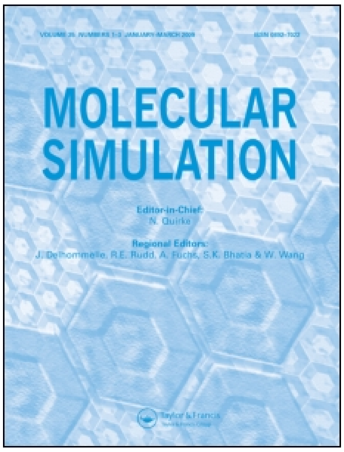

Molecular Simulation
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713644482

Recent Progress in First Principles $O(N)$ Methods

D. R. Bowler[a]; M. J. Gillan[a]
[a] Department of Physics and Astronomy, University College London, London

To cite this Article Bowler, D. R. and Gillan, M. J.(2000) 'Recent Progress in First Principles $O(N)$ Methods', Molecular Simulation, 25: 3, 239 — 255

To link to this Article: DOI: 10.1080/08927020008044128

URL: http://dx.doi.org/10.1080/08927020008044128

Molecular Simulation, 2000, Vol. 25, pp. 239-255
Reprints available directly from the publisher
Photocopying permitted by license only

Published by license under
the Gordon and Breach Science
Publishers imprint.
Printed in Malaysia.

RECENT PROGRESS IN FIRST PRINCIPLES $\mathcal{O}(N)$ METHODS

D. R. BOWLER and M. J. GILLAN*

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

(Received April 1999; accepted May 1999)

We describe recent progress in developing practical first principles methods for which the computer effort is proportional to the number of atoms: linear scaling or $\mathcal{O}(N)$ methods. It is shown that the locality property of the density matrix gives a general framework for constructing such methods. We then outline some of the main technical problems which must be solved in order to develop a practical $\mathcal{O}(N)$ method based on density functional theory and the pseudopotential method. Recent progress in solving these problems is presented, and we show that the spatial cut-off distances needed to achieve good accuracy are small enough to make the calculations feasible. Parallel implementation of the $\mathcal{O}(N)$ methods in the CONQUEST code is outlined, and it is shown that the code exhibits excellent linear-scaling behaviour on test systems of several thousand atoms. It is pointed out that the most important remaining problem concerns the optimal strategy for seeking the ground state. It is argued that there are three different mechanisms of ill-conditioning which cause present search methods to be inefficient, and some partial solutions are suggested.

Keywords: Principles; $\mathcal{O}(N)$ methods; CONQUEST code; functional theory

1. INTRODUCTION

First principles simulation based on density functional theory (DFT) and the pseudopotential method is now in widespread use in many areas of physics and chemistry. However, current calculations are limited to systems of no more than a few hundred atoms, because the computer time needed increases at least as the square of the number of atoms N. There has recently been an intensive effort [1-11] to overcome this limitation by developing techniques for which the computer time is linearly proportional to N.

---

*Corresponding author.

We report here recent progress with these so-called $\mathcal{O}(N)$ or linear scaling techniques.

The reason for the N^2 dependence is easy to understand. The usual DFT techniques work with the eigenfunctions $\psi_i(\mathbf{r})$ of the Kohn-Sham Hamiltonian, which extend over the entire volume of the system. The number of ψ_i functions is proportional to N, and the amount of information in *each* ψ_i is proportional to the volume of the system (which is proportional to N), so that the total number of variables needed to describe the electrons increases as N^2. For very large systems, the dependence becomes N^3, because conventional methods require the calculation of the scalar products $\langle\psi_i|\psi_j\rangle$:

$$\langle\psi_i|\psi_j\rangle = \int d\mathbf{r}\psi_i(\mathbf{r})\psi_j(\mathbf{r})^*, \tag{1}$$

and each of these N^2 quantities needs a computational time proportional to the volume. These bad dependencies on N arise because the ψ_i orbitals extend over the whole system. But it has been recognised for a long time that the quantum state of the electrons in condensed matter can be described in terms of localised functions, and this insight provides the key to developing $\mathcal{O}(N)$ methods.

2. GENERAL FRAMEWORK FOR $\mathcal{O}(N)$

Several authors have stressed that the fundamental reason for the existence of $\mathcal{O}(N)$ techniques can be seen in the properties of the two particle density matrix [10], $\rho(\mathbf{r},\mathbf{r}')$. This can be defined in terms of the Kohn-Sham eigenfunctions as:

$$\rho(\mathbf{r},\mathbf{r}') = \sum_i f_i\psi_i(\mathbf{r})\psi_i(\mathbf{r}')^*, \tag{2}$$

where f_i is the occupation number of orbital i. But instead of considering the ψ_i as the primary quantities in terms of which $\rho(\mathbf{r},\mathbf{r}')$ is defined, we regard $\rho(\mathbf{r},\mathbf{r}')$ itself as the primary quantity [1]. It is known that DFT can be formulated perfectly well in terms of $\rho(\mathbf{r},\mathbf{r}')$ without any explicit mention of wavefunctions. The total energy E_{tot} can be expressed explicitly in terms of ρ, and the ground state is obtained by minimising the functional $E_{\mathrm{tot}}[\rho(\mathbf{r},\mathbf{r}')]$, subject to the conditions: (i) ρ is Hermitian; (ii) ρ is idempotent (*i.e.*, its eigenvalues are all 0 or 1); (iii) ρ yields the correct number of electrons, N_{el}.

$$N_{el} = 2\int d\mathbf{r}\rho(\mathbf{r},\mathbf{r}). \tag{3}$$

The key property of $\rho(\mathbf{r},\mathbf{r}')$ is that it decays to zero as $|\mathbf{r}-\mathbf{r}'|\to\infty$. The fundamental reason for this decay is the loss of quantum phase coherence between distant points. This means that the amount of information in $\rho(\mathbf{r},\mathbf{r}')$ is proportional to N, and so much of the apparent information contained in a wavefunction description must be redundant.

Given this property of $\rho(\mathbf{r}, \mathbf{r}')$, and also remembering that a density matrix must be idempotent (*i.e.*, $\rho\cdot\rho=\rho$), a general $\mathcal{O}(N)$ approach to the determination of the ground state within DFT is to minimise E_{tot} with respect to ρ with the additional constraints:

$$\rho(\mathbf{r},\mathbf{r}')=0, \quad |\mathbf{r}-\mathbf{r}'|>R_c, \tag{4}$$

$$\rho\cdot\rho=\rho, \tag{5}$$

where R_c is some cut-off radius. This will lead to an upper bound to the true ground state energy, which will converge to the true value as R_c is increased. In developing practical methods, one cannot proceed exactly like this, as the ρ depends on two vector positions. Instead, we introduce the additional approximation that ρ is separable, *i.e.*, that it can be expressed in the form:

$$\rho(\mathbf{r},\mathbf{r}')=\sum_{i\alpha,j\beta}\phi_{i\alpha}(\mathbf{r})K_{i\alpha,j\beta}\phi_{j\beta}(\mathbf{r}'); \tag{6}$$

this is equivalent to demanding that ρ have only a finite number of non-zero eigenvalues. The functions $\phi_{i\alpha}(\mathbf{r})$ are known as 'localised orbitals'. The notation indicates that $\phi_{i\alpha}$ is the αth localised orbital on atom i. The matrix $K_{i\alpha,j\beta}$, is the density matrix in the representation of $\phi_{i\alpha}(\mathbf{r})$.

In the procedure we have proposed, the localisation of $\rho(\mathbf{r}, \mathbf{r}')$ expressed in Eq. (4) is replaced by the condition that the localised orbitals be non-zero only within limited regions, which in practice are spheres of radius R_{reg} centred on the atoms. The $\phi_{i\alpha}$ can then be seen as célosely related to the atomic-like orbitals that would be used in a tight binding description. The localisation of $\rho(\mathbf{r},\mathbf{r}')$ also requires that the matrix $K_{i\alpha,j\beta}$ be subject to a spatial cutoff, so that $K_{i\alpha,j\beta}=$ when the separation of the atoms i and j exceeds some value.

The matrix $K_{i\alpha,j\beta}$ cannot be allowed to vary freely, since $\rho(\mathbf{r},\mathbf{r}')$ is required to be idempotent, or at least 'weakly idempotent' (*i.e.*, its eigenvalues must lie in the interval [0,1]). There are several ways of enforcing this condition, one of which is equivalent to the $\mathcal{O}(N)$ DFT scheme proposed by Mauri, Galli and Car [1]. The method we have proposed [10, 12] is based on the 'purification' technique of McWeeny [13], recently used in tight binding calculations by Li, Nunes and Vanderbilt [4]. As explained elsewhere [10, 12], this requires the matrix K to be written as:

$$K = 3LSL - 2LSLSL \tag{7}$$

where $L_{i\alpha,j\beta}$ is an 'auxiliary' density matrix and $S_{i\alpha,j\beta}$ is the overlap matrix:

$$S_{i\alpha,j\beta} = \int d\mathbf{r}\phi_{i\alpha}(\mathbf{r})\phi_{j\beta}(\mathbf{r}). \tag{8}$$

The localisation of $\rho(\mathbf{r},\mathbf{r}')$ is then imposed as a spatial cutoff on $L_{i\alpha,j\beta}$:

$$L_{i\alpha,j\beta} = 0, \; |\mathbf{R}_{i\alpha} - \mathbf{R}_{j\beta}| > R_L, \tag{9}$$

where R_i is the position of atom i and R_L is a cutoff radius.

We summarise the overall scheme: The ground state energy and density matrix of the system are determined by minimising E_{tot} with respect to the $\phi_{i\alpha}(\mathbf{r})$ functions and the auxiliary density matrix $L_{i\alpha,j\beta}$, subject to spatial cutoffs R_{reg} and R_L. This gives an upper bound to the true E_{tot}, which is expected to go to the true value as R_{reg} and R_L are increased.

3. PRACTICAL IMPLEMENTATION

3.1. Conquest

Within CONQUEST, the search for the ground state is organised as a triple loop. At the innermost level, the $L_{i\alpha,j\beta}$ matrix is varied so as to minimise the total energy, while keeping the potential and localised orbitals fixed. Outside that loop, a self-consistent charge density and potential are found (with the concommitant re-minimisation of the energy with respect to $L_{i\alpha,j\beta}$). The outermost loop is the variation of the localised orbitals $\phi_{i\alpha}(\mathbf{r})$ to minimise the energy.

As a result of this structure, CONQUEST can operate in three different modes. If the localised orbitals and Kohn-Sham potentials are kept fixed at all times, then it operates as a non-self-consistent tight binding code (though with the form of the localised orbitals determining the Hamiltonian). If the potential is allowed to vary to achieve self-consistency, then an *ab initio* tight binding code results. Finally, by varying the localised orbitals, a full *ab initio* DFT code is achieved. One of the strengths of the method is that these different regimes can be combined in the same sample seamlessly – so, for instance, a full *ab initio* calculation focussing on reacting species could be embedded in a much larger (and computationally less demanding) *ab initio* tight binding calculation.

We have noted that the inner loop is identical to the ground state search in a tight binding calculation. It has been stressed recently [15] that in this context invariance with respect to linear transformations of the $\phi_{i\alpha}(\mathbf{r})$ imposes certain natural constraints, which should be respected. The implication of this is that in a non-orthogonal basis set, the correct metric must be applied to ensure correct gradients (*i.e.*, a fully contravariant gradient $\partial E_{\mathrm{tot}}/\partial L_{i\alpha,j\beta}$ must be added to the contravariant density matrix; this involves applying the metric S^{-1} to the gradient calculated by Nunes and Vanderbilt [16]). Our practical experience with the CONQUEST code is that this procedure is markedly more efficient that a naive search along the gradients of $\partial E_{\mathrm{tot}}/\partial L_{i\alpha,j\beta}$.

Although we find that these search methods generally work, we are not satisfied with their efficiency, and we return to this question in Section 4.

3.2. Dependence on the Spatial Cutoffs

We have already reported preliminary results on the dependence of the calculated ground state energy on the cutoffs R_{reg} and R_L, which suggested that accurate results are obtained with quite modest cutoffs. But these were mainly based on a model local pseudopotential and so were not fully realistic. We report here new tests using standard non-local pseudopotentials [17, 18]. The tests are done on perfect crystals of silicon.

In testing the dependence of E_{tot} on R_{reg}, we set R_L equal to infinity, which is equivalent to exact diagonalisation. For comparison, we have also done calculations with the standard plane wave code CASTEP [19] using precisely the same pseudopotential and other parameters. Figure 1 shows the calculated total energy as a function of R_{reg} for Si and Ge. The results show that E_{tot} converges to the correct value extremely rapidly once R_{reg} is greater than 4 Å. For this radius, each localisation region contains 17 neighbouring atoms, and the calculations are perfectly manageable.

Our tests on R_L were done with $R_{\mathrm{reg}} = 2.715$ Å, and the results for Si are shown in Figure 2. Rather accurate convergence to the $R_L = \infty$ value is obtained for $R_L \geq 8$ Å, which again is acceptable. No value is shown for exact diagonalisation because of technical difficulties in performing comparisons.

The conclusion from these tests is that the practical values of the spatial cutoffs needed to achieve good accuracy are encouragingly small. Above, we asked for what system size $\mathcal{O}(N)$ becomes more efficient than conventional

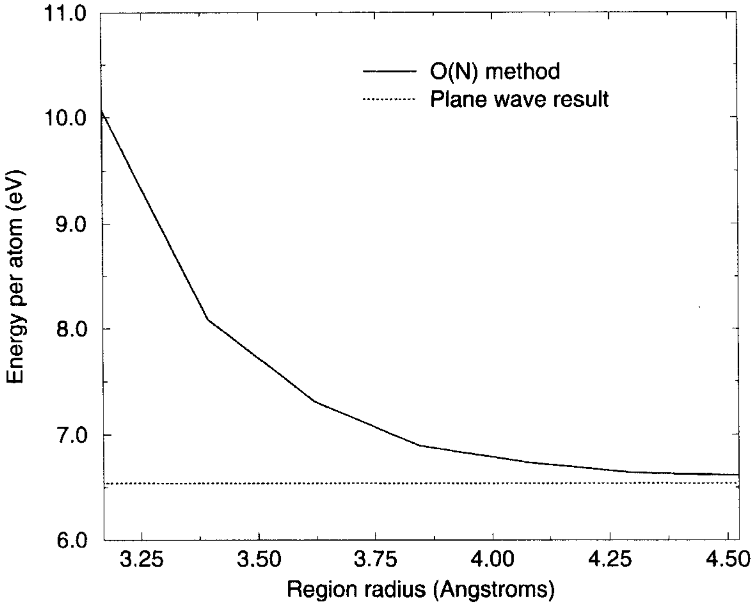

FIGURE 1 The total energy per atom versus region radius (R_{reg}) for silicon as calculated by CONQUEST. The value from the plane wave code CASTEP is shown for comparison.

methods. We believe that it is too soon to answer this question. The main reason is that the efficiency of the ground state search technique in $\mathcal{O}(N)$ is still rather poor, as discussed in Section 4.

3.3. Representation of the Localised Orbitals

In considering this question, it is helpful to remember the lessons that have been learnt from the use of plane wave basis sets in conventional pseudopotential calculations. Two of the major advantages of plane waves are that: (i) Systematic convergence of the total energy with respect to basis set completeness is achieved by increasing a single parameter-the plane wave cutoff energy, E_{cut}; (ii) they are free of bias, *i.e.*, they are completely flexible, and no judgement has to be made about the kind of chemical bonding that will occur. If possible, the basis set used in $\mathcal{O}(N)$ calculations should retain these advantages. In any case, we must certainly aim to achieve the accuracy normally expected of current plane wave calculations.

One possibility is to use the spherical analogue of plane waves, namely the product of spherical Bessel functions $j_l(r)$ and spherical harmonics

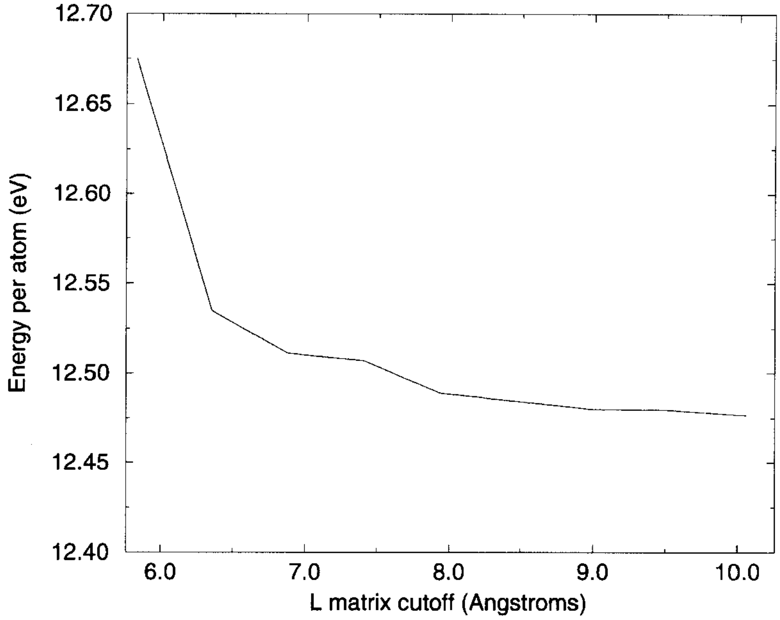

FIGURE 2 The total energy per atom versus density matrix cutoff (R_L) as calculated by CONQUEST for silicon. The region radius was 2.715 Å.

$Y_l^m(\hat{\mathbf{r}})$, in each localisation region. The advantages of this representation have been discussed by Haynes and Payne [20]; very preliminary results have recently been reported [21]. An alternative is simply to represent the $\phi_{i\alpha}(\mathbf{r})$ by their values on a grid and to calculate matrix elements of the kinetic energy by finite differences. This method is well established to be practicable in non- $\mathcal{O}(N)$ calculations [22] and has been studied in the $\mathcal{O}(N)$ context by us [23] and very recently by Hoshi and Fujiwara [24]. A third method is the B-spline basis set used in the CONQUEST code. In this scheme, the basis functions $\chi_s(\mathbf{r})$ in the expansion:

$$\phi_{i\alpha}(\mathbf{r}) = \sum_s b_{i\alpha s}\chi_s(\mathbf{r}) \tag{10}$$

are piecewise polynomial functions strictly localised on the points of a grid which is rigidly attached to each atom. Details of the B-spline scheme are reported elsewhere [10].

In the near future, we will move to a scheme where the localised orbitals are represented by a combination of B-splines and pseudo-atomic orbitals (PAOs) (such as those used in a well-known *ab initio* tight binding scheme

[25]). The advantage of PAOs is that they are well adapted; however a *systematic* improvement of basis set is hard. By combining B-splines and PAOs we expect to use the strengths of both methods while minimising the problems. This will be reported elsewhere.

3.4. Parallel Implementation

The essence of $\mathcal{O}(N)$ is that the system can be separated into independent spatial regions. This means that $\mathcal{O}(N)$ is ideally suited to parallel implementation, with each processor being responsible for a set of atoms and/or spatial regions. The way this is done in the parallel CONQUEST code is described in detail elsewhere [14], so we give here only a brief summary.

Each processor is given a three-fold responsibility. First, it is in charge of a certain group of atoms. This means that it holds the basis coefficients $b_{i\alpha s}$, the derivatives $\partial E_{\mathrm{tot}}/\partial b_{i\alpha s}$, and also the rows of all matrix elements such as $S_{i\alpha,j\beta}$ corresponding to these atoms. It is responsible for performing the transforms from basis coefficients to grid values $\phi_{i\alpha}(\mathbf{r}_l)$ for these atoms, and also for doing matrix multiplications needed to produce matrix rows associated with its atoms. Second, each processor is in charge of a domain of integration grid points $\mathbf{r}_l$, and has the job of calculating contributions to matrix elements coming from sums over this domain. It also has responsibility for the electron density and Kohn-Sham potential on its domain of points. Third, the processor is responsible for doing part of the spatial Fourier transforms needed in calculating the Hartree potential. In practice, this means that it deals with a set of columns of grid points in the x, y or z directions. The processors switch between their responsibilities in a synchronised manner, and communication of data between them is needed when this happens.

We have made extensive tests of the scaling properties of the CONQUEST code. It is important to stress that there are two completely different types of scaling. The first concerns the way the CPU time increases as the size of the simulated system increases for a *fixed* number of processors; we call this 'intrinsic scaling'. The second concerns the way the CPU time changes when a *given* simulated system is treated on varying numbers of processors.

In an implementation of CONQUEST on the Cray T3D, both types of scaling turn out to be excellent. As an illustration of the intrinsic scaling, we show in Figure 3 the total CPU time required per iteration for silicon crystals containing from 64 to 6144 atoms. (Here, total CPU time means the number of processors multiplied by the CPU time per processor). The

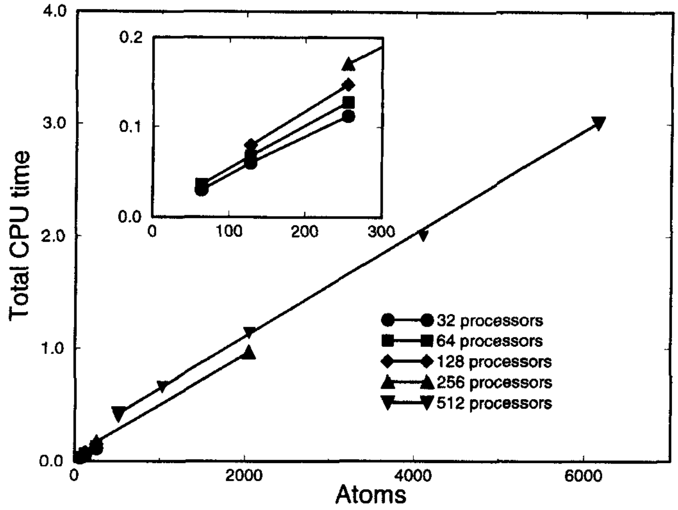

FIGURE 3 The total CPU time taken for silicon systems of differing size, using the CONQUEST code on the Cray T3D.

results show that for a fixed number of processors, the total CPU time is almost exactly proportional to the number of atoms, and this confirms the correctness of the underlying theory. We also find that for a given simulated system, the total CPU time increases only very weakly with the number of processors, which means that the fraction of time taken in communications is very small.

3.5. Forces on Atoms

In the conventional plane wave technique, the force on each atom is simply the Hellmann-Feynman force, *i.e.*, the force exerted by the electrons in the ground state associated with the current ionic positions (plus, of course, the Coulombic interaction between ionic cores). This relies on the fact that the basis set does not depend on the ionic positions. In the $\mathcal{O}(N)$ technique used in CONQUEST, the B-spline basis functions $\chi_s(\mathbf{r})$ move with the ions, and this gives rise to an additional contribution to the force, known as the Pulay contribution [26]. If the calculation is well converged with respect to basis set completeness, then the Pulay correction is small, but it is nonetheless essential to include it, in order to ensure exact consistency between the total energy and the forces. As will be described in more detail elsewhere, the Pulay contribution is straightforward to calculate. This means that the relaxation of the system to mechanical equilibrium and the generation of time-dependent ionic trajectories will be feasible in $\mathcal{O}(N)$ DFT calculations.

3.6. Achieving idempotency for *L*

As has already been emphasised, the density matrix of the true electronic ground state is idempotent. This important property is hard to impose during a minimisation; however, a number of schemes which drive the matrix towards idempotency have been suggested. For simplicity, we will consider these schemes in the framework of orthogonal tight binding theory; the extension to the non-orthogonal case and DFT is simple enough. The first scheme was proposed by McWeeny [13], who noted that if a matrix ρ is close to idempotency, then the matrix $\tilde{\rho}$ given by:

$$\tilde{\rho} = 3\rho^2 - 2\rho^3 \tag{11}$$

will be more nearly idempotent. It has the effect of driving the eigenvalues of ρ towards zero and one (this can be seen by considering the function $3\lambda^2 - 2\lambda^3$, which is shown in Fig. 4). If this transformation (often called

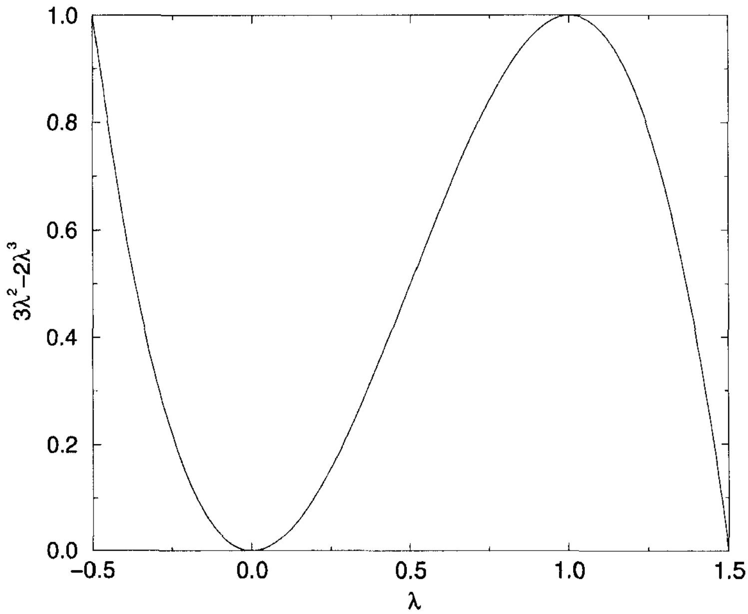

FIGURE 4 The McWeeny purification function, $f(\lambda) = 3\lambda^2 - 2\lambda^3$.

the McWeeny transformation or purification transformation) is applied iteratively (writing $\rho_{n+1} = 3\rho_n^2 - 2\rho_n^3$ for iteration $n+1$), then the sequence of matrices generated will converge on an idempotent matrix. In fact, this transformation is *quadratically* convergent (*i.e.*, if the idempotency error in ρ is $\delta\rho$, then the idempotency error in $\tilde{\rho}$ is $\delta\rho^2$).

Palser and Manolopoulos [27] have recently suggested using this iterative scheme in an $\mathcal{O}(N)$ manner. They point out that if the initial density matrix is a linear function of the Hamiltonian, with eigenvalues between zero and one, then the iteration will converge to the correct ground state density matrix (given by $\theta(\mu - H)$, where $\theta(x)$ is the Heaviside step function ($\theta = 1$ for $x > 0$ and $\theta = 0$ for $x < 0$) and μ is the chemical potential for electrons, or the Fermi energy), and that the energy, $E = 2\mathrm{Tr}[\rho_n H]$, will decrease monotonically at each step. This procedure has the advantage of being fast (it only requires two matrix multiplies per iteration) and efficient (it converges quadratically). Unfortunately, when a localisation criterion (also called a truncation) is applied to the density matrix to achieve linear scaling, the monotonic decrease of energy will fail at some point in the iterative search. This can be taken as an indication that truncation errors are dominating the calculation, and that the search should be stopped [27]; indeed, if it is not stopped, there is no guarantee that it will continue to converge towards an idempotent matrix. This is a heuristic criterion for stopping the iteration, and has the drawback that the method will not be variational, so that analytic forces will not be in agreement with the numerical gradient of the energy.

The Li, Nunes and Vanderbilt (LNV) scheme for achieving the ground state density matrix [4,16] also uses the McWeeny transformation, though in a rather different manner. Here, the energy is written as $E = 2\mathrm{Tr}[\tilde{\rho}H]$, with $\tilde{\rho}$ given by Eq. (11). Then the energy is minimised with respect to the elements of ρ, typically using a scheme such as conjugate gradients [28] to generate a sequence of search directions. The localisation of the density matrix is achieved by applying a spatial cutoff to the elements of ρ. This scheme has at least two advantages: first, each line minimisation can be performed analytically, as the energy is cubic in ρ; second, it is variational, so that the energy found is always an upper bound to the ground state, and forces obey the Hellmann-Feynman theorem and are in exact agreement with the numerical derivative of the energy.

However, there are drawbacks to the LNV technique. First, it is unclear what the best initial value should be for the density matrix; typically, it is taken to be $(1/2)\mathbf{I}$, or $(1/2)\mathbf{S}^{-1}$ in a non-orthogonal basis set. Second, as the McWeeny transformation is a cubic, it is unbounded from below, and a poor starting choice for the minimisation can lead to runaway solutions; a sign of this is typically that the cubic has complex extrema. Third, the scheme can be poorly convergent, and is not guaranteed the quadratic convergence of the McWeeny method.

We have recently proposed a hybrid between the McWeeny and LNV schemes which builds on the complementarity between these two methods [12]. It is based on the observation that the sequence of matrices generated during a McWeeny iterative search converges quadratically towards idempotency, and that the LNV search direction maintains idempotency in the density matrix to first order. Thus the McWeeny scheme is used as an initialisation to find an idempotent density matrix (but one which is not the ground state matrix because of truncation errors); this matrix is then used as the input to the LNV scheme, which maintains the idempotency while searching efficiently for the ground state density matrix. The combination of the two methods is both variational and robust – two highly desirable attributes for the inner loop of an $\mathcal{O}(N)$ DFT method.

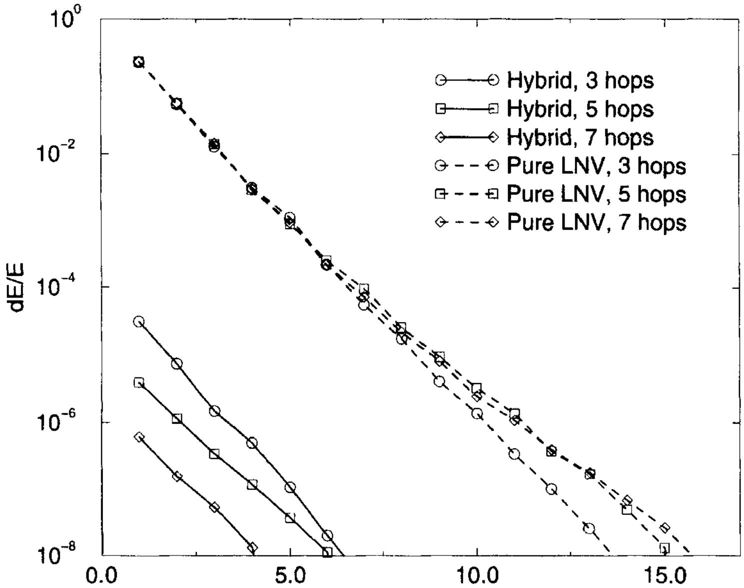

FIGURE 5 The difference between the cohesive energy at a given iteration and the final cohesive energy for the LNV stage of the hybrid scheme (solid lines) and the pure LNV scheme (dashed lines) for diamond structure carbon. Results are shown for different cut-off radii: 3 hops (circles), 5 hops (squares) and 7 hops (diamonds), where a hop is a Hamiltonian nearest neighbour distance.

As an example of the improved speed of convergence given by the hybrid scheme, Figure 5 shows the convergence to the ground state energy in diamond carbon for the LNV stage of the hybrid scheme and pure LNV (initialised from $\rho = (1/2)\mathbf{I}$). These results show that the McWeeny stage of the hybrid scheme gets closer to the ground state as the radius is increased, as expected, and that it acts as an excellent initial density matrix for the LNV scheme. The method has also been tested on a vacancy in diamond C, the Si(001) surface and liquid Si, as reported elsewhere [12].1

4. ILL-CONDITIONING PROBLEMS

In spite of the very encouraging findings summarised above, there remains a question that is still not completely solved: the strategy for finding the ground state. The search procedure outlined in Section 3.1 is generally successful, but it sometimes requires many iterations, particularly when the region radius R_{reg} is large. This means that practical calculations are rather inefficient. The problem concerns variations of the $\phi_{i\alpha}$, since convergence with respect to the $L_{i\alpha,j\beta}$ matrix in the inner loop is generally rapid. The large number of iterations is needed in the outer loop where the $\phi_{i\alpha}$ are varied.

The cause of the problem is ill-conditioning. This is an extremely generally phenomenon which afflicts minimisation problems in many areas of science, and occurs when the function being minimised has a wide range of curvatures. Suppose we need to locate the minimum of some general function $f(x_1,x_2,\ldots,x_N)$ which depends on the set of variables $\{x_i\}$. Let $C_{ij} \equiv \partial^2 f / \partial x_i \partial x_j$ be the curvature matrix (sometimes called the Hessian) evaluated at the minimum. If the eigenvalues λ_n of C_{ij} span a wide range – the ratio between the largest and smallest eigenvalues $\lambda_{\max}/\lambda_{\min}$ is large – then the surfaces of constant f are very elongated (see Fig. 6), and conventional techniques such as conjugate gradients become very inefficient. In fact it is known that the number of iterations needed by conjugate gradients is proportional to $(\lambda_{\max}/\lambda_{\min})^{1/2}$.

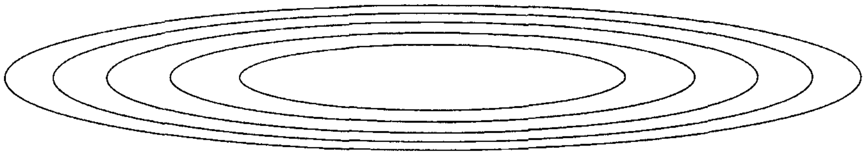

FIGURE 6 A function with elongated surfaces of constant f.

Although ill-conditioning is a very widespread phenomenon, its causes are specific to the problem at hand. In order to overcome the problem, it is essential to understand these causes. In fact, it has been recognised for many years that conventional first principles calculations can suffer from ill-conditioning, and its causes are already well understood. It turns out that the ill-conditioning we encounter in $\mathcal{O}(N)$ calculations is closely related to that found in conventional calculations, so it will be useful to spend a few moments recalling some well known facts.

In the usual plane wave techniques, the total energy E_{tot} has to be minimised with respect to the Kohn-Sham orbitals ψ_i, which are represented in a plane wave expansion:

$$\psi_i = \sum_{\mathbf{G}} c_{i\mathbf{G}} \exp(i\mathbf{G} \cdot \mathbf{r}). \tag{12}$$

For many systems, particularly metals, it is common to allow partial occupation of the orbitals, so that orbital ψ_i has occupation number f_i. We then have $E_{\mathrm{tot}} = E_{\mathrm{tot}}(\{c_{i\mathbf{G}}\},\{f_i\})$. The function E_{tot} has high curvatures associated with variation of $c_{i\mathbf{G}}$ at high wavevector $\mathbf{G}$. The reason for this is simply that the kinetic energy E_{kin} of the electrons is

$$E_{\mathrm{kin}} = 2 \sum_i f_i \sum_{\mathbf{G}} \frac{\hbar^2 G^2}{2m} |c_{i\mathbf{G}}|^2, \tag{13}$$

so that the curvature is proportional to G^2. Since this kind of ill-conditioning comes from the variation of curvature with length scale, we refer to this as 'length scale ill-conditioning'. The problem can be cured by conventional preconditioning methods, and an explicit, efficient scheme has been implemented and tested within CONQUEST [29] with great success.

Conventional techniques can also suffer from a second type of ill-conditioning, associated with the invariance of E_{tot} under unitary transformations of the orbitals. If all f_i are zero or unity, the E_{tot} is exactly invariant under transformations:

$$\psi_i \rightarrow \psi_i' = \sum_j U_{ij} \psi_j \tag{14}$$

where U_{ij} is unitary. This invariance implies that some of the eigenvalues of the Hessian vanish. But if the f_i deviate from zero or unity, the exact invariance is broken, and the vanishing eigenvalues of the Hessian acquire small positive values. It is their smallness that causes the ill-conditioning.

We refer to this mechanism as 'superposition ill-conditioning'. In conventional techniques, this is usually cured by a method known as sub-space rotation.

There is yet another cause of ill-conditioning. When variable occupation numbers are employed, orbitals whose energies are well above the Fermi energy will have very small values of f_i. Variations of the corresponding ψ_i will therefore have little effect on E_{tot}, so that the curvatures will again be small. Since orbitals having small f_i are almost redundant, we call this mechanism 'redundancy ill-conditioning'.

All these three types of ill-conditioning can also cause trouble in $\mathcal{O}(N)$ techniques. It is clear, for example, that the localised orbitals $\phi_{i\alpha}$ can vary on different length scales, so that length scale ill-conditioning is inevitable. This will not cause serious problems, and will be overcome by conventional preconditioning techniques.

Superposition ill-conditioning, associated with linear mixing of the localised orbitals, is more interesting. Two kinds should be distinguished. The first consists of mixing of different $\phi_{i\alpha}$ on the same atom. It is readily shown that this leaves E_{tot} exactly invariant, and cannot cause trouble. The second consists of mixing of $\phi_{i\alpha}$ on different atoms. Since the $\phi_{i\alpha}$ are constrained to be zero outside their regions, this second kind is, strictly, impossible. However, for large region radii there are variations which respect this constraint, while consisting almost exactly of linear superpositions of $\phi_{i\alpha}$ on different atoms. The small curvatures of E_{tot} arising from these variations give rise to superposition ill-conditioning. Our present belief is that this problem will not be difficult to overcome. The reason is that the characteristic variations responsible for the ill-conditioning can be calculated, and this will make it possible to precondition them.

Finally, we comment on redundancy ill-conditioning. We have noted that in conventional techniques this occurs when the number of orbitals exceeds the sum of the occupation numbers (*i.e.*, half the electron number in spin-paired calculations). An analogous problem will afflict $\mathcal{O}(N)$ when the number of $\phi_{i\alpha}$ is greater than half the electron number. This will not always happen, because for many systems the number of $\phi_{i\alpha}$ can be taken equal to half the electron number. But for other systems it will be essential, or at least desirable, to work with a larger number of $\phi_{i\alpha}$. Systems consisting of group IV elements are a case in point, because it will generally be natural to take four $\phi_{i\alpha}$ on each atom, one corresponding to the valence s-orbital and the other three to p-orbitals. Once again, we believe that preconditioning will allow us to overcome this problem, but detailed techniques have yet to be formulated.

5. PROSPECTS

The developments presented here give reason for great optimism about the future potential of $\mathcal{O}(N)$ DFT techniques. We have shown how the properties of the density matrix allow one to give a very general framework for constructing such techniques. The detailed methods we have implemented in the CONQUEST code represent only one possible way of doing this, and other ways will need to be investigated. We have pointed to a number of technical problems that must be overcome in constructing practical $\mathcal{O}(N)$ techniques, and we have shown that solutions to most of these problems are now available. However, some of these may only be interim solutions. We believe, for example, that the question of how best to represent the localised orbitals will need considerable further investigation before any consensus will be reached. The same goes for parallel implementation. We have outlined one way of doing this, and have shown that this works well for systems of several thousand atoms. However, we have done this on only one kind of machine (the Cray T3D), and it may well be that implementation on other machines (*e.g.*, vector parallel machines) will raise new questions. Finally, we have pointed out that there are unsolved questions about the right way to search for the ground state. The ill-conditioning problems that we have described will need deeper study.

Perhaps the most important conclusion is that $\mathcal{O}(N)$ DFT calculations definitely *work*. The spatial cutoff distances required are small enough to make the calculations perfectly feasible. Moreover, the $\mathcal{O}(N)$ behaviour of the calculations is actually demonstrated in practice. Encouraged by these results, our group is now working towards the application of these techniques to complex large-scale problems, including nanostructures on semiconductor surfaces.

References

[1] Mauri, F., Galli, G. and Car, R. (1993). *Phys. Rev. B*, **47**, 9973.
[2] Mauri, F. and Galli, G. (1994). *Phys. Rev. B*, **50**, 4316.
[3] Kim, J., Mauri, F. and Galli, G. (1995). *Phys. Rev. B*, **52**, 1640.
[4] Li, X.-P., Nunes, R. W. and Vanderbilt, D. (1994). *Phys. Rev. B*, **47**, 10891.
[5] Kohn, W. (1995). *Int. J. Quant. Chem.*, **56**, 229.
[6] Kohn, W. (1996). *Phys. Rev. Lett.*, **76**, 3168.
[7] Ordejón, P., Drabold, D. A., Martin, R. M. and Grumbach, M. P. (1995). *Phys. Rev. B*, **51**, 1456.
[8] Ordejón, P., Artacho, E. and Soler, J. M. (1996). *Phys. Rev. B*, **53**, 10441.
[9] Goedecker, S. and Colombo, L. (1994). *Phys. Rev. Lett.*, **73**, 122.
[10] Hernández, E., Gillan, M. J. and Goringe, C. M. (1996). *Phys. Rev. B*, **53**, 7147.

[11] Bowler, D. R., Aoki, M., Goringe, C. M., Horsfield, A. P. and Pettifor, D. G. (1997). Modell. Simul., In: *Mater. Sci. Eng.*, **5**, 199.
[12] Bowler, D. R. and Gillan, M. J. (1999). *Comput. Phys. Commun.*, **120**, 95.
[13] McWeeny, R. (1960). *Rev. Mod. Phys.*, **32**, 335.
[14] Goringe, C. M., Hernández, E., Gillan, M. J. and Bush, I. J. (1997). *Comp. Phys. Commun.*, **102**, 1.
[15] White, C. A., Maslen, P., Lee, M. S. and Head-Gordon, M. (1997). *Chem. Phys. Lett.*, **276**, 133.
[16] Nunes, R. W. and Vanderbilt, D. (1994). *Phys. Rev. B*, **50**, 17611.
[17] Kerker, G. (1980). *J. Phys. C*, **13**, 189.
[18] Kleinman, L. and Bylander, D. M. (1982). *Phys. Rev. Lett.*, **4**, 1425.
[19] Payne, M. C., Teter, M. P., Allan, D. C., Arias, T. A. and Joannopoulos, J. D. (1992). *Rev. Mod. Phys.*, **64**, 1045.
[20] Haynes, P. D. and Payne, M. C. (1997). *Comp. Phys. Commun.*, **102**, 17.
[21] Haynes, P. D. and Payne, M. C. (1998). *Solid State Commun.*, **108**, 737–741 and (1999). *Phys. Rev. B*, **59**, 12173.
[22] Chelikowsky, J. R., Troullier, N. and Saad, Y. (1994). *Phys. Rev. Lett.*, **72**, 1240; Briggs, E. L., Sullivan, D. J. and Bernholc, J. (1996). *Phys. Rev. B*, **54**, 14326.
[23] Hernández, E. and Gillan, M. J. (1995). *Phys. Rev. B*, **51**, 10157.
[24] Hoshi, T., Arai, M. and Fujiwara, T. (1995). *Phys. Rev. B*, **52**, R5459; Hoshi, T. and Fujiwara, T. (1997). unpublished.
[25] Sankey, O. F. and Niklewski, D. J. (1989). *Phys. Rev. B*, **40**, 3979.
[26] Pulay, P. (1969). *Mol. Phys.*, **17**, 197; Scheffler, M., Vigneron, J. P. and Bachelet, G. B. (1985). *Phys. Rev. B*, **31**, 6541.
[27] Palser, A. R. H. and Manolopoulos, D. M. (1998). *Phys. Rev. B*.
[28] Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. (1992). *Numerical Recipes* (Cambridge University Press, Cambridge).
[29] Bowler, D. R. and Gillan, M. J. (1998). *Comput. Phys. Commun.*, **112**, 103.