2.3

INTERATOMIC POTENTIAL MODELS FOR IONIC MATERIALS

Julian D. Gale
Nanochemistry Research Institute, Department of Applied Chemistry,
Curtin University of Technology, Perth, 6845, Western Australia

Ionic materials are present in many key technological applications of the modern era, from solid state batteries and fuel cells, nuclear waste immobiliza tion, through to industrial heterogeneous catalysis, such as that found in automotive exhaust systems. With the boundless possibilities for their utilization, it is natural that there has been a long history of computer simulation of their structure and properties in order to understand the materials science of these systems at the atomic level.

The classification of materials into different types is, of course, an arbitrary and subjective decision. However, when a binary compound is composed of two elements with very different electronegativities, as is the case for oxides and halides in particular, then it is convenient to regard it as being an ionic solid. The implication is that, as a result of charge transfer from one element to the other, the dominant binding force between particles is the Coulombic attraction between opposite charges. Such materials tend to be characterized by close-packed, dense structures that show no strong directionality in the bonding.

Typically, most ionic materials possess a large band gap and are therefore insulating. As a consequence, the notion that the solid is composed of spherical ions whose interactions can be represented by simple distance-dependent functional forms is quite a reasonable one, since overtly quantum mechanical effects are lesser than in materials where covalent bonding occurs. Thus it is possible to develop force fields that are specific for ionic materials, and this approach can be surprisingly successful considering the simplicity of the interatomic potential model.

When considering how to construct a force field for ionic materials, the starting point, as is the case for all types of system, is to assume that the total

S. Yip (ed.),
Handbook of Materials Modeling, 479–497.

energy, U_{tot}, can be decomposed into interactions between different numbers of atoms:

$$U_{\text{tot}} = \frac{1}{2!}\sum_i \sum_j U_{ij} + \frac{1}{3!}\sum_i \sum_j \sum_k U_{ijk} + \frac{1}{4!}\sum_i \sum_j \sum_k \sum_l U_{ijkl} + \cdots$$

Here, U_{ij} is the energy of interaction between a pair of atoms, i and j, or so-called two-body interaction energy; U_{ijk} is the extra interaction that arises (beyond the sum of the three two-body energy components for the pairs $i - j$, $j - k$, and $i - k$) when a triad of atoms is considered, and so forth for higher order terms. Note that the inverse factorial prefactor is required to avoid double counting of interactions between particles.

In principle, the above decomposition is exact if carried out to terms of high enough order. However, in practice it is necessary to truncate the expansion at some point. For many ionic materials it is often sufficient to only include the two-body term, though the extensions beyond this will be discussed later.

Imagining an ionic solid as being composed of cations and anions whose electron densities are frozen, which represents the simplest possible case, the physical interactions present can be intuitively understood. There will obviously be a Coulombic attraction between ions of opposite charge, with a corresponding repulsive force between those of like nature. Because ions are arranged such that the closest neighbours are of opposite sign, this gives rise to a strong net attractive energy that will tend to contract the solid in order to lower the energy. In order that an equilibrium structure is obtained there must be a counterbalancing repulsive force. This arises from the overlap of the electron densities of two ions, regardless of the sign of their charge, and has its origin in the Pauli repulsion between electrons. Hence, we can write the breakdown of the two-body energy in general terms as:

$$U_{ij} = U_{ij}^{\text{Coulomb}} + U_{ij}^{\text{repulsive}}$$

While real spherical ions will have a radial electron density distribution, it is convenient to treat the ions as point charges – i.e., as though all the electron density is situated at the nucleus. Within this approximation, the electrostatic interaction of two charged particles is just given by Coulomb's law;

$$U_{ij}^{\text{Coulomb}} = \frac{q_i q_j}{4\pi\epsilon_0 r_{ij}}$$

or, if written in atomic units, as will subsequently be done, we can drop the constant factor of $4\pi\epsilon_0$:

$$U_{ij}^{\text{Coulomb}} = \frac{q_i q_j}{r_{ij}}$$

The error in the electrostatic energy arising from the point charge approximation is usually subsumed into the repulsive energy contribution, since this latter term is usually derived by a fitting procedure, rather than from direct theoretical considerations.

1. Calculating the Electrostatic Energy

Not only is the electrostatic energy the dominant contribution to the total value, but it turns out that it is actually the most difficult to evaluate. While it is easy to write down that the electrostatic energy is the sum over all pairwise interactions, including all periodic images of the unit cell, the complication arises because the sum must be truncated for actual computation. Unfortunately, the summation is an example of a conditionally convergent series, i.e., the value of the sum depends on how the truncation is made. The reason for this can be understood by considering the interactions of a single ion with all other ions within a given radius, r. The convergence of the energy of interaction, U^r_{tot}, is given by the number of interactions, N_r, multiplied by the magnitude of the interaction, U^r:

$$U^r_{\text{tot}} = \sum_r N_r U^r$$

As r increases, the number of interactions rises in proportion to the surface area of the cut-off sphere:

$$N_r \propto 4\pi r^2.$$

However, the interaction itself only decreases as the inverse power of r, as has been shown previously. Consequently, the magnitude of interaction potentially increases as the cut-off radius is extended. The fact that the magnitude converges in practice relies on the fact that there is cancelation between interactions with cations and anions.

It turns out that the electrostatic energy of a system actually depends on the macroscopic state of a crystal due to the long-ranged effect of Coulomb fields. In other words, it is not purely a property of the bulk crystal, but also depends, in general, on the nature of the surfaces and of the crystal morphology [3]. To make it feasible to define an electrostatic energy that is useful for the simulation of ionic materials, it is conventional to impose two conditions on the Coulomb summation:

1. The sum of the charges within the system must be equal to zero:

$$\sum_i q_i = 0$$

2. The total dipole moment of the system in all directions must also be equal to zero:

$$\mu_x = \mu_y = \mu_z = 0$$

If these conditions are satisfied, the electrostatic energy will always converge to the same value as the cut-off radius is incremented.

It is also possible to define the electrostatic energy when the dipole moments along the three Cartesian axes differ from zero. This Coulomb energy is related to the value obtained when the dipole moment is zero, U^0, according to the following expression;

$$U = U^0 + \frac{2\pi}{3V}\left(\mu_x^2 + \mu_y^2 + \mu_z^2\right)$$

where V is the volume of the unit cell. Considering the expression for the dipole moment in a given direction, α;

$$\mu_\alpha = \sum_i q_i r_{i\alpha}$$

where $r_{i\alpha}$ is the position of the ith ion projected on to this axis, then there is a complication. Because there are multiple images of the same ion, due to the presence of periodic boundary conditions, the dipole contribution of any given ion is an ambiguous quantity. The only way to determine the true dipole moment is to perform the sum over all ions within the entire crystal, which includes those ions at the surface. This is the origin of the electrostatic energy being a macroscopic property of the system.

While it has been stated that the electrostatic energy is convergent if the above conditions are obeyed, it is not obvious how to achieve this in practice for a general crystal structure. Various methods have been proposed, the most reknown of which is that of Evjen who constructed charge neutral shells of ions about each interacting particle. However, this is more difficult to automate for a computational implementation and is best for high symmetry structures.

Apart from the need to converge to a defined electrostatic energy, there is also the issue of how rapidly the sum converges, since it is required that the calculation be fast for numerical evaluation. By far the dominant approach to evaluating the electrostatic energy is through the use of the summation method due to Ewald which aims to accelerate the convergence by partially transforming the expression into reciprocal space. While the details of the derivation are beyond the scope of this text, and can be found elsewhere [2, 9], the concepts behind the approach and the final result will be given below.

In Ewald's approach, a Gaussian charge distribution of equal magnitude, but opposite sign, is placed at the position of every ion in the crystal. Because the charges cancel, all but for the contribution from the differing shape of the distribution, the resulting electrostatic interaction between ions is now rapidly convergent when summed out in real space and converges to the energy U^{real}. In order to recover the original electrostatic energy it is then necessary to compute two further terms. Firstly, the interaction of the Gaussian charge distributions with each other must be subtracted. Because of the smooth nature of the electrostatic potential arising from such a distribution, it is possible to efficiently evaluate this term, U^{recip}, by expanding the charge density in planewaves with the periodicity of the reciprocal lattice. Again, the energy contribution is rapidly convergent with respect to the cut-off radius within reciprocal space. Finally, there is the self-energy, U^{self}, that arises from the interaction of the Gaussian with itself.

Mathematically, the Ewald sum is derived by a Laplace transform of the Coulomb energy and the final expressions are given below;

$$\begin{aligned}
U^{\mathrm{Coulomb}} &= U^{\mathrm{real}} + U^{\mathrm{recip}} + U^{\mathrm{self}} \\
U^{\mathrm{real}} &= \frac{1}{2}\sum_{R}\sum_{i}\sum_{j}\frac{q_i q_j}{r_{ij}} erfc\left(\eta^{\frac{1}{2}} r_{ij}\right) \\
U^{\mathrm{recip}} &= \frac{1}{2}\sum_{G'}\sum_{i}\sum_{j}\frac{4\pi}{V^2} q_i q_j \ \exp\left(iG.r\right)\frac{\exp -(G^2/4\eta)}{G^2} \\
U^{\mathrm{self}} &= -\sum_{i} q_i^2 \left(\frac{\eta}{\pi}\right)^{\frac{1}{2}}
\end{aligned}$$

where R denotes a real space lattice vector, G represents a reciprocal lattice vector and η is a parameter that determines the width of the Gaussian charge distribution. Note that the summation over reciprocal lattice vectors excludes the case when $G = 0$.

The key to rapid convergence of the Ewald sum is to choose the optimal value of η. If the value is small, then the Gaussians are narrow and so the real space expression converges quickly, while the reciprocal space sum requires a more extensive summation due to the higher degree of curvature of the charge density. Choosing a large value of η obviously leads to the inverse situation. One approach to choosing the convergence parameter is to derive an expression for the total number of terms to be evaluated in real and reciprocal space for a given accuracy and then to find the stationary point where this quantity is at a minimum. The choice of η_{opt} is then given by;

$$\eta_{\mathrm{opt}} = \left(\frac{N\pi^3}{V}\right)^{\frac{1}{3}}$$

where N is the number of particles within the unit cell. If the target accuracy, A, is represented by the given fractional degree of convergence (e.g.,

$A = 0.001$ would imply that the energy is converged to within 0.1%), then the cut-off radii in real and reciprocal space are given as follows:

$$r_{\text{opt}}^{\max} = \left(\frac{-\ln A}{\eta}\right)^{\frac{1}{2}}$$

$$G_{\text{opt}}^{\max} = 2(-\eta \ln A)^{\frac{1}{2}}$$

Before leaving the evaluation of the electrostatic energy, it is important to comment on other dimensionalities than three-dimensional (3-D) periodic boundary conditions. There is also an analogous approach involving a partial reciprocal space transformation in two dimensions, due to Parry, which can be employed for slab or surface calculations [6]. For the 1-D case of a polymer, the Coulomb sum is now absolutely convergent for a charge neutral system. However, it is still beneficial to use methods that accelerate the convergence, though there is less concensus as to the most efficient technique.

2. Non-electrostatic Contributions to the Energy

While the electrostatic energy often accounts for the majority of the binding, the non-Coulombic contributions are equally critical since they determine the position and shape of the energy minimum. As previously mentioned, there must always be a short-ranged repulsive force between ions to counter the Coulomb attraction and therefore prevent the collapse of the solid. Most work has followed the pioneering work in the field, as embodied in the Born–Meyer and Born–Lande equations for the lattice energy, by utilizing either an exponential or inverse power-law repulsive term. This gives rise to two widely employed functional forms, namely the Buckingham potential;

$$U_{ij}^{\text{short-ranged}} = A_{ij} \exp\left(-\frac{r_{ij}}{\rho_{ij}}\right) - \frac{C_{ij}}{r_{ij}^{6}}$$

and that due to Lennard–Jones:

$$U_{ij}^{\text{short-ranged}} = \frac{B_{ij}}{r_{ij}^{m}} - \frac{C_{ij}}{r_{ij}^{n}}$$

For the Lennard–Jones potential, the exponents m and n are typically 9–12 and 6, respectively. This latter potential can also be recast in many different forms by rewriting in terms of the well-depth, ε, and either the distance at which the potential intercepts the $U_{ij}^{\text{repulsive}} = 0$ axis, r_0, or the position of the minimum, r_{eq}. Both the Buckingham and Lennard–Jones potentials have the same common features – a short-ranged repulsive term and a slightly longer-ranged attractive term. The latter contribution, often referred to as the $C6$ term, arises as the leading term in the expansion of the dispersion energy between two non-overlapping charge densities.

When choosing between the use of Buckingham and Lennard–Jones potentials, there are arguments for and against both. Physically, the exponential form of the Buckingham potential should be more realistic because electron densities of ions decay with this shape and so it would seem natural that the repulsion follows the magnitude of the interacting ion densities, at least for weak overlap. However, in the limit of $r_{ij} \rightarrow 0$ the repulsive Buckingham potential tends to A_{ij}, i.e., a constant value that is unphysically low for nuclear fusion! Worse still, if the coefficient C_{ij} is non-zero, then the potential, while initially repulsive, goes through a maximum and then tends to $-\infty$ – a result that is physically absurd. In contrast, the Lennard-Jones potential behaves sensibly and tends to $+\infty$ as long as $m > n$. While the false minimum of the Buckingham potential is not usually a problem for energy minimization studies, it can be an issue in molecular dynamics where there is a finite probability of the system gaining sufficient kinetic energy to overcome the repulsive barrier.

There is a further solution to the problems with the Buckingham potential at small distances. The problems arise due to the simple power-law expression for the dispersion energy. However, this is also incorrect at short-range since the electron densities begin to overlap leading to a reduction of the dispersion contribution. This can be accounted for by explicitly damping the $C6$ term as the distance tends to zero, and the most widely used approach for doing this is to adopt the form proposed by Tang and Toennies:

$$U_{ij}^{C6} = -\left[1 - \left\{\sum_{k=0}^{6} \frac{\left(br_{ij}\right)^k}{k!}\right\} \exp\left(-br_{ij}\right)\right] \frac{C_{ij}}{r_{ij}^6}$$

Occasionally other short-ranged, two-body potentials are choosen, such as the Morse or a harmonic potential. However, these are normally selected when acting between two atoms that are bonded. In this situation, the potential is usually Coulomb-subtracted too, in order that the parameters can be directly equated with the bond length and curvature.

All the above short-ranged potentials are pairwise in form. However, there are instances where it is useful to include higher order contributions. For example, in the case of semi-ionic materials, such as silicates, where there is a need to reproduce a tetrahedral local coordination geometry, it is common to include three-body terms that act as a constraint on an angle:

$$U_{ijk} = \frac{1}{2} k_3 \left(\theta_{ijk} - \theta_{ijk}^0\right)^2$$

There are also many variants on this, such as including higher powers of the deviation of the angle from the equilibrium value and the addition of an exponential dependence on the bond lengths so that the potential becomes smooth and continuous with respect to coordination number changes. For systems containing particularly polarizable ions, there is also the possibility of including the three-body contribution to the dispersion energy, as embodied in the Axilrod–Teller potential.

As with all materials, it is necessary to select the most approriate force field functional form based on the physical interactions that are likely to dominate in an ionic material. While this will often consist of just the electrostatic term and a two-body short-ranged contribution for dense close-packed materials, it may be necessary to contemplate adding further terms as the degree of covalency and structural complexity increases.

3. Ion Polarization

Up to this point we have considered ions to have a frozen spherical electron density that may be represented by a point charge. While this is a reasonable representation of many cations, it is not as accurate a description for anions which tend to be much more polarizable. This can be readily appreciated for the oxide ion, O^{2-} in particular. In this case, the first electron affinity of oxygen is favourable, while the second electron affinity is endothermic due to the Coulomb repulsion between electrons. Consequently, the second electron is only bound by the electrostatic potential due to the surrounding cations, and therefore the distribution of this electron will be strongly perturbed by the local environment. It is therefore natural to include the polarizability of anions, and even some larger cations, in ionic potential models when reliable results are required.

While polarization may occur to arbitrary order, here the focus will be on the dipole polarizability, α, which is typically the dominant contribution. In the presence of an electric field, E, the dipole moment, μ, generated is given by;

$$\mu = \alpha E$$

and the polarization energy, U^{dipolar}, that results is:

$$U^{\text{dipolar}} = -\tfrac{1}{2}\alpha E^2$$

The electric field at an ion is given by the first derivative of the electrostatic potential with respect to the three Cartesian directions, and therefore can be calculated from the Ewald summation for a bulk material. In principle, it is then straightforward to apply the above point ion polarizability correction to the total energy of a simulation. However, it introduces extra complexity since the induced dipole moments will also generate an electric field at all other ions in the system. Hence, it is necessary to consider the charge–dipole and dipole–dipole interactions as well. The whole procedure involves iteratively solving for the dipole moments on the ions until self-consistency is achieved in a manner analogous to the self-consistent field procedure that occurs in quantum mechanical methods.

There is one disadvantage to the use of point ion polarizabilities, as described above, which is that the value of α is a constant. Physically, the more polarized an ion becomes, the harder it should be to polarize it further, and so the induced dipole is prevented from reaching extreme values. If the polarizablity is a constant, a so-called polarization catastrophe can occur in which the total electrostatic energy becomes exothermic faster than the repulsive energy increases leading to the collapse of two ions onto each other. This is particularly problematic with the Buckingham potential since the energy at zero distance tends to $-\infty$.

An alternative description of dipolar ion polarization that addresses the above problem is the shell model introduced by Dick and Overhauser [4]. Their approach is to create a simple mechanical model for polarization by dividing each ion into two particles, known as the core and the shell. Here the core can be conceptually thought of as representing the nucleus and core electrons, while the shell represents the more polarizable valence electrons. Thus the core is often positively charged, while the shell is negatively charged, though when utilizing a shell model for a cation it is not uncommon for both core and shell to share the positive charge. Both particles are Coulombically screened from each other and only interact via a harmonic restoring force:

$$U^{\text{core}-\text{shell}} = \tfrac{1}{2}k_{\text{cs}}r_{\text{cs}}^2$$

where r_{cs} is the distance between the core and shell.

There are two important consequences of the shell model approach. Firstly, because the shell enters the simulation as a point particle, the achievement of electrostatic self-consistency is transformed into a minimization of the shell coordinates. Consequently, this is achieved concurrently with the optimization of the real atomic positions (namely the core positions), though at the cost of doubling the number of variables. While this significantly increases the time required to invert the Hessian matrix, assuming Newton–Raphson optimization is being employed, the convergence rate is also enhanced through all the information on the coupling of coordinates with the polarization being utilized. Secondly, it is the usual convention for the short-ranged potentials to act on the shell of a particle, rather than on the core, which leads to the polarizability becoming environment dependent. If the force constant (second derivative) of the short-range potential acting on the shell is k_{SR} and the shell charge is

q_{shell}, the polarizability of the ion is equal to:

$$\alpha = \frac{q_{\text{shell}}^2}{k_{\text{cs}} + k_{\text{SR}}}$$

Special handling of the shell model is required in some simulations. In particular, for molecular dynamics the presence of a particle with no mass potentially complicates the solution of Newton's equations of motion. However, there are two solutions to this that parallel the techniques found in electronic structure methods. One approach is to divide the atomic mass so that a small fraction is attributed to the shell instead of the core. If chosen to be small enough, the frequency spectra for the shells is higher than any mode of the real material, such that the shells are largely decoupled from the nuclear motions. The disadvantage of this is that a smaller timestep is required in order to achieve an accurate integration. Alternatively, the shells can be minimized at every timestep in order to follow the adiabatic surface. Although the same timestep can now be used as per core-only dynamics, the cost per move is greatly increased. Similarly in lattice dynamics, it is also necessary to consider the contribution from relaxation of the shell positions to the dynamical matrix, which will act to soften the energy surface.

Both point ion polarizabilities and the shell model have benefits for interatomic potential simulations of ionic materials. Firstly, they act to stabilize lower symmetry structures and hence it would not be possible to reproduce the structural distortion of various materials without their inclusion. Secondly, they make it possible to determine many materials properties that intrinsically have a strong electronic component. For instance, both the low and high frequency dielectric constant tensors may be calculated, where the former is determined by both the electronic and nuclear contributions, while the latter is purely dependent on the contribution from the polarization model.

4. Derivation of Ionic Potentials

So far, the typical functional form of the interaction energy in ionic materials has been described, without discussing how the parameter values are arrived at within the model. Many aspects are similar to general forcefield derivation as practiced for organic and inorganic systems, be they ionic or not. However, there are a few differences also that will be highlighted below.

Given the dominance of the electrostatic contribution for ionic materials, the starting point for any force field is to determine the nature of the point charges to be employed. There are two broad approaches – either to employ the formal valence charge or to chose smaller partial charges. The main advantages of formal charges are that they remove a degree of freedom from the fitting process and also ensure wide compatability of force fields, in that parameters from binary compounds can be combined to model ternary or more complex phases where the cations do not have the same formal valence charge. Furthermore, when studying defects in materials the vacancy, interstitial or impurity will be guaranteed to carry the correct total charge. On the other hand, for materials with a formal valence of greater than $+2$ it is argued that formal charges are unrealistic and so partial charges must be used. Indeed, Mulliken charges from *ab initio* calculations do suggest that such materials are not fully ionic. However, the Mulliken charge is only one of several charge partitioning schemes. Arguably more pertinent measures of ionicity are the Born effective charges that describe the response of the charge density to an electric field. For a solid, where it is not possible to determine the charges that best reproduce the external electrostatic potential, as would be the case for molecules, considering the dipolar response is the next best thing. It is often the case that formal charges, in combination with a shell model for polarization, yield very similar Born effective charges to periodic density functional calculations [6]. Consequently, for low symmetry structures at least, both formal and partial charges can be equally valid in a well derived model.

Having determined the charge states of the ions, it is then necessary to derive the short-range and other parameters for the force field by fitting. Parameter derivation falls into one of two classes, either being based on the use of theoretical or experimental data. While truly *ab initio* parameter derivation is desirable, most theoretical procedures are subject to systematic errors and so empirical fitting to experimental information has tended to be prevalent.

Fitting consists of specifying a training set of observable quantities, that may be derived theoretically or experimentally, and then varying the parameters in a least squares procedure in order to minimize the discrepancy between the calculated and observed values [5]. Typically, the training set would consist of one or more structures that represent local energy minima (i.e., stable states with zero force) and data that provide information as to the curvature of the energy surface about these minima, such as bulk moduli, elastic constants, dielectric constants, phonon frequencies, etc. Ideally, multiple structures and as much data as possible should be included in the procedure in order to maximize transferability and to constrain the parameters to physically sensible values. Because it is possible to weight the observables according to their reliability or importance there can never be a single unambiguous fit.

In the above brief statement of what fitting is, it is given that the structural data is to be used as an observable. However, there are several distinct ways in which this can be done. If the force field is a perfect fit then the forces calculated at the observed experimental, or theoretically optimized, structure should be zero. Hence it is common to use the forces determined at this point as the observable for fitting, rather than the structure per se, since they are straight forward to calculate. In practice, the quality of the fit is usually imperfect and so there will be residual forces. Lowering the forces does not guarantee that the discrepancy in the optimized structural parameters will be minimized though, since this also depends on the curvature. Assuming that the system is within the harmonic region, the errors in the structure, Δx, will be related to residual force vector, f_{resid}, according to

$$\Delta x = H^{-1} f_{\mathrm{resid}}$$

where H is the Hessian matrix containing the second derivatives. Thus one approach to directly fitting the structure is to use the above expression for the errors in the structure. Alternatively, the structure can be fully optimized for each evaluation of the fit quality, which is considerably more expensive, but guaranteed to be reliable regardless of whether the energy surface is quadratic or not. This latter method, referred to as relaxed fitting, also possesses the advantage that any curvature related properties can be evaluated for the structure of zero force, such that the harmonic expressions employed are truly valid.

The case of a shell model fit deserves special mention here, since the issues do not usually arise during fits to other types of model. Because of the mapping of dipoles to a coordinate space representation there is the question of how to handle the shell positions during a fit. Given that the cores are equated with the nuclear position, and that it is difficult to ascribe atom-centered dipoles in a crystal, there is rarely any information on where the shells should be sited. In a relaxed fit the issue disappears, since the shells just optimize to the position of minimum force. For a conventional force-based fit then the shells must either still be relaxed explicitly at each evaluation of the sum of squares, or their coordinates can be included as variable parameters such that the relaxation occurs concurrently with the fitting process.

Theoretical derivation of parameters can either closely resemble empirical fitting, by inputing calculated observables, or alternatively an energy hypersurface can be utilized. In this latter case many different structures, usually sampled from around the energy minima, are specified along with their corresponding energies. As a result, the curvature of the energy surface is fitted directly rather than by assuming harmonic behavior about the minimum. Again the issue of weighting is particularly important since it tends to be more crucial to ensure a good quality of fit close to the minimum at the expense of points that are further away. To date it has been more common to utilize quantum mechanical data for finite clusters in potential derivation, rather than directly fitting solid state *ab initio* information. However, this introduces uncertainties, since it is not clear how transferable the gas phase cluster data will be to bulk materials since they are dominated by surface effects.

There are two further theoretical methods for parameter derivation that deserve a mention, namely electron gas methods and rule-based methods. The first is particularly significant since it was a popular approach in the early days of the computer simulation of ionic materials at the atomistic level. In the electron gas method, the energy of overlapping frozen ion electron densities is calculated according to density functional theory as a function of distance. These energies can then be used directly via splines or fitted to a functional form. Given that not all ions, such as O^{2-}, are stable *in vacu*, the ion densities were usually determined in an appropriate potential well to mimic the lattice environment. The results obtained directly from this procedure where not always accurate, given the limitations of density functional theory, so often the distance dependence was shifted to improve the position of the minimum. The second alternative theoretical approach is to use rules that encapsulate how to determine interactions from atomic properties, such as the polarizability and atomic radius, in order to generate force fields of universal applicability. Of course, this compromises the accuracy of the results for any given system, but can be useful for systems were there is little known data to fit to.

5. Applications of Ionic Potentials

Having defined the appropriate force field for a material, it is then possible to calculate many different properties in a very straight forward fashion. Simulations can be broadly divided into two categories – static and dynamic. In a static calculation, the structure of a material is optimized to the nearest local minimum, which may represent one desired polymorph of a system, as opposed to the global minimum, and then the properties are derived by consideration of the curvature about that position. For example, many of the mechanical, vibrational and electrical response properties are all functions of the second derivatives of the energy with respect to atomic coordinates and lattice strains. For pair potentials, the determination of these properties is not dramatically more expensive than the evaluation of the forces, with the exception of matrix inversions that may be required once the second derivative matrix has been calculated. This is in contrast to quantum mechanical methods where the determination of the wavefunction derivatives makes analytical property calculations almost as expensive as finite difference procedures. In a dynamical simulation, the probability distribution, composed of many different nuclear configurations, is sampled to provide averaged properties that depend on temperature. This usually involves performing either molecular dynamics (in which case the time correlation between data is known) or Monte Carlo (where configurations are selected randomly according to the Boltzmann distribution). Fundamentally static and dynamic methods differ because the former are founded within the harmonic approximation, while the latter allow for anharmonicity.

For the purposes of this section, the focus will be placed on the static information that can be obtained from ionic potentials, but stoichastic simulations would also be equally as applicable. The first information to be yielded by an energy minimization is the equilibrium structure. Given that many potentials are fitted to such data, it is not surprising that the quality of structural reproduction, at least for simple binary materials, is usually high. Many force fields are derived with out explicit reference to temperature, so consequently the structure that is calculated may contain implicit temperature effects even though the optimization was performed nominally at zero Kelvin.

As an example of the application of the formal charge, shell model potential a set of parameters has been derived for alumina. The observables used consisted of the structure of corundum and its elastic and dielectric constants. As a starting model, the parameters originally derived by Catlow *et al.* [1] were used and subjected to the relax fitting approach. Alumina is a material that has been much studied already, so the aim here is just to illustrate typical results yielded by a fit to such a material and some of the related issues. Values of the calculated properties for corundum, α-Al_2O_3 are given in Table 1, along with the comparison against experiment, using the potentials derived, which are given in Table 2.

Before considering the results, let us consider the parameters that resulted from the fit since they highlight a number of points. Firstly, by looking at the shell charges and spring constants it can be seen that the oxide ion is responsible for most of the polarizability of the system as would be expected. This is a natural result of the fitting process since the charge distribution between core and shell, as well as the spring constant, was allowed to vary. Secondly, in accord with this picture the attractive dispersion term for Al–O is set to zero, though even if allowed to vary it remains small. Finally, the oxygen–oxygen

Table 1. Calculated versus experimental structure and properties for aluminium oxide in the corundum structure based on a shell model potential fitted to the same experimental data

Observable	Experiment	Calculated
a (Å)	4.7602	4.9084
c (Å)	12.9933	12.9778
Al z (frac)	0.3522	0.3597
O x (frac)	0.3062	0.2987
C_{11} (GPa)	496.9	567.1
C_{12} (GPa)	163.6	224.6
C_{13} (GPa)	110.9	158.1
C_{14} (GPa)	−23.5	−54.3
C_{33} (GPa)	498.0	453.3
C_{44} (GPa)	147.4	127.6
C_{66} (GPa)	166.7	171.2
ε_{11}^{0}	9.34	8.70
ε_{33}^{0}	11.54	13.38
$\varepsilon_{11}^{\infty}$	3.1	2.88
$\varepsilon_{33}^{\infty}$	3.1	3.06

Table 2. Interatomic potential parameters derived for alumina based on relax fitting to the experimental observables given in Table 1. The starting parameters were taken from Catlow *et al.* and a two-body cut-off distance of 16.0 Å was employed, while that for the core-shell interaction was 0.8 Å. All non-Coulombic interactions not explicitly given are implicitly zero. The shell charges for A1 and O were −0.0395 and −2.0816, respectively

Species 1	Species 2	A (eV)	ρ (Å)	C (eV/Å^6)	k_{cs} (eV/Å^2)
A1 shell	O shell	1012.17	0.32709	0.0	–
O shell	O shell	22764.00	0.14900	22.368	–
A1 core	A1 shell	–	–	–	331.958
O core	O shell	–	–	–	24.625

repulsive term is particularly short-ranged and only makes a minute contribution at the equilibrium structure. Consequently, the A and ρ values are rarely varied from the original starting values.

The rhombohedral corundum structure is sufficiently complex that even though the potential was empirically fitted to this particular system it is still not possible to achieve a perfect fit. While for many dense high symmetry ionic compounds it is possible to obtain accuracy of better than 1% for structural parameters, the moment there are appreciable anisotropic effects it becomes more difficult. This is illustrated by corundum where it is impossible with the basic shell model to accurately describe the behavior in the *ab* plane and along the *c* axis simultaneously, leading to an error of 3% in the *a* and *b* cell parameters. Not only is this true for the structure, but it is even more valid for the curvature related properties. If the values of C_{11} and C_{33} are compared, which are indicative of the elastic behavior in the two distinct directions, the calculated values have to achieve a compromise by one value being higher than experiment, while the other is lower. In reality, alumina is elastically fairly isotropic, but a dipolar model cannot capture this. The above results for alumina also illustrate the fact that while it is usually possible to reproduce structural parameters to within a few percent, the errors associated with other properties can be considerably greater.

As pointed out earlier, although a formal charge model for alumina was employed, the ions in fact behave as though the system is less than fully ionic due to the polarizability. The calculated Born effective charges show that aluminium has a reduced ionicity with a charge of +2.32 in the *ab* plane and a slightly higher value of +2.55 parallel to the *c* axis. These magnitudes are in good agreement with assessments of the degree of ionicity of corundum obtained from *ab initio* calculations.

There are many more bulk properties that can be readily determined from interatomic potentials than those given above. For instance, phonon frequencies, dispersion curves and densities of states, acoustic velocities, thermal expansion coefficients, heat capacities, entropies and free energies can all be obtained from determining the dynamical matrix about an optimized structure [6]. Other important quantities can also be determined by creating defects in the system, such as vacancies, interstitials and grain boundaries, or by locating other stationary points, in particular transition states for ion diffusion. The possibilities are as boundless as the number of physical processes that can occur in a real material.

6. Discussion

So far, the basic ionic potential approach to the modeling of solids has been portrayed. While this is very successful for many of the materials for which it was intended, and that composed the majority of the earlier studies, there are increasingly many situations where extensions and modifications are required in order to broaden the scope of the technique. These enhancements recognize the fact that many systems comprise atoms that are less than fully ionic and often non-spherical.

One of the most limiting aspects of the ionic model is the use of fixed charges. It is often the case that potential parameters are derived for the bulk material alone where a compound is at its most ionic. However, the ideal force field should also be transferable to lower coordination environments, such as surfaces and even gas phase clusters. Fundamentally, the problem with any fixed charge model, be it formally or partially charged, is that it cannot reproduce the proper dissociation limit of the interaction. Ultimately, if sufficiently far removed from each other, an ionic structure should transform into separate neutral atoms.

There is a more sophisticated way of determining partial charges within a force field that addresses the above issue, which is to calculate them as an explicit function of geometry. While this has only been sparsely utilized to date, due to the extra complexity, it has the potential to capture, through charge-transfer, many of the higher order polarizabilities beyond the dipole level, as well as yielding the proper dissociation behavior. The predominant approach to determining the charges has been via electronegativity equalization [8]. Here the self energy of an ion is expressed as a quadratic function of the charge in terms of the electronegativity, χ, and hardness, μ:

$$U_i^{\text{self}}(q) = U_i^{\text{self}}(0) + \chi_i q + \frac{1}{2}\mu_i q^2$$

When coupled to the electrostatic energy of interaction between the ions, and solved subject to the condition of charge neutrality for the unit cell, this determines the charges on the ions. The main variation between schemes is the form selected for the Coulomb interaction between ions. While some workers have used the limiting point-charge interaction of $\frac{1}{r}$ at all distances, it has been argued that damped interactions should be used that more realistically mimic the nature of two-centre integrals (i.e., tend to a constant value as $r \longrightarrow 0$).

Variable charge schemes have shown some promise, and have clear advantages since they allow multiple oxidation states to be treated with a single set of parameters, at least in principle. This simplifies the study of materials where the same cation occurs in multiple oxidation states, since no prior assumption needs to be made as to the charge ordering scheme. However, there are still many challenges in this area since it appears that choosing the more formally correct screened Coulomb interaction leads to the electrostatics only contributing weakly to the interionic forces to an extent that is unrealistic.

Looking beyond dipolar polarizability, which is a limitation of the most widely used form of ionic model, there are instances where higher order contributions are important. Here, we consider two examples that highlight the issues. Experimentally it is observed that many cubic rock salt structured materials exhibit a so-called Cauchy violation in that the elastic constants C_{12} and C_{44} are not equivalent. It has been demonstrated that two-body potential models are unable to reproduce this phenomenon, and inclusion of dipolar polarizability fails to improve the situation. The Cauchy violation actually requires a many-body coupling of the interactions through a higher order polarization. This can be handled through the inclusion of a breathing shell model. Here the shell is given a finite radius that is allowed to vary with a harmonic restoring force about an equilibrium value, with the repulsive short-ranged potential also acting on it. This non-central ion force generates a Cauchy violation, though always of one particular sign ($C_{44} > C_{12}$), while the experimental values can be in either direction.

A second example of the role of polarization, is in the stability of polymorphs of alumina. If the relative energies of alumina adopting different possible M_2O_3 structures is examined using most standard interatomic potential models, including that given in the previous section, then it is found that the corundum structure (which is the experimental ground state under ambient conditions) is not the most stable, with the bixbyite form being preferred. Investigations have demonstrated that the inclusion of quadrupolar polarizability is essential here [7]. This can be readily achieved within the point ion approach, but is more difficult in the shell model case. While an elliptical breathing shell model can capture the effect, it highlights the fact that the extension of this mechanical approach to higher order terms becomes increasingly cumbersome.

While most alkali and alkaline earth metals conform reasonably well to the ionic model, there are substantial problems with describing many of the remaining cations in the periodic table. In particular, transition metals ions are often non-spherical due to the partial occupancy of the d-orbitals. The classic example of this is when the anti-bonding e_g^* orbitals of an octahedral ion are half-filled for a particular spin, giving rise to a Jahn–Teller distortion, as is the case for Cu^{2+}. To describe this effect with a simple potential model is impossible, except by constructing a highly specific model with different short-ranged potentials for each metal–ligand interaction, regardless of the fact that they may be acting between the same species. So far, the only solution to the problem of ligand–field effects has been to resort to approaches that mimic the underlying quantum mechanics, but in an empirical fashion. Hence, most work has utilized the angular overlap model to describe a set of energy levels that are subsequently populated according to a Fermi–Dirac distribution, where the states are determined by diagonalizing a 5 × 5 matrix determined according to the local environment [11]. This approach has been successfully used to describe the manganate (Mn^{3+}, d^4) cation, as well as other systems within a molecular mechanics framework.

At the heart of the ionic potential method is the electrostatic energy, normally evaluated according to the Ewald sum when working within 3-D boundary conditions. However, this approach possesses the disadvantage that it scales at best as $N^{\frac{3}{2}}$, where N again represents the number of atoms within the simulation cell. In an era when very large scale simulations are being targeted, it is necessary to also reassess the underlying algorithms to ensure the optimal efficiency is attained. Consequently, the fundamental task of calculating the Coulomb energy is still an area of active research. Approaches currently being employed include the particle-mesh and cell multipole methods. The desirable characteristics of an algorithm are now that it should both scale linearly with system size and also be amenable to parallel computation. Both of these can be achieved as long as the method is local in real space, in some cases with complementary linear-scaling in reciprocal space, or if a hierarchical scheme is utlized within the cell multipole method to make the problem increasing coarse-grained the greater the distance of interaction is. Methods have been proposed that use a spherical cut-off in real space alone, which naturally satisfies both desirable criteria [10]. However, it becomes difficult to achieve the defined Ewald limiting value without a considerable prefactor.

7. Outlook

The state of the art in force fields for ionic materials looks set for a gradual evolution that sees it take on board many concepts from other types of system, while retaining the aim of an accurate evaluation of the electrostatic energy at the core. For the very short-ranged interactions it is likely that bond order models, widely used in the semiconductor and hydrocarbon fields, and also closely related to the approach taken for metallic systems, will be blended with schemes that capture the variation of the charge and higher order multipole moments as a function of structure. The result will be force fields that are capable of simulating not only one category of material, but several distinct ones. Development of solid state quantum mechanical methods to increased levels of accuracy will increasingly provide the wealth of information required for parameterisation of more complex interatomic potentials for systems, especially where there is a paucity of experimental data. Ultimately, this will lead to a seamless transition to models capable of reliably describing interfaces between ionic and non-ionic systems – currently one of the most challenging problems in materials science.

References

[1] C.R.A. Catlow, R. James, W.C. Mackrodt, and R.F. Stewart, "Defect energetics in α-Al_2O_3 and rutile TiO_2," *Phys. Rev. B*, 25, 1006–1026, 1982.

[2] C.R.A. Catlow and W.C. Mackrodt, "Theory of simulation methods for lattice and defect energy calculations in crystals," *Lecture Notes in Phys.*, 166, 3–20, 1982.

[3] S.W. de Leeuw, J.W. Perram, and E.R. Smith, "Simulation of electrostatic systems in periodic boundary conditions. i. lattice sums and dielectric constants," *Proc. R. Soc. London, Ser. A*, 373, 27–56, 1980.

[4] B.G. Dick and A.W. Overhauser, "Theory of the dielectric constants of alkali halide crystals," *Phys. Rev.*, 112, 90–103, 1958.

[5] J.D. Gale, "Empirical potential derivation for ionic materials," *Phil. Mag. B*, 73, 3–19, 1996.

[6] J.D. Gale and A.L. Rohl, "The general lattice utility program (GULP)," *Mol. Simul.*, 29, 291–341, 2003.

[7] P.A. Madden and M. Wilson, " 'Covalent' effects in 'ionic' systems," *Chem. Soc. Rev.*, pp. 339–350, 1996.

[8] W.J. Mortier, K. van Genechten, and J. Gasteiger, "Electronegativity equalization: applications and parameterization," *J. Am. Chem. Soc.*, 107, 829–835, 1985.

[9] M.P. Tosi, "Cohesion of ionic solids in the Born model," *Solid State Phys.*, 16, 1–120, 1964.

[10] D. Wolf, P. Keblinski, S.R. Philpot, and J. Eggebrecht, "Exact method for the simulation of Coulombic systems by spherically truncated, pairwise r^{-1} summation," *J. Chem. Phys.*, 110, 8254–8282, 1999.

[11] S.M. Woodley, P.D. Battle, C.R.A. Catlow, and J.D. Gale, "Development of a new interatomic potential for the modeling of ligand field effects," *J. Phys. Chem. B*, 105, 6824–6830, 2001.