

With the rapid expansion of computing power comes a new problem, namely how to manage the deluge of data that will follow. Scientists will not merely do bigger and better calculations: they will also be empowered to run very detailed multi-dimensional parameter sweeps involving many thousands (or more) individual simulations as part of a single study. Thus we have to face the challenge of enabling simulation scientists to manage both their data and the information content of their data with accuracy and reliability. Up to now most scientists who are beginning to exploit the power of computing to run large-scale parameter sweeps are using bespoke data management methods, but these run the risk of being inaccurate (how do you know that any extracted number corresponds to the run you think it does?) and make it hard for anyone else to find their way through the data.

We cite here work we have been carrying out on data management as part of our work on grid computing for simulation scientists. Basically there are three parts to this. First is the need to be able to manage data within a write-infrequently read-often (WIFRO) system with an organisation of the data being controlled by the job submission tools (Dove et al., 2007). The traditional practice of scientists incorporating their data within their normal file system is not going to cope with the amounts of data we can now generate, both because the data structures are not properly defined, and because the data are easily corrupted. For example, many of us too easily switch between test and production mode, without the file system properly reflecting this switch. In short, it is too easy to not organise data management with long-term data curation in mind, and a scientist's normal desktop file system is not the ideal location for data curation. There are now data curation and data grid solutions appearing in routine use. For example, the Storage Resource Broker (Rajasekar et al., 2003) is an early example of such a system. The important challenge is to make data curation an integral part of the job submission and data analysis process, and not a task that is carried out post-analysis. Tools that handle this process include RMCS (Dove et al., 2007) and the Application Hosting Environment (Coveney et al., 2006).

The second piece of work on data management reflects the need to have machine-readable representation of the data. We have found that XML is particularly useful for this, and the Chemical Markup Language (CML) is one XML language that has been designed with the needs of our community in mind (White et al., 2006a). Briefly, XML files have each piece of data marked with tags to make the meaning of the data clear defined. For example, the energy computed by a simulation can be marked up as

```
<property dictRef="ossia:Energy" title="Energy">  
  <scalar dataType="fpx:real" units="cmlUnits:eV">7.22e0</scalar>  
</property>
```

In this example, energy is classified as a property. It is given a title, and associated with an entry in an XML dictionary specific for the simulation code (OSSIA) that defines the meaning of the variable (“energy” is a particular example of a term that can take on many shades of meaning). The actual number is specified as being of scalar data type real floating point number, and the units are specified in a separate XML units dictionary.

XML/CML writing is now possible from simulation codes, and examples such as GULP (Gale and Rohl, 2003), DL\_POLY\_3 (Todorov and Smith, 2004), CASTEP (Segal et al., 2002), SIESTA (Soler et al., 2002) and OSSIA (Warren et al., 2001) now write CML output files. One advantage of using XML is that it is easy to transform XML files to formats that highlight the information content of data are which are easy for non-experts to read (White et al., 2006b). XHTML is one such useful format, with tables of data being transformed to graphs on-the-fly (e.g. using the Scalar Vector Graphics, SVG<sup>3</sup>, format, another XML language tailored for graphics) or atomic configurations rendered within the browser for easy viewing (eg using the Java Jmol program<sup>4</sup>). Another advantage of XML is that tools exist to automatically extract key information from files in useful formats. We note that a number of mainstream simulation and codes have very poor textual data representation, requiring the user to know in detail what the output means, often by reference back to separate input files.

The third piece of work on data management concerns the use of metadata, which is basically defined as “data about data”. The former ideal was that all data were documented in a laboratory notebook, but this is highly impractical with the huge quantities of data that can be generated. Each data file (or set of files) must now be accompanied with a set of metadata that includes information on the details of the job (e.g. when and where run, and by whom), of the code, of the input parameters, of the

---

<sup>3</sup> Scalable Vector Graphics (SVG): XML Graphics for the Web. <http://www.w3.org/Graphics/SVG/> (URL active November 2007)

<sup>4</sup> Jmol: an open-source Java viewer for chemical structures in 3D. <http://www.jmol.org/> (URL active November 2007)

core output values, and of the context of the job. A good set of metadata allows data files to be unambiguously identified, and is essential for any search tool. Moreover, by extracting a rich set of metadata that includes both input parameters and core output data values such as final or average values, metadata can provide an easy interface to data (Tyer et al., 2007).

### 9.3. Integration of job and data management

I remarked above that data management needs to be integrated into the job submission process. Briefly I note here that our RMCS tool for job submission includes the creation of a data archive, the upload of output files to the data archive, and the creation and storage of a rich set of metadata associated with each job (Dove et al, 2007). Additional tools can make the creation and submission of hundreds of RMCS jobs as part of a parameter sweep easy, and metadata tools can make the collation of tables of output data simple. For example, the data shown in Figures 2 and 3 were obtained using two simple shell commands to set up and submit all the jobs represented in the data (one point per job), and a third command was used to collate the data for plotting. The automation of the job submission and data management tasks now releases the time of the scientist for doing science.

### Acknowledgements

I am pleased to acknowledge the help of many collaborators in aspects of this work, including Emilio Artacho, Kat Austen, Richard Bruin, Mark Calleja, Andrew Goodwin, Peter Murray-Rust, Erika Palin, Kostya Trachenko, Andrew Walker, Andrew Walkingshaw, Toby White (Cambridge), Kerstin Kleese van Dam, Ilian Todorov and Rik Tyer (Daresbury Laboratory), Alfonso Hernandez-Laguna and Ignacio Sainz-Diaz (Granada), Dan Wilson and Björn Winkler (Frankfurt) and Ian Swainson (Chalk River, Ontario). Financial support from NERC, Royal Society, CMI and EPSRC is gratefully acknowledged.

### References

- Ackland, G.J., Warren, M.C., and Clark, S.J. (1997) Practical methods in ab initio lattice dynamics. . *Journal of Physics: Condensed Matter* 9, 7861–7872.
- Allan, N.L., Barrera, G.D., Barron, T.H.K., and Taylor, M.B. (2001) Evaluation of thermodynamic properties of solids by quasiharmonic lattice dynamics. *International Journal of Thermophysics* 22, 535–546.

- Allen, M.P. and Tildesley, D.J. (1988) *Computer Simulation of Liquids*. Oxford University Press.
- Archer, T.D., Birse, S.E.A., Dove, M.T., Redfern, S.A.T., Gale, J.D., and Cygan, R.T. (2003) An interatomic potential model for carbonates allowing for polarization effects. *Physics and Chemistry of Minerals* 30, 416–424.
- Austen, K.F., White, T.O.H., Marmier, A., Parker, S.C., Artacho, E., and Dove, M.T. (2008) Electrostatic versus polarisation effects in the adsorption of aromatic molecules of varied polarity on an insulating hydrophobic surface. *Journal of Physics: Condensed Matter* 20, Art n° 03521.
- Baroni, S., Giannozzi, P., and Testa, A. (1987a) Elastic-constants of crystals from linear-response theory. *Physical Review Letters* 59, 2662–2665.
- Baroni, S., Giannozzi, P., and Testa, A. (1987b) Green-function approach to linear response in solids. *Physical Review Letters* 58, 1861–1864.
- Binder, K. and Heermann, D.W. (2002) *Monte Carlo Simulation in Statistical Physics, 4th edition*. Springer.
- Born, M., and Mayer, J.E. (1932) Zur Gittertheorie der Ionenkristalle. *Zeitschrift für Physik* 75, 1–18.
- Bosenick, A., Dove, M.T., Myers, E.R., Palin, E.J., Sainz-Diaz, C.I., Guiton, B., Warren, M.C., Craig, M.S., and Redfern, S.A.T. (2001) Computational methods for the study of energies of cation distributions: applications to cation-ordering phase transitions and solid solutions. *Mineralogical Magazine* 65, 193–219.
- Bródka, A. (2002) Ewald type summation method for electrostatic interactions in computer simulations of a three-dimensional system periodic in one direction. *Chemical Physics Letters* 363, 604–609.
- Burnham, C.W. (1990) The ionic model: Perceptions and realities in mineralogy. *American Mineralogist* 75, 443–463.
- Catlow, C.R.A. (2000) Perspective on "Conduction in polar crystals. I. Electrolyte conduction in solid salts" - Mott NF, Littleton MJ (1938) *Trans Faraday Soc* 34: 485. *Theoretical Chemistry Accounts* 103, 205–208.
- Coveney, P.V., Sadiq, S.K., Saksena, R., Thyveetil, M., Zasada, S.J., McKeown, M., and Pickles, S. (2006) A lightweight application hosting environment for grid computing. *Proceedings of the UK e-Science All Hands Meeting 2006*, 217–224.
- Dove, M.T. (1989) On the computer modelling of diopside: towards a transferable potential for silicate minerals. *American Mineralogist* 74, 774–779.
- Dove, M.T. (1993) *Introduction to lattice dynamics*. Cambridge University Press.
- Dove, M.T. (2003) *Structure and Dynamics: an atomic view of materials*. Oxford University Press.
- Dove, M.T., Cool, T., Palmer, D.C., Putnis, A., Salje, E., and Winkler, B. (1993) On the role of Al/Si ordering in the cubic–tetragonal phase transition in leucite. *American Mineralogist* 78, 486–492.
- Dove, M.T., and Redfern, S.A.T. (1997) Lattice simulation studies of the ferroelastic phase transitions in (Na,K)AlSi<sub>3</sub>O<sub>8</sub> and (Sr,Ca)Al<sub>2</sub>Si<sub>2</sub>O<sub>8</sub> feldspar solid solutions. *American Mineralogist* 82, 8–15.
- Dove, M.T., Walker, A.M., White, T.O.H., Bruin, R.P., Austen, K.F., Frame, I., Chiang, G-T., Murray-Rust, P., Tyler, R.P., Couch, P.A., Kleese van Dam, K., Parker,

- S.C., Marmier, A., and Arrouvel, C. (2007) Usable grid infrastructures: practical experiences from the e-Minerals project. *Proceedings of the UK e-Science All Hands Meeting 2007*, 48–55.
- Evans, D.J. (1986) Thermal conductivity of the Lennard-Jones fluid. *Physical Review A*, 34, 1449–1453.
- Gale, J.D. (1996) Empirical potential derivation for ionic materials. *Philosophical Magazine B* 73, 3–19.
- Gale, J.D. (1997) GULP – a computer program for the symmetry adapted simulation of solids. *Journal of the Chemical Society: Faraday Transactions* 93, 629–637.
- Gale, J.D. and Rohl, A.L. (2003) The General Utility Lattice Program, *Molecular Simulation* 29, 291–341.
- Geisler, T., Trachenko, K., Rios, S., Dove, M.T., and Salje, E.K.H. (2003) Impact of self-irradiation damage on the aqueous durability of zircon ( $\text{ZrSiO}_4$ ): implications for its suitability as a nuclear waste form. *Journal of Physics: Condensed Matter*, 15, L597–L605.
- Gillan, M.J., Alfè, D., Brodholt, J., Vocadlo, L., and Price, G.D. (2006) First-principles modelling of Earth and planetary materials at high pressures and temperatures. *Reports on Progress in Physics* 69, 2365–2441.
- Gonze, X., Beuken, J.M., Caracas, R., Detraux, F., Fuchs, M., Rignanese, G.M., Sindic, L., Verstraete, M., Zerah, G., Jollet, F., Torrent, M., Roy, A., Mikami, M., Ghosez, P., Raty, J.Y., and Allan, D.C. (2002) First-principles computation of material properties: the ABINIT software project. *Computational Materials Science* 25, 478–492.
- Harris, F.E. (1998) Ewald summations in systems with two-dimensional periodicity. *International Journal of Quantum Chemistry* 68, 385–404.
- Laio, A. and Parrinello, M. (2002) Escaping free-energy minima. *Proceedings of the National Academy of Sciences of the United States of America* 99, 12562–12566.
- Ohsumi, T., Slater, B., Gao, F.F., Yu, J.H., Sakamoto, Y., Zhu, G.S., Terasaki, O., Vaughan, D.E.W., Qiu, S.L., and Catlow, C.R.A. (2004) Fine structures of zeolite-linde-L (LTL): Surface structures, growth unit and defects. *Chemistry – A European Journal* 10, 5031–5040.
- Palin, E.J., Dove, M.T., Redfern, S.A.T., Bosenick, A., Sainz-Diaz, C.I., and Warren, M.C. (2001) Computational study of tetrahedral Al-Si ordering in muscovite. *Physics and Chemistry of Minerals* 28, 534–544.
- Patel, A., GD Price, G.D., and MJ Mendelssohn, M.J. (1991) A computer-simulation approach to modeling the structure, thermodynamics and oxygen isotope equilibria of silicates. *Physics and Chemistry of Minerals* 17, 690–699.
- Payne, M.C., Teter, M.P., Allan, D.C., Arias, T.A., and Joannopoulos, J.D. (1992) Iterative minimization techniques for abinitio total-energy calculations – molecular dynamics and conjugate gradients. *Reviews of Modern Physics* 64, 1045–1097.
- Post, J.E. and Burnham, C.W. (1986) Ionic modeling of mineral structures and energies in the electron gas approximation:  $\text{TiO}_2$  polymorphs quartz, forsterite, diopside. *American Mineralogist* 71, 142–150.

- Price, G.D., Parker, S.D., and Leslie, M. (1987) The lattice dynamics and thermodynamics of the  $\text{Mg}_2\text{SiO}_4$  polymorphs. *Physics and Chemistry of Minerals* 15, 181–190.
- Rajasekar, A., Wan, M., Moore, R., Schroeder, W., Kremenek, G., Jagatheesan, A., Cowart, C., Zhu, B., Chen, S-Y., and Olschanowsky, R. (2003) Storage Resource Broker - Managing distributed data in a grid. *Computer Society of India Journal*, 33, 42-54.
- Rapaport, D.C. (2004) *The Art of Molecular Dynamics Simulation*, 2nd edition. Cambridge University Press.
- Sainz-Díaz, C.I., Hernandez-Laguna, A., and Dove, M.T. (2001) Modelling of dioctahedral 2:1 phyllosilicates by means of transferable empirical potential. *Physics and Chemistry of Minerals* 28, 130–141.
- Sainz-Díaz, C.I., Palin, E.J., Hernández-Laguna, A., Dove, M.T. (2003) Octahedral cation ordering of illite and smectite. Theoretical exchange potential determination and Monte Carlo simulations. *Physics and Chemistry of Minerals*, 30, 382–392.
- Sanders, M.J., Leslie, M. and Catlow, C.R.A. (1984) Interatomic potentials for  $\text{SiO}_2$ . *Journal of the Chemical Society Chemical Communications*, 1271–1273.
- Schwarz, K. and Blaha, P. (2003) Solid state calculations using WIEN2k. *Computational Materials Science* 28, 259–273.
- Segall, M.D., Lindan, P.J.D., Probert, M.J., Pickard, C.J., Hasnip, P.J., Clark, S.J., and Payne, M.C. (2002) First-principles simulation: ideas, illustrations and the CASTEP code. *Journal of Physics: Condensed Matter* 14, 2717–2744.
- Smith, W (2006) A short description of DL\_POLY. *Molecular Simulation* 32, 935–943.
- Soler, J.M., Artacho, E., Gale, J.D., Garcia, A., Junquera, J., Ordejon, P., and Sanchez-Portal, D. (2002) The SIESTA method for ab initio order-N materials simulation. *Journal of Physics: Condensed Matter* 14, 2745–2779.
- Thain, D., Tannenbaum, T., and Livny, M. (2003) Condor and the Grid. In: *Grid computing: making the global infrastructure a reality*, (F. Berman, A.J.G. Hey, and G. Fox, Eds.), Wiley, Chapter 11.
- Todorov, I.T. and Smith, W. (2004) DL\_POLY\_3: the CCP5 national UK code for molecular-dynamics simulations. *Royal Society Philosophical Transactions* 362, 1835–1852.
- Trachenko, K., Pruneda, M., Artacho, E., and Dove, M.T. (2004) Radiation damage effects in the perovskite  $\text{CaTiO}_3$  and resistance of materials to amorphization. *Physical Review B* 70, Art n° 134112.
- Tsuenyuki, S., Tsukada, M., Aoki, H., and Matsui, Y. (1988) 1st principles interatomic potential of silica applied to molecular dynamics. *Physical Review Letters* 61, 869–872.
- Tyer, R.P., Couch, P.A., Mortimer-Jones, T.V., van Dam, K., Todorov, I.T., Bruin, R.P., White, T.O.H., Walker, A.M., Austen, K.F., and Dove, M.T. (2007) Metadata management and grid computing within the e-Minerals project. *Proceedings of the UK e-Science All Hands Meeting 2007*, 415–422.

- Voter, A.F., Montalenti, F., and Germann, T.C. (2002) Extending the time scale in atomistic simulation of materials. *Annual Review of Materials Research* 32, 321–346.
- Walker, A.M., Sullivan, L.A., Trachenko, K., Bruin, R.P., White, T.O.H., Dove, M.T., Tyer, R.P., Todorov, I.T., and Wells, S.A. (2007) The origin of the compressibility anomaly in amorphous silica: A molecular dynamics study. *Journal of Physics: Condensed Matter* 19, Art n° 275210.
- Walker, A.M., Wright, K., and Slater, B. (2003) A computational study of oxygen diffusion in olivine. *Physics and Chemistry of Minerals* 30, 536–545.
- Warren, M.C., Dove, M.T., Myers, E.R., Bosenick, A., Palin, E.J., Sainz-Diaz, C.I., Guiton, B., and Redfern, S.A.T. (2001) Monte Carlo methods for the study of cation ordering in minerals. *Mineralogical Magazine* 65, 221–248.
- White, T.O.H., Murray-Rust, P., Couch, P.A., Tyer, R.P., Bruin, R.P., Todorov, I.T., Wilson, D.J., Dove, M.T., Austen, K.F., and Parker, S.C. (2006A) Application and uses of CML within the e-Minerals project. *Proceedings of the UK e-Science All Hands Meeting 2006*, 606–613.
- White, T.O.H., Dove, M.T., Bruin, R.P., Austen, K.F., Walker, A.M., Artacho, E., Murray-Rust, P., Walkingshaw, A.D., Marmier, A., Parker, S.C., Couch, P.A., Tyer, R.P., Todorov, I.T., and Wilson, D.J. (2006B) Information delivery in computational mineral science: the e-Minerals data handling system. *Proceedings of Proceedings of the Second IEEE International Conference on e-Science and Grid Computing 2006*, 59.
- Winkler, B., Dove, M.T., and Leslie, M. (1991) Static lattice energy minimization and lattice dynamics calculations on minerals using three-body potentials. *American Mineralogist* 76, 313–331.
- Wright, K., Freer, R., and Catlow, C.R.A. (1994) The energetics and structure of the hydrogarnet defect in grossular – a computer-simulation study. *Physics and Chemistry of Minerals* 20, 500–503.