

Recent progress in linear scaling *ab initio* electronic structure techniques

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2002 J. Phys.: Condens. Matter 14 2781

(<http://iopscience.iop.org/0953-8984/14/11/303>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 171.66.16.27

The article was downloaded on 17/05/2010 at 06:18

Please note that [terms and conditions apply](#).

Recent progress in linear scaling *ab initio* electronic structure techniques

D R Bowler¹, T Miyazaki² and M J Gillan¹

¹ Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK

² National Institute for Materials Science, 1-2-1 Sengen, Tsukuba, Ibaraki 305-0047, Japan

E-mail: david.bowler@ucl.ac.uk, miyazaki.tsuyoshi@nims.go.jp and m.gillan@ucl.ac.uk

Received 8 November 2001

Published 8 March 2002

Online at stacks.iop.org/JPhysCM/14/2781

Abstract

We describe recent progress in developing linear scaling *ab initio* electronic structure methods, referring in particular to our highly parallel code CONQUEST. After reviewing the state of the field, we present the basic ideas underlying almost all linear scaling methods, and discuss specific practical details of the implementation. We also note the connection between linear scaling methods and embedding techniques.

1. Introduction

Ab initio methods based on density functional theory (DFT) have made an immense impact on the modelling of condensed matter over the past 20 years. When combined with the pseudopotential technique, they have made it possible to study the energetics and dynamics of large systems containing hundreds of atoms. The scientific applications have been wide-ranging, covering fields as diverse as surface catalysis, high-pressure mineralogy, nanotechnology, aqueous solution chemistry, and semiconductor physics. A major effort is now under way to extend *ab initio* modelling across the length scales and timescales. This review is concerned with the treatment of very large systems containing thousands or tens of thousands of atoms. The crucial challenge here is to devise techniques for which the computational load, in terms of both cpu cycles and memory, is simply proportional to the number of atoms N , rather than increasing as some higher power. Such techniques are called linear scaling or $\mathcal{O}(N)$ (order- N) techniques. Many research groups have been working to construct practical $\mathcal{O}(N)$ codes [1–14, 16–19, 22–25], but this review will concentrate mainly on recent progress with our CONQUEST code [15, 20, 21, 26, 27], while referring to related developments by others.

With traditional DFT methods, the computational load increases at least as N^2 , and ultimately as N^3 , for reasons that have long been recognized. DFT is usually formulated in terms of the eigenfunctions (electron orbitals) $\psi_n(\mathbf{r})$ of the Kohn–Sham Hamiltonian.

Since these extend over the entire system, the amount of information contained in each ψ_n is proportional to N . The number of occupied orbitals $n = 1, 2, \dots$ is also proportional to N , so the memory requirement must go as N^2 . But traditional methods also demand the evaluation of the $\frac{1}{2}N(N+1)$ scalar products $\langle \psi_m | \psi_n \rangle$, each of which takes a number of cpu cycles proportional to N , so the computational load goes as N^3 . But this rapid increase with N must be unnecessary, since condensed-matter physics and chemistry both tell us that the properties of matter are *local*. For example, the energy of a covalent bond in an organic molecule adsorbed on a semiconductor surface does not depend on the behaviour of distant atoms. The whole notion of intensive quantities in thermodynamics rests on the validity of a local description. So it must be possible to formulate DFT—and indeed other electronic structure techniques—in a local manner, so that the cpu effort and memory are proportional to the number of atoms.

There are many reasons that $\mathcal{O}(N)$ methods are important. The first is practical: we need to be able to model large complex systems. In the rapidly emerging field of nanotechnology, realistic and accurate modelling will be vital, because some of the key processes, such as the growth and self-assembly of nanowires, quantum dots, and other nanostructures, are difficult to observe directly. The way in which quantum-based modelling can help the understanding of semiconductor growth is already clear from earlier work on, for example, the growth of the Si(001) surface from disilane [28]. But many of the large-scale complex structures that are spontaneously formed during surface growth involve systems of thousands of atoms, and are beyond traditional DFT simulation methods. The growth of Ge overlayers on Si(001) is one example [29, 30]. $\mathcal{O}(N)$ methods will also be crucial in the modelling of biomolecules, such as proteins and DNA. Until ten years ago, the large biomolecular modelling effort was based entirely on empirical interaction models, but the last few years have seen major new initiatives applying DFT-based modelling in this area (see e.g. [31]). Again, systems of thousands of atoms are often needed, and $\mathcal{O}(N)$ methods will be indispensable. The increasing convergence of biomolecular science and nanotechnology only reinforces the point. Another completely different reason that $\mathcal{O}(N)$ methods are important comes from their relation with the embedding problem—the problem of modelling accurately a limited region of a system, which is ‘embedded’ in surrounding material. The notion of locality is crucial in embedding, and it is almost certain that progress in the understanding of the $\mathcal{O}(N)$ problem will lead to progress in the embedding problem (this theme is developed further in section 2.7 and [50]). In addition, one might mention the relevance of $\mathcal{O}(N)$ methods to entirely different electronic structure techniques such as quantum Monte Carlo.

We now summarize briefly the principles on which CONQUEST is based. Quantum locality in large systems can be expressed within DFT by saying that the density matrix $\rho(\mathbf{r}, \mathbf{r}')$ decays to zero with increasing distance between points \mathbf{r} and \mathbf{r}' . This decay reflects the loss of electronic phase coherence at large distances [14, 19, 32–35]. Locality is the unifying theme of $\mathcal{O}(N)$ methods, and many of the techniques, including those embodied in CONQUEST, are based explicitly on the locality of the density matrix. In terms of the Kohn–Sham orbitals $\psi_n(\mathbf{r})$, this can be written as

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_i f_i \psi_i^*(\mathbf{r}) \psi_i(\mathbf{r}') \quad (1)$$

where f_i are the occupation numbers. The Kohn–Sham total energy can be rewritten in terms of the density matrix [20]: the Hartree and exchange–correlation energies, which are written in terms of the charge density ($n(\mathbf{r}) = \rho(\mathbf{r}, \mathbf{r})$) do not change; the kinetic and pseudopotential energies become

$$E_{\text{KE}} = -\frac{\hbar^2}{2m} \int d\mathbf{r} (\nabla_r^2 \rho(\mathbf{r}, \mathbf{r}'))_{r=r'} \quad (2)$$

$$E_{\text{ps}} = 2 \int d\mathbf{r} d\mathbf{r}' V_{\text{ps}}(\mathbf{r}, \mathbf{r}') \rho(\mathbf{r}, \mathbf{r}'). \quad (3)$$

The fact that $\rho(\mathbf{r}, \mathbf{r}')$ tends to zero as $|\mathbf{r} - \mathbf{r}'| \rightarrow \infty$, exponentially for insulators or algebraically for metals [33–37], implies that the amount of information scales linearly with the size of the system; an $\mathcal{O}(N)$ method can be created by making an approximation, and *enforcing* locality: $\rho(\mathbf{r}, \mathbf{r}') = 0$, $|\mathbf{r} - \mathbf{r}'| > R_c$. While the scaling of both memory and computational effort will allow large systems to be simulated on a workstation, for very large systems containing thousands or tens of thousands of atoms, the codes need to run efficiently on parallel computers; this aspect is discussed later in the article.

However, the six-dimensional quantity $\rho(\mathbf{r}, \mathbf{r}')$ is not the ideal variable to work with, so in CONQUEST the assumption is made (with only the restriction that the original quantity had a finite number of non-zero eigenvalues) that it can be written in *separable* form:

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_{i\alpha, j\beta} \phi_{i\alpha}(\mathbf{r}) K_{i\alpha j\beta} \phi_{j\beta}(\mathbf{r}') \quad (4)$$

where $\phi_{i\alpha}(\mathbf{r})$ is a *support function* (or a *localized orbital*) centred on atom i , and $K_{i\alpha j\beta}$ is the density matrix in the basis of the support functions. Then locality can be enforced by applying separate cut-offs to $K_{i\alpha j\beta}$ and the set of *support functions* $\{\phi_{i\alpha}(\mathbf{r})\}$ [20]. The minimization of the energy with respect to each of these quantities then drives the system towards the ground state.

The support functions are *non-orthogonal*, and forced to be confined within a *localization region*, of radius R_c . They are themselves represented in terms of other basis functions (which will be described in detail in section 3), and it is important that they be freely varied in the search for the electronic ground state. By increasing the cut-offs and improving the completeness of the basis set representing the support functions, an $\mathcal{O}(N)$ method can be made to reproduce traditional methods with plane-wave accuracy [21, 27].

The minimization in CONQUEST involves three different variables: the elements of the density matrix, K ; the support functions, $\phi_{i\alpha}$; and the charge density, $n(\mathbf{r})$, which must be consistent with the potential $V(\mathbf{r})$. We have chosen to decouple these variables, fixing first the support functions and the charge density, and minimizing with respect to the density matrix, then achieving self-consistency (while minimizing with respect to the density matrix every time that the potential changes), and finally varying the support functions [27]. This scheme has various advantages: first, it decouples the degrees of freedom associated with the support functions and the density matrix; second, it allows us to treat regions of the system with fixed support functions (in an *ab initio* tight-binding manner) while freely varying others; third, on a practical level, it allows us to optimize the procedures for the different minimizations independently.

The rest of the article is arranged as follows: in the next section, we describe minimization of the energy with respect to the density matrix, K ; then we consider possible basis sets for the support functions, and describe our choice; we then describe some practical details relating to the implementation of our scheme and conclude the article.

2. Finding the density matrix

Within the atom-centred basis of support functions, the density matrix $K_{i\alpha j\beta}$ (also called the *kernel*) is clearly equivalent to the density matrix in a non-orthogonal tight-binding formulation. There has been a great deal of work investigating effective $\mathcal{O}(N)$ methods for finding the density

matrix in tight-binding form [1, 2, 4–10, 17, 18, 21, 22, 26] which gives us a strong position to start from. However, there is an important issue, which will be addressed fully below: in the formulation described up to this point, the basis in which the density matrix, $K_{i\alpha j\beta}$, is written is *non-orthogonal*, while most tight-binding methods have been primarily formulated with an orthogonal basis set. We shall first describe the methods, and then address this important question of non-orthogonality.

Tight-binding techniques which have been extended to *ab initio* techniques fall broadly into four categories: recursion [10]; density matrix minimization (DMM) [5, 6, 12, 26]; orbital minimization [7, 9]; and penalty functions [25]. These have been described in some detail by Goedecker [1], so we shall only outline the methods.

2.1. The Fermi operator expansion

The Fermi operator expansion technique [10, 38–40] is a conceptually and computationally simple way of obtaining the density matrix, at the expense of introducing a finite electronic temperature and losing a variational principle. The Fermi operator (a finite-temperature density matrix) can be defined as

$$\hat{F}_{\mu,T} = f((\hat{H} - \mu)/kT) \quad (5)$$

where $f(x) = 1/(1 + \exp(x))$ is the Fermi function.

Now, the Fermi function only has to cover a finite energy range, namely the width of the density of states for the system in question (or the difference between the minimum and maximum eigenvalues of the Hamiltonian). Within this range, it can be represented by a polynomial in the energy:

$$f(x) = \sum_{p=0}^{n_{pl}} C_p E^p \quad (6)$$

which means that the Fermi matrix $F_{\mu,T}$ can be represented as a polynomial in the Hamiltonian:

$$\hat{F}_{\mu,T} = \sum_{p=0}^{n_{pl}} C_p \hat{H}^p. \quad (7)$$

This then gives the expression for one element of the Fermi matrix as

$$\langle i\alpha | \hat{F}_{\mu,T} | j\beta \rangle = \sum_{p=0}^{n_{pl}} C_p \langle i\alpha | \hat{H}^p | j\beta \rangle. \quad (8)$$

Conceptually, then, the method works by fitting a polynomial to the Fermi function over the range of the eigenvalues. Then, using the coefficients of this polynomial and moments of the Hamiltonian, elements of the finite-temperature density matrix, or the Fermi matrix, are constructed. To make the method $\mathcal{O}(N)$, the Fermi operator can be truncated beyond a certain cut-off radius. In practice, for stability, a Chebyshev polynomial is used [38], which leads to a recursion relation for the coefficients:

$$p_{\mu,T}(H) = \frac{c_0}{2} + \sum_{j=1}^{n_{pl}} c_j T_j(H) \quad (9)$$

and

$$\begin{aligned} T_0(H) &= I \\ T_1(H) &= H \\ T_{j+1} &= 2HT_j(H) - T_{j-1}(H). \end{aligned} \quad (10)$$

Once the Fermi operator has been truncated, the forces are not exactly equal to the derivative of energy; there is, however, a formalism which gives a force which is the exact derivative of the energy [39]. In some cases the error in energy due to the high electronic temperature may be significant; a scheme is available [41] for extrapolating the $T = 0$ energy from a high temperature which can correct this.

2.2. Density matrix minimization

DMM [5, 12, 23, 24, 26] seeks to find the density matrix by minimizing the energy with respect to the density matrix elements (since $E_{\text{band}} = 2 \text{Tr}[\rho H]$). However, two constraints must be applied to the minimization: (i) either constant electron number or constant Fermi energy; (ii) idempotency of ρ . The first is relatively easy to address; maintaining constant Fermi energy is as simple as minimizing $2 \text{Tr}[(\rho - \mu I)H]$, with μ the Fermi energy, while various schemes exist for maintaining the electron number constant [42, 43].

Idempotency of ρ (which is equivalent to requiring that the eigenvalues of ρ all be either zero or one, or that $\rho^2 = \rho$) is a much harder constraint to impose during a variational minimization, and instead use is made of McWeeny's purification transformation [44]:

$$\rho = 3\tilde{\rho}^2 - 2\tilde{\rho}^3. \quad (11)$$

Provided that the eigenvalues of $\tilde{\rho}$ lie between $-\frac{1}{2}$ and $\frac{3}{2}$, the eigenvalues of ρ will lie between zero and one. McWeeny first proposed this as an iterative procedure (so that if $\rho_n = \tilde{\rho}$ then $\rho_{n+1} = \rho$), but another technique [5] is to minimize $E_{\text{band}} = 2 \text{Tr}[\rho H]$ with respect to the elements of $\tilde{\rho}$. If this approach is taken, then we must assume that $\tilde{\rho}$ is also separable (as ρ is in equation (4)): $\tilde{\rho} = \sum_{i\alpha j\beta} \phi_{i\alpha}(\mathbf{r}) L_{i\alpha j\beta} \phi_{j\beta}$; then we can write $K = 3L^2 - 2L^3$. Approaches to DMM are varied: pure McWeeny iteration can be used [45] (also with a modified cubic form which preserves electron number); pure minimization can be used [5, 6, 42]; minimization followed by [24] or interspersed with [23] McWeeny purification is also pursued; finally, McWeeny purification (which is not variational) followed by minimization (which is variational) [26]. It can be shown [26] that the McWeeny purification approaches a manifold of idempotent density matrices perpendicularly in its final stages, while the minimization yields a gradient tangential to this surface (and preserves idempotency to first order).

2.3. Orbital minimization

Another method for finding the ground-state density matrix is orbital minimization, which was proposed from two different routes, leading to essentially the same formalism [7, 9]. Consider a system of N electrons, described by $N/2$ non-interacting states $\{|\psi_i\rangle\}$. In order to avoid an *explicit* orthonormalization step (which scales with N^2 in a localized basis set and N^3 in a plane-wave basis), the following functional is defined:

$$E = 2 \text{Tr}[QH] - \eta [2 \text{Tr}[QS] - N] \quad (12)$$

$$Q = \sum_{n=0}^N (I - S)^n \quad (13)$$

where $S_{ij} = \langle \psi_i | \psi_j \rangle$ is the overlap matrix, and the expansion for Q is truncated (typically at $n = 1$). The orbitals $\{|\psi_i\rangle\}$ are represented by a basis $\{|\phi_\mu\rangle\}$, so

$$|\psi_i\rangle = \sum_{\mu} C_{i\mu} |\phi_\mu\rangle. \quad (14)$$

Then as the energy is minimized with respect to the coefficients $C_{i\mu}$, the overlap matrix will tend to the identity. The method can be made $\mathcal{O}(N)$ by localizing the orbitals ψ_i , and only allowing contributions from ϕ_μ within a specified volume.

The drawback with the method is that it is subject to many local minima if the minimal set of orbitals ($N/2$) is used; it has been extended to more orbitals, which to some extent corrects this problem [46].

2.4. Penalty functionals

Penalty functionals apply a similar idea to orbital minimization techniques (which effectively penalize non-orthogonality), but instead seek to penalize the deviation away from idempotency [19]; these are methods which actively seek to impose idempotency directly. The original technique [19] defined the following functional:

$$Q[\rho; \mu, \alpha] = E_{\text{NI}}[\rho] - \mu N[\rho] + \alpha P[\rho] \quad (15)$$

$$P[\rho] = \{\rho^2(1 - \rho^2)\}^{\frac{1}{2}} \quad (16)$$

where E_{NI} is the non-interacting energy. Then the ground state was found by varying ρ and seeking the lowest value of Q , in a standard manner.

However, it was found [47] that the branch point introduced by the square root prevented minimization using techniques such as conjugate gradients. When performing a variational minimization, this constitutes a large handicap. Instead, a general functional was introduced [25]:

$$Q[\rho] = E[\rho] + \alpha P^2[\rho] \quad (17)$$

which allows use of minimization techniques.

2.5. Non-orthogonality

All of the above methods have been described using orthogonal bases, but the formalism above relies on non-orthogonal bases. Each of the methods can be reformulated in a non-orthogonal basis, but this raises various issues:

- (i) Matrix expressions become more complex (e.g. we now have $K = 3LSL - 2LSLSL$, where S is the overlap matrix).
- (ii) Compound matrices (such as LSL) are longer ranged.
- (iii) When using variational methods, the inverse of the overlap, S^{-1} , is required to correct the gradients [48].

Some choose to work within the non-orthogonal basis, and derive an approximate S^{-1} [27,40], while others convert to an orthogonal basis, using a variety of methods (incomplete inverse Cholesky factorization [24], Cholesky decomposition [49]). All of these approaches require an approximation (either in the S^{-1} -factorization or in the Cholesky factorization), with the former having the advantage that no basis set changes are being used, and the latter that the formalism is much simpler and shorter ranged.

2.6. Examples from Conquest

We illustrate the effectiveness of $\mathcal{O}(N)$ methods, and the convergence with density matrix cutoff, by exploring two systems with Conquest: a carbon vacancy and a silicon interstitial. For both these systems, we are interested only in the convergence of the energy with respect to the density matrix, and so fix the support functions. The energy convergence (shown as a deviation from the large radius converged result) for the carbon vacancy is shown in figure 1, which shows that the energy is converged to within 0.1 eV with a radius of 6 Å, while the energy convergence for the silicon interstitial is shown in figure 2, showing




Figure 1. Convergence of self-consistent ground state energy with density matrix radius for a vacancy in carbon (63 atom cell), calculated using CONQUEST on an NEC-SX5. Energies are given relative to the converged, large radius L matrix result.




Figure 2. Convergence of self-consistent ground state energy with density matrix radius for an interstitial in silicon (65 atom cell), calculated using Conquest on an NEC-SX5. Energies are given relative to the converged, large radius L matrix result.

convergence to within 0.1 eV with a radius of 9 Å. Both of these results demonstrate that for real systems, good energy convergence can be achieved for relatively modest cutoffs.

2.7. The connection with embedding

We finish by noting that the localization inherent in $\mathcal{O}(N)$ schemes makes them ideal for use in embedding [50]. By this, we mean embedding of one system within another (e.g. a point

defect into perfect bulk), not embedding of one technique within another. If the system is divided into region I (the region of interest—e.g. the point defect and its surroundings) and region II (the embedding matrix—e.g. the perfect bulk) then the amount of region II required is equivalent to the range of the density matrix, and we expect the energy convergence with size of region I to scale just as the energy convergence scales with density matrix range. This is illustrated in figure 3, which shows the convergence of the energy for a Ge substitutional defect in diamond Si with radius of region I for a tight-binding implementation of the embedding technique [50].




Figure 3. Convergence of energy with region I radius for a Ge substitutional impurity in Si expressed relative to the result for an infinite region I radius.

3. Localized orbitals

3.1. Representing the localized orbitals

As emphasized in section 1, the localized orbitals $\phi_{i\alpha}(\mathbf{r})$ have to be freely varied in the search for the DFT ground state. This raises the technical problem common to all quantum calculations—the representation of the orbitals, i.e. basis sets. A peculiar feature of the basis sets needed in the scheme that we have outlined is that the orbitals to be represented vanish outside the localization region, and the basis functions clearly need to have the same property. Before discussing the approaches to this that have been adopted, it is useful to consider briefly the conflicting requirements that basis functions have to satisfy.

First, they should ideally be well adapted to the function to be represented, which means that only a *few* basis functions should be needed to represent the orbitals. Second, the representation should be systematically improvable; this means not just that the basis set must be complete, but also that the convergence should be rapid as the size of the basis set is increased. As a rider to this, the computational effort should not increase too fast as one increases the number of basis functions, and ideally no faster than linearly. Third, the operations to be done with the basis functions should be mathematically simple, so that the number of computer operations is small. Fourth, since some parts of a total-energy calculation have to be

done on spatial grid, which generally causes problems like breaking of translational symmetry, the basis must be designed so that as little as possible has to be done on the grid.

The plane-wave basis sets generally used with the DFT/pseudopotential approach fail miserably under the first criterion: they are poorly adapted to the representation of orbitals, and even for a simple element like Si, typically 100 plane waves per atom are needed. In spite of this, they are the basis of choice for high-precision work because they satisfy the second and third criteria outstandingly well, and for ease of use these are more important than the first. Localized basis sets like Gaussians are better under the first criterion (typically, 15–20 would be needed for Si), and are satisfactory under the third, but run into serious troubles under the second. Pseudo-atomic orbitals satisfy the first criterion even better, but have similar problems with the second. In spite of these drawbacks, localized basis states are important and widely used when high precision is not needed, because they require less memory and can be performed with modest computer resources.

Several types of basis set have been used or proposed for linear scaling DFT calculations [51, 53, 55]. When CONQUEST was first written, the support functions were represented as numerical values on a grid [15]. A grid basis set gives a natural way of representing orbitals that vanish outside specified regions. Nevertheless, it is the ultimate in maladaptation, and needs fine grids and large amounts of memory in order to achieve good precision. The kinetic energy is particularly troublesome. In more conventional DFT/pseudopotential schemes based on grid basis sets, high-order finite-difference methods are used to give an accurate representation of the Laplacian operator [52]. Because of this, grid basis sets were replaced in CONQUEST by a scheme akin to finite elements.

This finite-element scheme represents the $\phi_{i\alpha}(\mathbf{r})$ in terms of piecewise-continuous polynomials, using a technique sometimes referred to as *B*-splines. Full details of the scheme, with demonstrations of its effectiveness, are presented in a published report [51], so here we give only a brief summary. Suppose first that we have a continuous function $f(x)$ in one dimension, which we wish to represent. The *B*-spline basis consists of localized functions $\theta_s(x)$, centred on the points of a grid, whose nodes are at positions $X_s = sa$, where a is the grid spacing. The basis functions are all images of each other, displaced by an integer number of grid spacings, so $\theta_s(x) = \theta_0(x - X_s)$. The basis function $\theta_0(x)$ vanishes identically outside the range $-2a < x < 2a$. Inside this range, it is put together from cubic polynomials:

$$\theta_0(x) = \begin{cases} 1 - \frac{3}{2}x^2 + \frac{3}{4}|x|^3 & \text{if } 0 < |x| < a \\ \frac{1}{4}(2 - |x|)^3 & \text{if } a < |x| < 2a \\ 0 & \text{if } 2a < |x| \end{cases} \quad (18)$$

and has the property that it and its first two derivatives are continuous everywhere. In fact, the only discontinuities are in the third derivative at the points $|x| = 0, a, \text{ and } 2a$. The representation of a continuous function

$$f(x) \simeq \sum_s b_s \theta_s(x) \quad (19)$$

can be made arbitrarily precise by systematically reducing the grid spacing a . This is exactly analogous to increasing the plane-wave cut-off G_{max} when taking a plane-wave calculations to convergence.

In fact, there is a close relationship between *B*-spline and plane-wave basis sets. The $\theta_s(x)$ basis functions can be used to form Bloch-like functions $\chi_k(x)$ by the unitary transformation

$$\chi_k(x) = \sum_s e^{ikX_s} \theta_s(x). \quad (20)$$

To obtain the full set of distinct χ_k -functions, k should be restricted to the range $-\pi/a < k < \pi/a$. As $|k| \rightarrow 0$, the functions become identical to plane waves, and in fact they

rather precisely reproduce plane waves except near the ends of the interval $(-\pi/a, \pi/a)$. This means that B -splines with grid spacing a are nearly equivalent to plane waves with cut-off $G_{\max} = \pi/a$.

In practice, of course, we work in three dimensions, and the three-dimensional B -splines $\Theta_s(\mathbf{r})$ are defined as Cartesian products:

$$\Theta(\mathbf{r} - \mathbf{R}_s) = \theta(x - X_s)\theta(y - Y_s)\theta(z - Z_s) \quad (21)$$

where (X_s, Y_s, Z_s) are the Cartesian components of \mathbf{R}_s , and the support functions are represented as

$$\phi_{i\alpha}(\mathbf{r}) = \sum_s b_{i\alpha s} \Theta_s(\mathbf{r} - \mathbf{R}_{is}). \quad (22)$$

In the current scheme, the blip grid on which the $\Theta_s(\mathbf{r})$ are sited is defined separately for each atom, and moves with that atom. To enforce the vanishing of $\phi_{i\alpha}(\mathbf{r})$ outside the support region, we include in equation (22) only those $\Theta_s(\mathbf{r})$ that are non-zero only for points within the region. The reason for making the blip grid move with the atom is that this ensures that each $\phi_{i\alpha}(\mathbf{r})$ is represented always in terms of the same set of basis functions.

Blip functions therefore give us a scheme that is closely related to plane waves, but at the same time respects the strict localization of the support functions. It also shares another feature with plane waves, and that is that as the blip spacing is decreased, the computational effort grows linearly only with the number of blip functions. This is because the number of blip functions that are non-zero at each point in space does not increase as a decreases.

But this is certainly not the only spatially localized basis set that is closely related to plane waves. An alternative is the spherical-wave basis proposed by Haynes and Payne [53]. Spherical waves are the energy eigenfunctions of a particle confined to a spherical box, where the radius of the box is the support region radius R_{reg} . These eigenfunctions have definite angular momentum quantum number l and magnetic quantum number m , and can be written as

$$\Omega_{nlm}(\mathbf{r}) = j_l(k_n r) Y_l^m(\theta, \phi) \quad (23)$$

where $Y_l^m(\theta, \phi)$ is the appropriate spherical harmonic, depending on the polar and azimuthal angles θ and ϕ , and $j_l(k_n r)$ is the spherical Bessel function depending on radial distance r , with the allowed wavevectors k_n fixed by the condition that the basis functions go to zero at the boundary of the support region: $j_l(k_n R_{\text{reg}}) = 0$. The properties of this basis set have been explored in detail in [53, 54], where tests on molecules (H_2 , HCl , Cl_2 , SiH_4) and bulk silicon show reasonable convergence properties relative to plane-wave results.

This scheme has the nice feature that a total-energy calculation can be converged in exactly the same way as a plane-wave calculation, by systematically increasing the plane-wave cut-off wavevector G_{\max} . However, it has a serious objection, which may make it difficult to use in practical calculations. This is that as the size of the basis set is increased, the computational effort grows as the square of the number of basis functions. This is because the number of basis functions that are non-zero at any point in space is proportional to the number of basis functions, so the matrix elements $\langle \phi_{i\alpha} | O | \phi_{j\beta} \rangle$ of any operator O for any two atoms i and j grow very rapidly. Since we know in advance that at least 100 spherical waves will be needed per atom, the number of operations needed to calculate each such matrix element is likely to be at least 10^4 , and this may make the calculations too slow.

Finally, we mention the pseudo-atomic basis sets used in SIESTA [55], keeping our remarks brief, since a separate article in this special issue is devoted to SIESTA. The basic philosophy is similar to that originally developed by Sankey and Niklewski [56]. The basis functions are the atomic orbitals obtained from a self-consistent DFT calculation on the free

atom, except that, in order to ensure that the functions vanish identically outside the region radius, the free-atom calculation is done in the presence of a confining potential. In the original Sankey–Niklewski scheme, the confining potential is simply an infinite potential beyond the radius R_{reg} . This makes the localized orbitals go to zero linearly as $r \rightarrow R_{\text{reg}}$, so there is a discontinuity in the first derivative, which may well exacerbate the breaking of translational symmetry in the parts of the calculation done on a spatial grid. This gives a motivation for making the confining potential go to infinity more continuously as $r \rightarrow R_{\text{reg}}$, and this freedom is being exploited in the latest SIESTA basis sets [57]. In addition, in order to obtain satisfactory precision, it is essential to go beyond minimal basis sets, and to include at least two basis functions for each angular momentum (so-called ‘double-zeta’ basis sets), and to include also polarization functions.

The different approaches to the problem of basis sets taken with linear scaling DFT schemes thus reflect the tension between the four criteria outlined at the start of this section, and particularly the tension between good adaptation of the basis set to the form of the orbitals, and resulting economy in memory use, on the one hand, and systematic improvability on the other. As in more conventional DFT methods, there cannot be a ‘best’ approach to basis sets, since the physical problem being addressed and the resources available will place different weights on the criteria. Our view is therefore that there is great merit in a flexible approach, in which different types of basis set are employed for different problems, or at different stages of a given problem. We also believe that it may be possible to combine different schemes, for example pseudo-atomic orbitals and *B*-splines, in the same way as mixed basis sets have long been used in conventional DFT/pseudopotential calculations.

3.2. Blip operations

As described above, the support functions are represented in a basis of blip functions (or *B*-splines), defined on a grid that moves with each atom. We need to perform integrals involving the support functions to generate matrix elements (such as $S_{i\alpha j\beta} = \int d\mathbf{r} \phi_{i\alpha}(\mathbf{r})\phi_{j\beta}(\mathbf{r})$). For the overlap matrix and the kinetic energy part of the Hamiltonian, the integration can be performed analytically in terms of the $b_{i\alpha s}$ -coefficients. For example, $S_{i\alpha, j\beta}$ can be expressed as

$$S_{i\alpha, j\beta} = \sum_{m,n} b_{i\alpha m} b_{j\beta n} s_{im, jn} \quad (24)$$

where

$$s_{im, jn} = \int d\mathbf{r} \Theta_{im} \Theta_{jn}. \quad (25)$$

However, some parts of the Hamiltonian matrix cannot be calculated analytically, and integration must be approximated by summation on a grid. This ‘integration grid’ is completely distinct from the blip grid, and is a single fixed grid covering the whole simulation cell. For a uniform cubic grid of spacing h_{int} , a matrix element such $S_{i\alpha, j\beta}$ would be approximated as

$$S_{i\alpha, j\beta} \simeq \delta\omega_{\text{int}} \sum_{\ell} \phi_{i\alpha}(\mathbf{r}_{\ell})\phi_{j\beta}(\mathbf{r}_{\ell}) \quad (26)$$

where $\delta\omega_{\text{int}} = h_{\text{int}}^3$ is the volume per grid point, and \mathbf{r}_{ℓ} is the position of the ℓ th grid point. As shown elsewhere [20], h_{int} should generally be about half of h_{blip} .

Analytic evaluation, if possible, is preferable, but the double summation required (see equation (24)) imposes a computational cost. Our strategy is to evaluate analytically the on-site ($i = j$) matrix elements of overlap and kinetic energy, and to use grid summation for all others. The thinking here is that the on-site terms are large, so accuracy is important; but there are few of them, so the cost of analytic evaluation is small.

The transformation from the coefficients $b_{i\alpha s}$ to the values of $\phi_{i\alpha}(r_l)$, where r_l is a grid point, is called a blip-to-grid transform [21, 51]—using the separable form of blips shown above in equation (21), this is extremely similar in concept to a FFT, and can be evaluated extremely efficiently. The integration can be performed efficiently by creating small blocks of integration grid points, and making partial contributions to matrix elements between all atoms touching the block with a single BLAS call, as discussed below.

4. Practicalities

In this section, we address the practical implementation of CONQUEST on massively parallel machines, as well as some of the practical problems which we have encountered in the course of writing CONQUEST, related both to efficiency on parallel computers and to robustness [21, 27, 58].

4.1. Implementation

The division of workload between processors is an important part of all parallel codes—indeed, load balancing is a large subject in its own right. Nominally, the computational effort and storage requirements of CONQUEST are divided up as follows:

- (i) Every processor has responsibility for a group of atoms (storing the blip coefficients for each atom and transforming their values onto the integration grid)—the primary set.
- (ii) Every processor has responsibility for rows of matrices of these atoms (storing values and performing multiplications for these rows).
- (iii) Every processor has responsibility for an area of the integration grid (storing data on this area and performing integrations)—the domain.

Practically, the assignment of groups of atoms and areas of the integration grid will have a large effect on the efficiency of the code—typically, we want the groups of atoms and grid points to be compact and local and we want the two groups to overlap as much as possible (to restrict communication). Below, we describe a key technique in achieving flexibility in load balancing as well as efficiency in computation—small groups.

4.2. The use of small groups

In CONQUEST, there are two key areas of effort: matrix multiplication, and grid operations (integration, blip-to-grid transforms, etc). In both of these areas, there are two natural levels of organization: individual (e.g. grid point or matrix element); and global (all atoms or grid points). We have found that it is vital for efficiency to create an intermediate level of organization—small groups of the entities (we call a small group of atoms a *partition* and a small group of integration grid points a *block*).

To understand the use of small groups, let us take an example of matrix multiplication, where we perform the summation

$$C_{ij} = \sum_k A_{ik} B_{kj}. \quad (27)$$

Here, each processor is responsible for calculating all elements C_{ij} for atoms i for which it is responsible—its primary set. It already has the values A_{ik} stored locally, but will have to fetch the values B_{kj} for the atoms k outside its primary set. Following the two natural levels mentioned above, there are two extremes that we can consider for fetching the elements B_{kj} (and hence interleaving communication and calculation): individually (fetch a single element,

compute the partial contribution to C_{ij} , repeat—fine interleaving); and globally (fetch all elements B_{kj} which will be required, then do all computations—coarse interleaving). The first of these will be overly expensive in communications (all communication involves a latency or start-up cost, so a significant gain is made by transferring long messages) while the second will potentially be expensive in memory. Somewhere between these two extremes there will be a good balance between memory required on-processor, and the latency of short communications.

Partitions of atoms (and hence matrix elements) are extremely useful as they simplify the task of finding the compromise between the two extremes given above. For example, if a processor is responsible for several partitions, then transferring the B_{kj} according to the partitions will split up the calculation in a natural way; this will be explored more in the next section, and has been extensively discussed in a recent paper [58]. Another area where these groups are helpful is in integration: we use the highly optimized BLAS routines on all integration grid points in a block to yield extremely efficient integration routines; this is touched on later.

4.3. Matrix multiplication

Recently, we have studied the efficiency of sparse matrix multiplication on highly parallel machines [58]. Each processor takes responsibility for several partitions of atoms, formed into its primary set. For a given matrix multiplication, we form a *halo* consisting of all atoms (or partitions) within range of *any* primary set atom, and loop over these atoms during the multiplication. It is also helpful to form a *covering set*: a super-set of different matrix haloes. This simplifies searches and indexing for various different matrix multiplications. We have also identified a multiplication *kernel*: a piece of code that is repeatedly called, and which can be optimized on different machines. We have achieved 10% of peak speed on a Cray T3E, which is respectable given the pattern of matrices.

We illustrate the practical performance of the CONQUEST code in figure 4. Here we show the time taken for various aspects of the calculation for increasing system sizes on a Cray T3E-1200. As the technique is variational, the time to self-consistency will improve after the slow, initial search. We see that for fixed number of atoms per processor, the time taken is almost constant, showing that the matrix multiplication (which forms the bulk of the workload for these calculations) scales extremely well in parallel.

In the work illustrated above, and in the detailed matrix multiplication code described in [58], we have assumed a superscalar, RISC-type architecture (such as that found in the Cray T3E or SGI Origin 2000 or 3000). However, there is another type of machine commonly used for large calculations, the so-called vector-parallel machine. Most of these machines have no cache for their vector units, which means that the fundamental assumption of our recent work (i.e. of cache re-use) is undermined. In fact, the code shown in [58] can only reach $\sim 1\%$ of the peak performance on an NEC-SX5. So we need to create a new code to achieve high performance on vector-parallel machines, while making the changes as limited as possible to maintain portability.

Accordingly, we have created a new matrix multiplication kernel (that small part of the code that is repeatedly performed to build up the product matrix) specifically for vector-parallel machines. This code can be divided into two parts: the calculation of the locations of the arrays (both the result matrix and the input matrices); and the actual matrix multiplication operations. We obtain efficiency by vectorizing the latter part; the former part is rather cheap, as shown elsewhere [58]. The resultant new kernel typically achieves $\sim 20\text{--}30\%$ of the peak speed of the NEC-SX5 (peak speed being 8 GFlops), which we believe is entirely respectable for

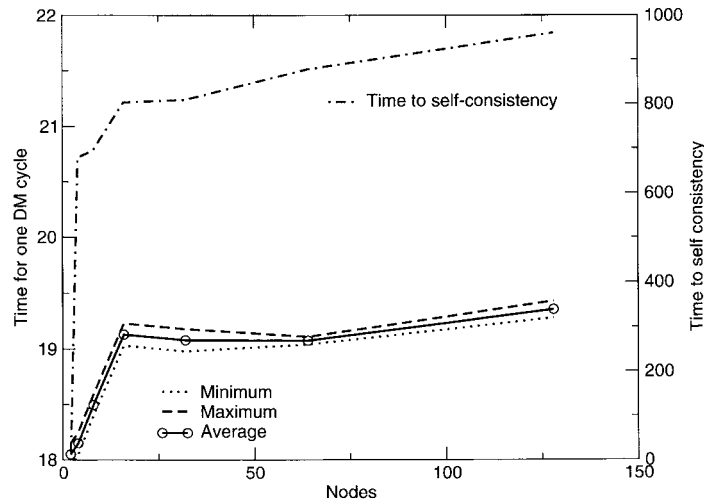


Figure 4. Time taken for CONQUEST calculations for a single DMM (left axis, bottom three curves) and for a complete search for self-consistency, starting from scratch, using fixed support functions.

distributed, sparse matrix multiplication.

There are two points worth making about this new kernel. First, the code has a different structure for ‘maximal’ and ‘minimal’ multiplication (in marked contrast to the original, superscalar code, which maintained the symmetry between these operations). This change is due to the difference in the operations which can be vectorized in the two cases. Second, a work array is required for vectorization which might be prohibitive for superscalar architectures, but not for vector-parallel machines, which typically have larger amounts of memory per processor. The size of this work array can be controlled by changing the size of the partitions, which thus play important roles in both superscalar and vector-parallel calculations, but in very different ways.

4.4. Integration and blip transforms

The new method of matrix multiplications shown in the last section has improved the efficiency of CONQUEST dramatically. It makes it possible to do much larger calculations than before. As a result, requirements for other parts of the code become harder. Recently we have changed a part of the code, which performs the calculation of matrix elements and blip–grid transforms explained in [21]. In this section we will briefly show the changes we made.

Let us first consider what kinds of operation that this requires. The matrix elements made from support functions, projector functions, or their derivatives such as $S_{i\alpha,j\beta} = \langle \phi_{i\alpha} | \phi_{j\beta} \rangle$, $P_{i\alpha,j\beta} = \langle \phi_{i\alpha} | \chi_{j\beta} \rangle$, and $\langle \nabla \phi_{i\alpha} | \nabla \phi_{j\beta} \rangle$, are calculated by a summation over integration grid points. For example, $S_{i\alpha,j\beta}$ is calculated from

$$S_{i\alpha,j\beta} = \delta\omega_{\text{int}} \sum_{r_l} \phi_{i\alpha}(r_l) \phi_{j\beta}(r_l). \quad (28)$$

Here, ω_{int} is the volume per grid point and r_l is integration grid which is common to the support region of i and that of j . The above summation for a given set of indices (i, α) and (j, β) can be regarded as a matrix multiplication, where r_l serves as the column index of the left matrix and as the row index of the right matrix. As explained in [21], we use a BLAS-3

Table 1. How small groups are made up for different members.

Members	Small groups	Primary sets
Atoms	Partitions	Bundles
Integration grid points	Blocks	Domains

routine ‘dgemm’ to do these matrix multiplications and we have introduced the term ‘blocks’ for the effective use of this routine. A block is an assembly of integration grid points and the above summation are divided into the summation of partial contributions from all blocks and the calculation of the partial contribution by dgemm for each block. As we will see later again, we refer to an atom whose support region contains at least one integration grid point in a block b as a neighbour atom of a block b for support functions. For these calculations, we must have the list of the pairs of atoms (i, j) for each block b , both of which are neighbour atoms of the block. In practice, a block is a cuboid³ containing $n_x \times n_y \times n_z$ points and its size should be determined by comparing the gain in the speed of dgemm by increasing the size of matrices and the loss by unnecessary operations from zero values in $\phi_{i\alpha}(r_l)$.

To perform the above summation, we also must know the values of support or projector functions on the integration grid—the blip-to-grid transform shown in (22). Blip-grid transforms are performed only for r_l in the blocks which include one or more integration grid points in the support region of the atom i . We refer to these blocks as neighbour blocks of the atom i .

In short, we need to consider the optimal way of performing the following tasks:

- (i) Making lists and tables to perform blip-grid transforms and the calculation of matrix elements: making lists of neighbour atoms of blocks; making lists of neighbour blocks of atoms; and so on.
- (ii) Blip-grid transformation.
- (iii) Calculation of matrix elements.

It should be noted that the matrices calculated in this way are used to calculate other matrices: LSL , $LSLSL$, $LSLH$, and so on.

In the new method, we use the same data structure (small groups) as shown in the previous sections. In the method for performing matrix multiplications, each node has a set of small groups called partitions, and each partition has its members, i.e. atoms. Obviously, we can regard a block as a small group of integration grid points. Each processor is responsible for a set of blocks which we call the domain. This is summarized in table 1.

Each processor has one domain and one bundle. Hereafter, we refer to a node which performs operations with respect to atoms as a bundle-responsible node, and the node doing operations relating to the integration grid as a domain-responsible node. In performing the three tasks described above, we have a lot of communications between bundle-responsible nodes and domain-responsible nodes. In blip-grid transforms, for example, bundle-responsible nodes first calculate $\phi_{i\alpha}(r_l)$ for integration grid points r_l in the support region of the atoms i , because bundle-responsible nodes have a set of blip coefficients $\{b_{i\alpha s}\}$. Then, they must send these values to domain-responsible nodes which have integration grid points r_l . In the calculation of matrix elements, each domain-responsible node accumulates the partial contributions to matrix elements from all blocks in the domain, and these contributions are sent to bundle-responsible nodes, which accumulate the contributions sent from their halo nodes to make their matrix

³ This is not necessary in principle, and non-orthorhombic cells and hence blocks will be addressed in the future.

elements. How to organize these communications efficiently is the key point for performing the operations in this section.

In searching neighbour atoms of blocks or neighbour blocks of atoms, and in the way of labelling, we can completely follow the scheme used in matrix multiplications. For each block b in the domain, there are atoms i in the system whose distance from b is less than the cut-off radius of support functions or projector functions. We refer to these atoms as neighbour atoms of b . The atoms i which are neighbours of at least one block in a domain form a set which we call halo atoms of the domain. The set of partitions containing at least one halo atom are referred to as halo partitions, and the set of nodes having at least one halo partition as halo nodes. Further, we can define a covering set made of partitions as the one which includes all neighbour atoms of one or more blocks in a domain. We call this a domain covering set (DCS) of partitions. Following this way of naming, a grand covering set used in matrix multiplications can be referred to as a bundle covering set (BCS) of partitions. Similarly, as we need a list of neighbour blocks of atoms in a bundle, we define the terms, such as neighbour blocks of the atom i , halo blocks of a bundle, and a BCS of blocks. Even in increasing the size of a simulation cell, if we increase the number of processors and keep the form of each domain, the number of members in covering sets is obviously constant. Thus, with the covering sets, the cost in searching neighbour atoms or blocks is proportional to N , not to N^2 . This advantage of the new code is important for the calculations on very large systems.

If we choose the domain as compact as possible, we can optimize the memory requirement for domain-responsible nodes in keeping $\{\phi_{i\alpha}(\mathbf{r})\}$ and $\{\chi_{i\alpha}(\mathbf{r})\}$. Further, if we increase the overlap of a bundle and a domain for each node, the cost of the communication between bundle- and domain-responsible nodes becomes small. It is obvious that we must have good load balancing and that we must also consider the cost of matrix multiplications, which is most expensive at present. Although we have already succeeded in achieving good load balancing for matrix multiplications, to optimize bundles and domains simultaneously is a more complicated problem. We are now working on this problem and will report the results in the near future.

5. Future prospects

In many ways, linear scaling DFT is now established as a viable technique. Within the CONQUEST project, we have shown how practical linear scaling performance can be achieved on systems of many thousands of atoms. Most of the practical problems presented by the search for the self-consistent ground state for such large systems are now solved, or are close to being solved. The practical challenges of implementing the algorithms on large parallel computers, including PC clusters, have also been addressed in detail. However, one issue that is not yet solved to our satisfaction is that of the basis sets for representing support functions, and this is the main reason that CONQUEST has not yet been applied to major scientific problems.

But the ideas embodied in CONQUEST can be seen as part of a larger current of thought that is being followed by many condensed-matter groups—a move away from extended orbitals and extended basis sets and towards a formulation in terms of localized orbitals and localized basis sets. The key issue of how to make effective localized basis sets, which is so crucial to CONQUEST, was explored in depth in the very recent CECAM workshop on ‘Localized orbitals and linear scaling calculations’, which was part-funded by Psi-k. A clear message from this workshop was that the new atomic-like basis functions being developed by several groups can often compete very effectively with plane waves, while demanding far less memory and often far fewer cpu cycles. By contrast, CONQUEST is currently based on finite-element methods that are deliberately related to the plane-wave approach. The exciting recent progress with atomic-like basis sets is already making an important impact

in the SIESTA [57] and PLATO [59] codes and in other codes, and will undoubtedly be important for the future of CONQUEST and for linear scaling DFT in general. We have high hopes that the SIESTA–CONQUEST collaboration, now in its early stages, will accelerate progress in this area.

Finally, we want to emphasize the very broad importance of linear scaling ideas. We have stressed the close relation between linear scaling and the ‘embedding’ problem, and we have presented simple examples that suggest the practical importance of this relation for future work. Another completely different reason for the broad importance of linear scaling is its implications for quantum Monte Carlo simulation. We believe that many of the current linear scaling ideas will be directly transferred to improve the system-size scaling of QMC simulation. By the same token, we expect that the long-standing problem of QMC embedding will also be helped forward by some of the new ideas. The future looks exciting!

Acknowledgments

We are happy to acknowledge useful discussions with D Manolopoulos, A Horsfield and S Goedecker, and assistance with optimization and parallelization from EPCC (Ian Bush) and CSAR (Martyn Foster and Stephen Pickles).

References

- [1] Goedecker S 1999 *Rev. Mod. Phys.* **71** 1085
- [2] Pettifor D G 1989 *Phys. Rev. Lett.* **63** 2480
- [3] Yang W 1991 *Phys. Rev. Lett.* **66** 1438
- [4] Galli G and Parrinello M 1992 *Phys. Rev. Lett.* **69** 3547
- [5] Li X-P, Nunes R W and Vanderbilt D 1993 *Phys. Rev. B* **47** 10 891
- [6] Daw M S 1993 *Phys. Rev. B* **47** 10 895
- [7] Ordejón P, Drabold D, Grumbach M and Martin R 1993 *Phys. Rev. B* **48** 14 646
- [8] Aoki M 1993 *Phys. Rev. Lett.* **71** 3842
- [9] Mauri F, Galli G and Car R 1993 *Phys. Rev. B* **47** 9973
- [10] Goedecker S and Colombo L 1994 *Phys. Rev. Lett.* **73** 122
- [11] Stechel E B, Williams A R and Feibelman P J 1994 *Phys. Rev. B* **49** 10 088
- [12] Nunes R W and Vanderbilt D 1994 *Phys. Rev. B* **50** 17 611
- [13] Hierse W and Stechel E B 1994 *Phys. Rev. B* **50** 17 811
- [14] Kohn W 1995 *Int. J. Quantum Chem.* **56** 229
- [15] Hernández E and Gillan M J 1995 *Phys. Rev. B* **51** 10 157
- [16] Kress J D and Voter A F 1996 *Phys. Rev. B* **53** 12 733
- [17] Horsfield A P 1996 *Mater. Sci. Eng. B* **37** 219
- [18] Horsfield A P, Bratkovsky A M, Fearn M, Pettifor D G and Aoki M 1996 *Phys. Rev. B* **53** 12 964
- [19] Kohn W 1996 *Phys. Rev. Lett.* **76** 3168
- [20] Hernández E, Gillan M J and Goringe C M 1996 *Phys. Rev. B* **53** 7147
- [21] Goringe C M, Hernández E, Gillan M J and Bush I J 1997 *Comput. Phys. Commun.* **102** 1
- [22] Baer R and Head-Gordon M 1997 *Phys. Rev. Lett.* **79** 3962
- [23] Daniels A D, Millam J M and Scuseria G E 1997 *J. Chem. Phys.* **107** 425
- [24] Challacombe M 1999 *J. Chem. Phys.* **110** 2332
- [25] Haynes P D and Payne M C 1999 *Phys. Rev. B* **59** 12 173
- [26] Bowler D R and Gillan M J 1999 *Comput. Phys. Commun.* **120** 95
- [27] Bowler D R, Bush I J and Gillan M J 2000 *Int. J. Quantum Chem.* **77** 831
- [28] Owen J H G, Miki K, Bowler D R, Goringe C M, Goldfarb I and Briggs G A D 1997 *Surf. Sci.* **394** 79
- [29] Goldfarb I, Owen J H G, Hayden P T, Bowler D R, Miki K and Briggs G A D 1997 *Surf. Sci.* **394** 105
- [30] Oviedo J, Bowler D R and Gillan M J 2001 *Phys. Rev. B* submitted
- [31] Rovira C and Parrinello M 2000 *Int. J. Quantum Chem.* **80** 1172
- [32] Bowler D R, Aoki M, Goringe C M, Horsfield A P and Pettifor D G 1997 *Modell. Simul. Mater. Sci. Eng.* **5** 199
- [33] Ismail-Beigi S and Arias T 1999 *Phys. Rev. Lett.* **82** 2127

- [34] Goedecker S 1998 *Phys. Rev. B* **58** 3501
- [35] He L and Vanderbilt D 2001 *Phys. Rev. Lett.* **86** 5341
- [36] Kohn W 1959 *Phys. Rev.* **115** 809
- [37] des Cloiseaux J 1964 *Phys. Rev. A* **135** 658
- [38] Goedecker S and Teter M 1995 *Phys. Rev. B* **51** 9455
- [39] Voter A F, Kress J D and Silver R N 1996 *Phys. Rev. B* **53** 12 733
- [40] Stephan U and Drabold D 1998 *Phys. Rev. B* **57** 6391
- [41] Gillan M J 1989 *J. Phys.: Condens. Matter* **1** 689
- [42] Qiu S-Y, Wang C Z, Ho K M and Chan C T 1994 *J. Phys.: Condens. Matter* **6** 9153
- [43] Goringe C M 1995 *DPhil Thesis* Oxford University
- [44] McWeeny R 1960 *Rev. Mod. Phys.* **32** 335
- [45] Palser A H R and Manolopoulos D E 1998 *Phys. Rev. B* **58** 12 704
- [46] Kim J, Mauri F and Galli G 1995 *Phys. Rev. B* **52** 1640
- [47] Haynes P D and Payne M C 1998 *Solid State Commun.* **108** 737
- [48] White C A, Maslen P, Lee M S and Head-Gordon M 1997 *Chem. Phys. Lett.* **276** 133
- [49] Millam J M and Scuseria G E 1997 *J. Chem. Phys.* **106** 5569
- [50] Bowler D R and Gillan M J 2001 *Chem. Phys. Lett.* submitted
- [51] Hernández E, Gillan M J and Goringe C M 1997 *Phys. Rev. B* **55** 13 485
- [52] Bernholc J, Briggs E L, Sullivan D L, Brabec C J, Nardelli M B, Rapcevicz K, Roland C and Wensell M 1997 *Int. J. Quantum Chem.* **65** 531
- [53] Haynes P D and Payne M C 1997 *Comput. Phys. Commun.* **102** 17
- [54] Gan C K, Haynes P D and Payne M C 2001 *Phys. Rev. B* **63** 205109
- [55] Sánchez-Portal D, Ordejón P, Artacho E and Soler J M 1997 *Int. J. Quantum Chem.* **65** 453
- [56] Sankey O F and Niklewski D J 1989 *Phys. Rev. B* **40** 3979
- [57] Artacho E, Sánchez-Portal D, Ordejón P, García A and Soler J M 1999 *Phys. Status Solidi b* **215** 809
- [58] Bowler D R, Miyazaki T and Gillan M J 2001 *Comput. Phys. Commun.* **321** 1000
- [59] Kenny S D, Horsfield A P and Fujitani H 2000 *Phys. Rev. B* **62** 4899