$\mathcal{O}(N)$ methods in electronic structure calculations

2012 Rep. Prog. Phys. 75 036503

(http://iopscience.iop.org/0034-4885/75/3/036503)

Rep. Prog. Phys. **75** (2012) 036503 (43pp) doi:10.1088/0034-4885/75/3/036503

$\mathcal{O}(N)$ methods in electronic structure calculations

D R Bowler[1,2,3] **and T Miyazaki**[4]

[1] London Centre for Nanotechnology, UCL, 17-19 Gordon St, London WC1H 0AH, UK
[2] Department of Physics & Astronomy, UCL, Gower St, London WC1E 6BT, UK
[3] Thomas Young Centre, UCL, Gower St, London WC1E 6BT, UK
[4] National Institute for Materials Science, 1-2-1 Sengen, Tsukuba, Ibaraki 305-0047, Japan

E-mail: david.bowler@ucl.ac.uk and MIYAZAKI.Tsuyoshi@nims.go.jp

Received 8 September 2011, in final form 3 November 2011
Published 15 February 2012
Online at stacks.iop.org/RoPP/75/036503

Abstract
Linear-scaling methods, or $\mathcal{O}(N)$ methods, have computational and memory requirements which scale linearly with the number of atoms in the system, N, in contrast to standard approaches which scale with the cube of the number of atoms. These methods, which rely on the short-ranged nature of electronic structure, will allow accurate, *ab initio* simulations of systems of unprecedented size. The theory behind the locality of electronic structure is described and related to physical properties of systems to be modelled, along with a survey of recent developments in real-space methods which are important for efficient use of high-performance computers. The linear-scaling methods proposed to date can be divided into seven different areas, and the applicability, efficiency and advantages of the methods proposed in these areas are then discussed. The applications of linear-scaling methods, as well as the implementations available as computer programs, are considered. Finally, the prospects for and the challenges facing linear-scaling methods are discussed.

(Some figures may appear in colour only in the online journal)

This article was invited by M Finnis.

Contents

1. Introduction 1
2. Real-space methods 3
2.1. Finite differences 4
2.2. Finite elements and local real-space bases 4
2.3. Atomic-like orbitals 5
2.4. Representing localized orbitals 7
2.5. Hamiltonian building 7
3. Linear-scaling methods 8
3.1. Density matrix properties and Wannier functions 8
3.2. Solving for the density matrix 11
4. Technical details and parallelization 23
4.1. Non-orthogonal basis functions 23
4.2. Preserving electron number 25
4.3. Parallelization and sparse matrices 25
4.4. Structure relaxation 29
5. Implementation and applications 29
5.1. Implementations 29
5.2. Applications 30
6. Conclusions 36
Acknowledgments 38
References 38

1. Introduction

Electronic structure calculation methods based on the density functional theory (DFT) have been playing important roles in condensed matter physics for more than 40 years. In the early stages, DFT calculations were employed mainly for the study of the electronic structure of simple solids, using a few atoms in a unit cell, with the use of periodic boundary conditions. Since then, there has been a huge effort to improve the accuracy and efficiency of the calculation techniques. In terms of efficiency, after the pioneering work by Car and Parrinello [1], the size of the target systems has increased dramatically, and

more and more examples of the DFT studies, especially on aperiodic systems like surface structures, have emerged. DFT calculations on systems containing hundreds of atoms are currently ubiquitous. As the system size for DFT calculations has become larger, the variety of materials and phenomena investigated by the method has increased. The information of the total energy and atomic forces calculated by DFT methods can provide reliable data independently from experiments, and the methods are nowadays considered as one of the established research tools in fields such as physics, chemistry, materials science and many others. Recently, there have been DFT studies in the complex fields of nano-structured materials and biological systems. In the study of these classes of materials, we need to treat systems containing at least thousands of atoms. However, as is well known, once the number of atoms N in a system reaches around one thousand, the cost of standard DFT calculations increases very rapidly as a cube of N. To overcome this problem, methods known as linear-scaling or $\mathcal{O}(N)$ DFT methods have been developed [2]. The progress of these methods in the last ten to fifteen years is remarkable, and the purpose of this review paper is to overview the recent progress of $\mathcal{O}(N)$ DFT methods.

We will start with an overview of the conventional DFT method and its advantages. In the normal DFT approach, we solve for the Kohn–Sham (KS) orbitals $\Psi_{\nu k}(\boldsymbol{r})$, which are the eigenstates of the KS equation [3]:

$$\hat{H}^{\mathrm{KS}}\Psi_{\nu\boldsymbol{k}}(\boldsymbol{r}) = \left[-\frac{\hbar}{2m}\nabla^2 + V_{\mathrm{ext}}(\boldsymbol{r}) + V_{\mathrm{H}}(\boldsymbol{r}) + V_{\mathrm{XC}}(\boldsymbol{r})\right]\Psi_{\nu\boldsymbol{k}}(\boldsymbol{r}) = \epsilon_{\nu\boldsymbol{k}}\Psi_{\nu\boldsymbol{k}}(\boldsymbol{r}). \quad (1)$$

Here $\hat{H}^{\mathrm{KS}}$ is the KS Hamiltonian, and ν and $\boldsymbol{k}$ are the band index and $\boldsymbol{k}$ points in the first Brillouin zone. Hereafter, we omit $\boldsymbol{k}$ for clarity because we consider large systems and the number of $\boldsymbol{k}$ points is small. $V_{\mathrm{ext}}(\boldsymbol{r})$ is the potential from nuclei, $V_{\mathrm{H}}(\boldsymbol{r})$ is the Hartree potential, and $V_{\mathrm{XC}}(\boldsymbol{r})$ is the exchange-correlation potential in the KS formalism. The most accurate DFT calculations often use a plane-wave basis set to express the KS orbitals:

$$\Psi_\nu(\boldsymbol{r}) = \sum_{|\boldsymbol{G}|<G_{\max}} c_\nu(\boldsymbol{G})\exp(\mathrm{i}\boldsymbol{G}\cdot\boldsymbol{r}). \quad (2)$$

A plane-wave basis set has two main advantages. First, the accuracy of the basis set can be systematically improved. In equation (2), $G_{\max}$ is obtained from the cutoff energy E_{cut} as $\frac{\hbar^2 G_{\max}^2}{2m} = E_{\mathrm{cut}}$. The number of plane-waves, N_G, is controlled only by the number E_{cut}. The accuracy of the basis set can be improved simply by increasing E_{cut}, and a variational principle with respect to E_{cut} is satisfied. The other advantage is that forces can be calculated easily without the Pulay correction term because the basis set is independent of atomic positions (though such basis-set-dependent corrections become necessary when changing the unit cell size or shape). These two advantages make it possible to calculate both energy and forces accurately with plane-wave basis sets.

In order to realize accurate plane-wave calculations, we need to introduce several theoretical techniques. First of all, plane-wave calculations rely on the idea of pseudopotentials [4]. With this method, it is possible to work only with valence electrons and their pseudo-wavefunctions, which are much smoother than the real wavefunctions which oscillate strongly in the core region, and to replace the nuclear potential and the core electrons with a pseudopotential. There have been several kinds of techniques proposed to make pseudo-wavefunctions smoother [5–7]. Using the method of ultra-soft pseudopotentials [8], even the cutoff energy for the localized 3d orbitals of transition metals can be reduced dramatically. With these improvements in theoretical techniques, the total energy converges quickly with respect to the cutoff energy, and this is essential to make accurate DFT calculations feasible. In addition, the major part of the error in the total energy usually comes from the expression of KS orbitals in the core region. Hence, the relative energetic stability of two states (e.g. two different atomic structures) can be reproduced without the absolute convergence because most of the errors are cancelled in the energy difference. Note that it is also possible to reduce the number of plane-waves by using augmentation for wavefunctions in the core region as in the linearized augmented plane-wave (LAPW) or the projector augmented wave (PAW) method [9].

It is essential that we reduce the number of plane-waves by introducing the pseudopotential or other similar techniques. However, even with very smooth pseudo-wavefunctions, N_G is typically one hundred times larger than the number of electrons. When we want to diagonalize $\langle\boldsymbol{G}|H^{\mathrm{KS}}|\boldsymbol{G}'\rangle$, the required memory scales as $\mathcal{O}(N_G^2)$ and CPU time as $\mathcal{O}(N_G^3)$. Hence it is impossible to employ direct (exact) diagonalization except for very small systems. Instead of using exact diagonalization, we can obtain the KS orbitals by minimizing the DFT total energy with respect to the coefficients $\{c_\nu(\boldsymbol{G})\}$, as shown in the work by Car and Parrinello [1]. Since we only need the occupied KS orbitals in such iterative methods, the memory requirement to store $\{c_\nu(\boldsymbol{G})\}$ is proportional to $N_B N_G$, which is roughly 100 times smaller than N_G^2. Then, we update the coefficients $\{c_\nu(\boldsymbol{G})\}$ by calculating the gradient of the total energy with a constraint to keep the KS orbitals orthogonal to each other. This is done by calculating $(H^{\mathrm{KS}} - \Lambda_{\nu,\nu'})$ or $H^{\mathrm{KS}} - \epsilon_\nu$ with Gram–Schmidt orthogonalization of $\{c_\nu(\boldsymbol{G})\}$. In the calculation of the KS Hamiltonian, we need to calculate the density $n(\boldsymbol{r})$. For this, we first calculate $n(\boldsymbol{r})$ as

$$n(\boldsymbol{r}) = \sum_\nu f_\nu \Psi_\nu^*(\boldsymbol{r})\Psi_\nu(\boldsymbol{r}). \quad (3)$$

If we perform this in a straightforward way, we need the operations of $\mathcal{O}(N_G^2)$ for each band ν, and the total number of operations needed for the transformation from $\{c_\nu(\boldsymbol{G})\}$ to $\Psi_\nu(\boldsymbol{r})$ is of the order of $N_B N_G^2$, which is quite expensive. However, we can dramatically reduce the cost of the calculation using the fast Fourier transform (FFT) method, and the number of operations in equation (3) becomes $N_B N_G \ln(N_G)$. Although equation (3) is still the most expensive part in the calculations of small systems for many plane-wave DFT codes, the reduction of the computational cost by the FFT method is essential for the success of plane-wave DFT calculations.

The orthogonalization of the KS orbitals is also an expensive operation, which includes the calculations $\int \mathrm{d}\boldsymbol{r}\Psi_\nu(\boldsymbol{r})\Psi_{\nu'}(\boldsymbol{r})$ for all pairs of band indices $\{\nu, \nu'\}$. The

total cost of the operations is $\mathcal{O}(N_B^2 N_G)$, but we can see that it is only proportional to N_G. As we have seen, in the iterative method with the plane-wave basis set, there are no operations where the cost increases as fast as N_G^2. This is the reason why we can do efficient calculations even with large N_G. The iterative diagonalization technique, FFTs and *ab initio* pseudopotentials used in the plane-wave calculations are the key factors which make it possible to employ accurate but efficient DFT calculations. Using these techniques, with the increase in the computer power, the time required for solving KS equations has become smaller and smaller, and the system size for the target of DFT studies has become larger and larger. There was a report in 2002 of DFT calculations on a DNA system including hydrating water molecules, which consisted of 1194 atoms, including 138 water molecules [10].

However, this situation changed about 5–10 years ago. Recently, the growth in computer power comes mainly from the increase in the number of processors or cores, while the speed of each core or processor remains unchanged. The number of cores of the biggest supercomputers is currently reaching sub-millions. The Jaguar machine at Oak Ridge in the US has 224 162 cores, and the new Japanese supercomputer 'K' already has more than 700 000 cores. To utilize such computing power, it is essential to determine whether or not a technique or approach has good *parallel* efficiency. In this respect, the FFT method has a serious drawback. As is well known, the FFT needs all-to-all communication (i.e. each core communicating with all other cores) and the time required for communications will grow rapidly with the increase in cores or processors. As explained above, we cannot perform efficient plane-wave DFT calculations without the FFT technique. Thus, we need to introduce a different type of basis set which will be more suitable for parallel calculations.

In addition, there is another serious obstacle to increasing the system size in DFT calculations. When the number of atoms exceeds a few hundred, the orthogonalization of the KS eigenstates becomes the most expensive operation instead of the FFT. The CPU time for the FFT part is proportional to $N_B N_G \log(N_G)$ and it increases as $\mathcal{O}(N^2)$, since both N_B and N_G are proportional to the number of atoms N. On the other hand, the CPU time for orthogonalization increases as $N_B^2 N_G$, which is $\mathcal{O}(N^3)$. Once this part becomes the most expensive part, it is very difficult to make the system size larger.

From our brief survey of the field, we can see two key points which must be overcome to realize efficient DFT calculations on extremely large systems and massively parallel computers.

- Develop a method to calculate electronic structure which is suitable for massively parallel calculations [11, 12].
- Solve the electronic structure with better scaling than $\mathcal{O}(N^3)$, ideally with linear scaling.

We note that there has been considerable effort in this area within standard DFT codes, often focusing on molecular dynamics (MD) [13, 14]. However, there are limits to this effort, and so we must consider alternatives. For the first point, real-space methods are considered to have an advantage. There has been a concerted effort to develop practical real-space methods and significant progress has been achieved in the last decade. Several DFT codes using this technique are now available to researchers. Here, it is essential to understand whether or not these new methods maintain the advantages of plane-wave methods, such as high accuracy, ease of calculating atomic forces and systematic convergence. Regarding the second point, there were already several proposals for $\mathcal{O}(N)$ methods more than ten years ago[5], where the cost of the calculation is only proportional to N. Within empirical tight-binding (TB) methods, there has been a significant number of applications using such linear-scaling techniques [15, 16]. However, there were almost no examples, until very recently, where linear-scaling DFT methods were used for the purpose of actual scientific research. To replace conventional DFT methods with a new method, it is necessary for the method to have the same accuracy and stability as the standard methods, and reasonable efficiency. Compared with empirical TB methods, DFT calculations are more complex and have many potential sources of instability, especially in large-scale calculations. In this respect, the success of the first point is important also for the second point. Plane-wave DFT methods have been under intense development for over 25 years and are widely used; competing for efficiency is therefore difficult for linear-scaling methods, except for very large (thousands of atoms) systems. Identifying problems which require systems of this size can be a challenge, particularly to researchers used to the constraints of cubic scaling codes.

The main purpose of this paper is to review the recent progress of the $\mathcal{O}(N)$ methods. However, following our discussion above, we first survey recent progress in real-space methods. We then turn to the localization of electronic structure, first for Wannier functions and then the density matrix. In the major part of the review, we survey linear-scaling methods and related developments in seven different areas and consider extensions to standard DFT. Technical details (including non-orthogonality, electron number, parallelization and sparse matrices) are dealt with in a separate section which is mainly intended for practitioners in the field; however, it is important to note that high parallel efficiency is a key criterion for a successful linear-scaling code. Finally, we describe various implementations of the methods, as well as applications of linear-scaling DFT before concluding with a survey of the challenges facing the field.

2. Real-space methods

As touched on above, the use of real-space methods both for efficient parallelization and for modelling larger systems is well established, and has been reviewed elsewhere in detail [17–19]. It is also used extensively for combined quantum mechanical/molecular mechanical (QM/MM) simulations, which is used both in solid-state systems [20] and more commonly in biological systems [21]. Approaches to a real-space implementation of DFT can, like Gaul, be roughly

[5] The recursion method, described in section 3.2.3, dates back to the 1970s, while the first linear-scaling approaches were proposed in the early 1990s, for instance divide-and-conquer (section 3.2.2) and DMM and OMM methods (section 3.2.1)

divided into three parts[6]: finite difference methods; finite element (FE) methods and the use of local basis functions. In these methods, the kinetic, pseudopotential and exchange-correlation energies are found in real space exclusively. The solution of the Laplace–Poisson equation for the electrostatic potential, however, sometimes retains the use of reciprocal space.

In all cases, the advantage of real-space methods stems from *spatial locality*, which in turn leads to sparsity of the Hamiltonian; these ideas will re-appear throughout section 3. Finite difference approaches represent the electronic states directly on a fixed grid in real space, with a finite difference operator for the kinetic energy (KE). FE methods [22] use the piecewise continuous local basis of FE analysis. Local basis functions represent the wavefunctions in terms of local orbitals, often centred on the atoms (e.g. Gaussians). This spatial locality can lead to linear-scaling Hamiltonian building, and is also at the heart of the linear-scaling electronic solvers described in section 3. We describe these approaches in outline below, but without the intention of giving a detailed review of the methods; other reviews are cited for the interested reader.

2.1. Finite differences

Finite difference methods do away with a basis set entirely, and represent the wavefunctions directly by their numerical values on a grid; the grid spacing is one parameter by which the convergence can be judged. This approach requires an approximation for the differential operator used in calculating the KE and in solving the Poisson equation; the simplest approximation is found by expanding a function in positive and negative directions:

$$\psi(x_{n+1}) = \psi(x_n) + \psi'(x_n)h + \tfrac{1}{2}\psi''(x_n)h^2 + \tfrac{1}{6}\psi'''(x_n)h^3 \ldots \tag{4}$$

$$\psi(x_{n-1}) = \psi(x_n) - \psi'(x_n)h + \tfrac{1}{2}\psi''(x_n)h^2 - \tfrac{1}{6}\psi'''(x_n)h^3 + \cdots \tag{5}$$

where h is the grid spacing and $\psi(x_n)$ is the value of the function $\psi(x)$ at a grid point x_n. By adding these two equations, an approximation for $\partial^2\psi(x_n)/\mathrm{d}x^2$ can be derived which is accurate to second order in h:

$$\frac{\partial^2\psi(x_n)}{\partial x^2} \simeq \frac{1}{h^2}\left[\psi(x_{n+1}) + \psi(x_{n-1}) - 2\psi(x_n)\right] - \frac{1}{12}\frac{\partial^4}{\partial x^4}\psi(x_n)h^2 + \mathcal{O}(h^4). \tag{6}$$

The errors can be of either sign, depending on the derivatives and value of ψ, which has the consequence that the FD method is not variational; the loss of variational nature can make it harder to converge the parameters used in computational methods. Naturally, there are higher order approximations for the Laplacian which can be generated (see, for example, the algorithm in appendix A [17]); the order of the expansion is the other parameter that defines the convergence of these methods. An alternative discretization, the Mehrstellen discretization [23, 24], has also been used, and introduces a non-Hermitian Hamiltonian in exchange for greater accuracy of representation for a given order. In all cases, a higher order discretization leads to a larger range for the operator, which impacts on efficiency. There have been proposals for variational representations of the KE operator in real space [25, 26] which have been compared in detail [26]; these representations would alleviate the variational problem in finite difference methods.

Once the Schrödinger equation has been discretized on the grid, it can be written in a matrix form, with a size proportional to the number of grid points. However, the resulting matrices are sparse, owing to the locality of the real-space representation. Indeed, the main source of spread in the matrices is the KE term (where the order of the approximation chosen will directly affect the locality). The sparsity of the matrices makes the method ideal for massive parallelization [17] and efficient solvers. Solution of the Poisson equation is often accomplished directly on the grid via multigrid techniques [27], without recourse to FFTs which can cause problems to parallel scaling of plane-wave codes.

The FD technique has been reviewed extensively [17–19]. It has been applied by a number of groups [12, 24, 28–33] with extension to PAWs [34]. The solution of the equations can be accelerated by the multigrid method [17, 35], which has been extensively applied in at least one real-space DFT code [27]. FD methods can also be combined with localized orbitals, either fixed to the atoms [27, 36] or with adaptive localization regions [37]; these localized orbitals are an integral part of linear-scaling approaches.

2.2. Finite elements and local real-space bases

FE methods are well known from the engineering field, and their application to electronic structure methods has a long history [22, 38–40]. The method uses basis functions which are chosen to be piecewise polynomials, local in real space. The simplest possible FE basis consists of linear functions which are one on the defining grid point and zero beyond its nearest neighbours; cubic functions are more common [38, 40] (and indeed the blip functions mentioned below as a basis for the CONQUEST $\mathcal{O}(N)$ code [41] are functionally equivalent to FE basis functions).

In the FE method, the unit cell is divided into elements (the simplest of which are cubes, but the shape is in principle arbitrary so long as the simulation cell is filled). Once the elements and basis functions are chosen, the Schrödinger equation can be written as a matrix equation, as was the case for the FD method (and the similarities and differences between the two techniques are elegantly described by Beck [17]). In particular, the FE method introduces non-orthogonal basis functions, which leads to a generalized eigenvalue equation [39] (though this is a familiar problem in electronic structure techniques seen also for ultra-soft pseudopotentials). The mesh fineness can be different in different areas of the simulation cell, leading to efficiencies when large areas of vacuum are considered, or allowing all-electron calculations to be performed with appropriate resolutions in different areas of the cell. The method shares the locality advantages of the

[6] 'All Gaul is divided into three parts', Julius Caesar, *De Bello Gallico*, Chapter 1.

FD method while also being variational, and has been applied in various areas [22, 42–46].

Similar to FEs are another class of real-space basis functions which are local in real space, with many of the well-defined properties which make plane-waves valuable (e.g. orthonormality and systematic convergence). Daubechies wavelets [47] are a specific class of wavelets with attractive properties for electronic structure (particularly massively parallel or linear-scaling implementations). They are orthogonal and local in real and Fourier space; the use of wavelets as a general approach to electronic structure has been discussed in detail elsewhere [48]. It is also possible to make a multiresolution implementation [49]. Wavelets have been implemented as a basis in one major *ab initio* code (ABINIT).

Discrete variable representations [50] are another of this class of basis, and have been successfully applied to *ab initio* MD calculations [51]. It is intriguing to note the close relationship between these basis functions and the psinc functions [52] described below and used for a linear-scaling code (and earlier used to calculate KE in a real-space DFT code) [53]. Both of these basis sets have the property that they are non-zero only on one grid point, and zero on all others (known as *cardinality*; wavelets are known as semi-cardinal as they are cardinal only at one resolution). However, despite their attractive properties, these basis sets are not yet in widespread use. As computational resources shift towards multi-core processors it may be that their properties make them more attractive than plane-waves.

Lagrange functions form another cardinal basis set which have been proposed for electronic structure calculations [54, 55]. Blip functions (or b-splines) [41], which can be shown to be a form of FE, are local, grid-based non-orthogonal bases also used in linear-scaling methods. They have also recently been used for linear-scaling quantum Monte Carlo calculations [56]. The psinc functions mentioned above, which are periodic bandwidth-limited delta functions, are another local basis set [52]; interestingly, almost the same functions were derived in the context of optimal local basis sets [57]. One of the first *ab initio* $\mathcal{O}(N)$ methods proposed [58] used a plane-wave basis to represent localized orbitals.

2.3. Atomic-like orbitals

Functions which mimic the atomic wavefunctions near the ionic core are a popular choice of basis function, which make sound computational sense: they provide an excellent solution for much of real space and are spatially local. However, to correctly model the changes which occur to electronic structure on formation of bonds, variational freedom is required, including both radial freedom for the valence electrons and angular freedom (often solved by adding orbitals with higher angular momentum than the valence states). Numerical atomic orbitals (NAOs—used for all-electron calculations) or pseudo-atomic orbitals (PAOs—used with pseudopotentials for convenience) are in wide use in both standard and linear-scaling codes [59–67], though this is by no means an exhaustive list (other bases include Gaussian-based orbitals [68, 69], muffin-tin orbitals and augmented plane-waves). NAOs have been recently reviewed [70].

These basis functions are written as a radial function multiplied by a spherical harmonic (normally the real spherical harmonics):

$$\chi_{nlm}(\boldsymbol{r}) = R_{nl}(r)Y_m^l(\hat{\boldsymbol{r}}). \tag{7}$$

The formalism allows for efficient evaluation of matrix elements. Analytic operations are possible for the angular terms, while the radial terms are performed in reciprocal space with a very fine mesh [59, 71], or analytically, for Gaussians. Typically, functions are confined within a sphere which removes extended tails and results in sparse matrices. Even for conventional codes, this can render the building of Hamiltonian matrices linear scaling (discussed further below, in section 2.5), reducing a significant cost.

The major drawback with these basis sets is the lack of systematic convergence: the number of radial functions in a given angular momentum channel can be increased (often known as multiple zeta or multiple valence, so that two radial channels are notated DZ or DV) and extra angular momentum channels can be included but there is no clear rule as to how functions should be added to systematically improve the energy. There have been studies which show that convergence can be achieved, and which suggest routes to creation of convergent basis sets [59, 60, 62] but these schemes lack the simplicity of basis sets with a single parameter (e.g. the KE cutoff for plane-waves, or grid spacings for analytic real-space methods); an example of the convergence with respect to basis set size is shown in figure 1(*a*).

The problem of confinement has generated a number of different solutions. The simplest approach is effectively to impose an infinite potential well of some radius on the atom [71], though this has the side effect that the derivative of the orbital is discontinuous at the boundary, which can cause problems with the calculation of forces and stresses. The confinement excites the atom slightly and mimics the effect of condensation into a molecule or condensed matter environment. However, it is not clear how to confine the atom, particularly as different orbitals will have different ranges.

To avoid the discontinuity produced by an infinite potential, a number of suggestions have been made for an alternative potential (surveyed, along with the methodology known as *ab initio* TB, in a review [73]). Confining potentials suggested include simple polynomials: (r^2 [74], r^6 [75]); smoothing the free atomic wavefunctions with an exponential using a cutoff and width over which the smoothing is applied [65]; an exponential potential applied between two points [72]; a cubic truncation between two points (applied to the bare atom potential) [62]; and a product of an exponential and $1/r^2$ on the free atom (for all-electron calculations) [60]. All these different schemes produce a smooth transition to zero in the tails of the orbitals.

As well as methods to confine the orbitals, there are many different methods to generate the basis sets themselves. These can be split into two approaches: first, how to generate a set of either pseudo-atomic orbitals or numerical atomic orbitals (depending on whether or not a pseudopotential is used); and second, how to use these as basis functions (either as they are or combined into other functions). We will consider these two problems in turn.

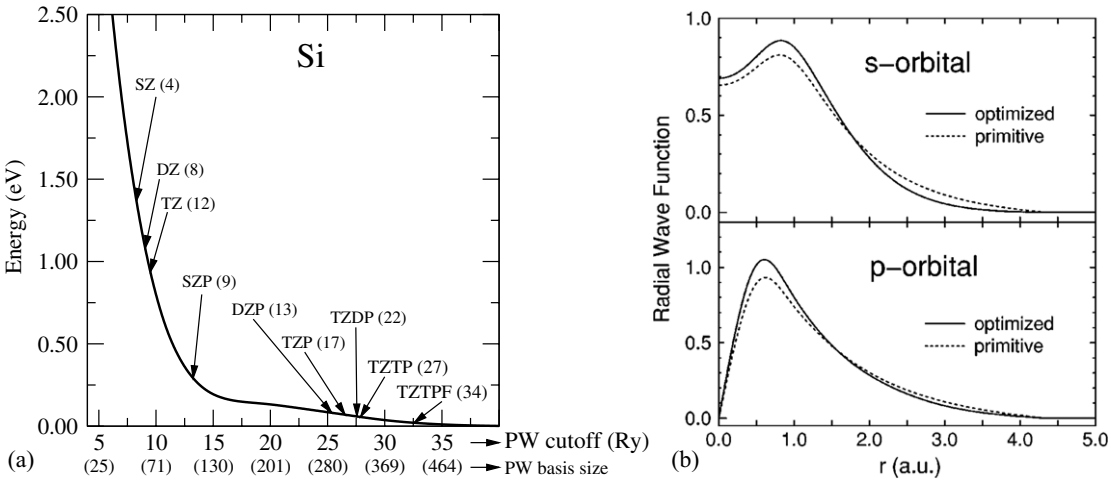

Figure 1. (*a*) Convergence of energy with PAO basis size for bulk silicon in SIESTA [72]. (*b*) Change in shape for atomic orbital following optimization within OPENMX [62]. Reprinted figures with permission from [72] and [62]. Copyright (2001) and (2003) by the American Physical Society.

As there is considerable flexibility in deciding, for instance, at what radius to cut off a function, or how many radial functions to use for each angular momentum channel, much effort has gone into deciding how to generate accurate basis sets with minimal computational effort and human intervention. An early approach to the cutoff problem [76] was to use the energy change found on confining individual orbitals (typically in the range 50–300 meV) as a single, balanced criterion. The advantage is that one parameter can be used for orbitals of different inherent sizes, and this is used extensively in the SIESTA code (see section 5.1 for more details). Another approach within SIESTA is to optimize the orbital shape relative to a highly converged plane-wave calculation by varying the confinement potential, but applying a fictitious pressure-like quantity [77] to stop expansion of the orbitals beyond a reasonable size. This idea of optimizing the confinement has also been applied to a damping function multiplying the orbitals [78].

Spillage [79, 80] is an idea which can be used when optimizing atomic orbitals. In a periodic system, it is defined as

$$S = \frac{1}{N_k}\frac{1}{N_B}\sum_{\boldsymbol{k}}\sum_{i=0}^{N_B}\langle\psi_i(\boldsymbol{k})|\left[1 - P(\boldsymbol{k})\right]|\psi_i(\boldsymbol{k})\rangle \qquad (8)$$

where N_k is the number of k points, N_B is the number of bands, $|\psi_i(\boldsymbol{k})\rangle$ is a Bloch state found with a plane-wave code and the projection operator $P(\boldsymbol{k})$:

$$P(\boldsymbol{k}) = \sum_{\mu\nu}|\phi_\mu(\boldsymbol{k})\rangle S^{-1}_{\mu\nu}\langle\phi_\nu(\boldsymbol{k})| \qquad (9)$$

with $S_{\mu\nu}$ the overlap between atomic orbitals (see section 4.1 for more details on non-orthogonality). The spillage measures how well atomic-like orbitals can reproduce the wavefunctions from another calculation (e.g. highly converged plane-wave basis calculations). The spillage is used to guide optimization of atomic-like orbitals and to investigate the generation of transferrable basis sets. For instance, basis sets using orbitals from both neutral and positively charged ions, and optimizing cutoff radius to minimize spillage [65] were found to be transferrable and accurate. Spillage has also been used to optimize PAOs with radial functions expanded in terms of Bessel functions [81]; the functions are fit to converged plane-wave calculations of dimers.

Generation of basis sets consisting of atomic-like orbitals can also be performed in a semi-systematic way. For instance, generation of successive solutions of a confined atom with increasing numbers of nodes for a given angular momentum channel [62], or building a large set of functions including different cutoffs, angular momenta, Rydberg functions and extend basis functions, and ordering the set by searching for the function which lowers energy most on addition to the existing set [60]. A method for the direct optimization of the radial function [82] within a self-consistent loop has also been given to yield an optimal minimal basis (this parallels to some extent the representation of support functions or generalized Wannier functions by b-splines or psincs mentioned above in section 2.2; it has also been suggested that localized functions can be expanded in terms of spherical waves [83] for free electrons, and recent analytic developments [84] have simplified and improved the scaling of this method). Methods for optimizing Gaussian basis sets are also available [69, 85].

While a large set of atomic-like orbitals can be used directly to represent the wavefunctions, it can be more efficient to combine them into a smaller set of functions. Polarized atomic orbitals [86] are one way to do this: a minimal set of polarized atomic orbitals is defined in terms of a large basis set of standard quantum chemistry orbitals. As is to be expected, the contraction results in a small increase in total energy, but the convergence is good, and the error is linear in system size (indicating size extensive behaviour, and local errors); moreover, the structural relaxation is reliable. The idea has been refined to extract polarized AOs from molecular orbitals [87] which is closely related to the extraction of Wannier functions described in section 3.1. It has also been extended [88] so that minimization of polarized AOs and the density matrix is separated; this was implemented within an $\mathcal{O}(N)$ method and shown to be effective. A related method [89]

uses Cholesky decomposition to extract localized molecular orbitals from the density matrix.

The OpenMX code mixes large numbers of PAOs (typically six per angular momentum channel) into a set of orbitals equivalent to a DZP basis [62] with the PAOs simply generated as orthonormal functions (by increasing the number of nodes) for a confined atom; the change in radial function following optimization is illustrated in figure 1(*b*). An optimized set of PAOs has been generated for calculations of biological molecules (covering H, C, N, O, P and S) [63], and it was suggested that optimization at each step of the minimization is not necessary. A similar approach is used for the representation of support functions with PAOs in Conquest [66]. It is important to note, though, that there are certain symmetry-imposed restrictions on the number of support functions that can be used: for instance, trying to represent four support functions only with $l = 0$, $l = 1$ and $l = 2$ will give incorrect answers in bulk crystals (breaking symmetry). A brief overview of approaches to PAO mixing, as well as a scheme for mixing two PAOs (neutral and 2^+ atom) has been given [90]. A comparison of the method used in OpenMX and using spillage [91] shows that comparison to plane-wave results with chemical accuracy can be achieved with localized orbitals, though two functions per angular momentum channel (DZDP) are generally needed (and up to $l = 2$). A method combining Gaussians from multiple sites into a minimal basis on each atom [92, 93] has been proposed. It uses a filtration algorithm (cf the FOE method in section 3.2.3) which removes unwanted high energy components to optimize the orbitals, and may allow large systems to be calculated on modest resources.

Overall, it should be clear that there is considerable effort being made to understand and optimize atomic-like orbital approaches. Given the history of the method, it is perhaps a little surprising that there is still so much work to do, but that simply reflects the impossibility of finding a perfect basis set for the diverse environments modelled by DFT.

2.4. Representing localized orbitals

When performing $\mathcal{O}(N)$ calculations, many codes represent the density matrix (described in section 3.2) in terms of localized orbitals, $\phi_{i\alpha}(\boldsymbol{r})$; for instance, onetep calls them non-orthogonal generalized Wannier functions (NGWFs) and Conquest calls them support functions. These functions can be simply individual PAOs, or more generally represented in terms of a basis set, either the atomic-like orbitals of section 2.3, or the local real-space basis sets described in section 2.2.

Using atomic-like orbitals is convenient and gives a relatively small basis size. In the limit of a single PAO per localized orbital, of course, there is no optimization required, which removes a level of complexity from the calculation; however, it is important to note that there can be problems inverting the overlap matrix in the case of double zeta basis sets (discussed further in section 4.1). Atomic-like orbitals also suffer from basis-set superposition error [94]; while the magnitude of the error reduces with basis set size, it is still significant [95], and significantly worse for more contracted basis sets [67], though can be corrected very successfully. These basis sets are widely used (e.g. in OpenMX, siesta and Conquest).

The local, real-space basis sets, such as the psincs used in onetep and the b-splines used in Conquest, can be converged systematically to the plane-wave limit [41, 96] and are free from basis-set superposition error [97]. The resulting orbitals will be more transferable as they are optimized *in situ*, and possibly more accurate. However, these basis sets require more computational effort to converge than PAOs, and calculating small energy differences between structures will require tight tolerances on the minimizations.

The real question when deciding on a basis set is that of accuracy versus computational cost. The minimization is variational, which means that less effort will be required once the initial functions have been converged, and also that library functions could be calculated and read in as a starting point. In terms of matrix multiplications, which lie at the heart of linear-scaling codes, real-space basis sets have a significant advantage: as a minimal number of orbitals can be used, the multiplies are significantly faster. For instance, for group IV elements such as C and Si, the computational cost to go from four orbitals (minimal) to nine orbitals (the smallest number possible when using polarization functions in a PAO basis) is a factor of 11, and going to thirteen orbitals (a double zeta plus polarization) is a factor of 34. These factors can offset the extra time required but there is no single correct answer.

2.5. Hamiltonian building

Many methods which are not linear scaling in the search for the ground state nevertheless use localized real-space basis sets, and rely on this locality to build the Hamiltonian in a linear-scaling manner [61, 98–100]. The computational effort required for Hamiltonian building is significant, and for a few hundred atoms with localized orbitals the resulting matrix can be exactly diagonalized efficiently, with most effort being spent in the Hamiltonian building. The Hamiltonian typically is made up of different terms, which can require different treatments: KE; electron–ion interaction (either via bare Coulomb term or pseudopotentials); Hartree energy; and exchange and correlation energies. The KE is inherently local (and only shows spread for high-order finite difference methods) and will not be considered further.

As individual basis functions are local in space, the integrals required to form the Hamiltonian can be reduced from $\mathcal{O}(N^3)$ scaling (the integrals between all pairs of basis function (N^2) must be evaluated over the whole system (N)) to $\mathcal{O}(N)$ scaling. For non-local pseudopotentials, the electron–ion interaction can be evaluated using a separable form to give integrals between localized orbitals and non-local projector functions, followed by a matrix multiplication. The local part is more complex: it is formally written as $H_{ij}^{\mathrm{local}} = \sum_k \langle \phi_i | V_k | \phi_j \rangle$ for the matrix element between atoms i and j, which involves three-centre integrals and hence poor efficiency. The standard solution involves integration on a grid (making the potential by summing over atoms k *before* the integration is performed), often using the charge and potential

from the neutral (sometimes free) atom [59, 71] to form a smoothly varying function (which is relatively insensitive to grid spacing). A method to make the neutral-atom potential separable has been proposed [101] which involves expanding the potential in terms of local functions; this allows the potential to be evaluated in the same way as the non-local potential. As the potential is spherically symmetric, it can be expanded with radial functions and spherical harmonics (up to $l = 6$ was found to be sufficient for convergence [101]). This procedure removes any grid dependence apart from a charge density difference $\delta n(\boldsymbol{r}) = n(\boldsymbol{r}) - \sum_i n_i(\boldsymbol{r})$ with $n_i(\boldsymbol{r})$ an atomic density.

The Hartree potential (found as the solution of the Poisson equation) is often solved using FFTs, which are strictly $\mathcal{O}(N \log N)$; the standard approach to the electrostatic solution of the ionic problem, the Ewald sum, scales as $\mathcal{O}(N^{3/2})$, though the use of a neutral-atom potential [71] removes the need for this step. The other method in common use for electrostatic problems is the fast multipole method, which may scale as $\mathcal{O}(N)$ asymptotically (see, for instance, discussions [98, 99, 102–104]). Other methods used include density fitting [105], FFTs combined with a wavelet solution for surface problems [106] and combining FEs and Gaussians for the direct solution of the Poisson equation with the fast multipole method for calculation of the boundary conditions [107].

The calculation of the exchange-correlation matrix is generally straightforward, but a number of different approaches have been given, both exact and approximate [71, 75,104, 108–110]. The question of calculating exact exchange within DFT (or Hartree–Fock) has been studied extensively, and a number of approaches which scale linearly with system size have been derived [111–119]. This ensures accuracy and efficiency are possible within local orbital, real-space codes. The route to linear-scaling construction of the Hamiltonian building is clear. Now we turn to consider the solution for the ground state of the system.

3. Linear-scaling methods

As we have noted above, significant savings can be made even for conventional eigenvalue solvers if a basis set which is local in real space is used to represent the wavefunctions. The Hamiltonian building process becomes linear scaling, leaving the solution for the eigenstates as the most expensive part, and the part with the worst scaling. It is natural to consider whether this can be improved as well; herein lies the heart of the development of linear-scaling codes.

A natural first point to consider is that the KS formulation of DFT introduced the wavefunctions as an aid to solution, not as an integral part of the formalism. Indeed the Hohenberg–Kohn theorems rely only on the charge *density* of the system; we might ask whether a search over charge densities might not be a route to finding the electronic ground state. This leads to the approach known as orbital-free DFT (OFDFT), discussed in detail in section 3.2.5; in brief, it requires an approximation for the KE which can reduce accuracy, but is used as an efficient method for calculations on large metallic systems.

Instead of the charge density, it is more helpful to work in terms of the *density matrix*, which is defined formally in terms of the eigenstates of the system as

$$\rho(\boldsymbol{r}, \boldsymbol{r}') = \sum_n f_n \psi_n(\boldsymbol{r}) \psi_n^\star(\boldsymbol{r}') \tag{10}$$

where n indexes the eigenstate and f_n gives the occupancy of the state. As we are still within the KS approach, this is the single particle, two-point density matrix (not to be confused with the many-body density matrix familiar from statistical mechanics and quantum information). In operator notation, the finite temperature density matrix can also be written as the Fermi function of the Hamiltonian [120]: $\hat{\rho} = 1/(1 + \exp[(\hat{H} - \mu)/k_\mathrm{B}T])$. Many of the properties of the density matrix were investigated and summarized by McWeeny [121]. A consequence of quantum interference effects on the density matrix is that it is ranged:

$$\rho(\boldsymbol{r}, \boldsymbol{r}') \to 0, |\boldsymbol{r} - \boldsymbol{r}'| \to \infty. \tag{11}$$

However, the functional details of how the decay proceeds are rather complex, and can be related to the localization of Wannier functions. We consider these ideas in the next section.

3.1. Density matrix properties and Wannier functions

The Bloch states found when solving the Schrödinger equation for a periodic system (as is done in most electronic structure codes) are delocalized and spread throughout the unit cell, and hence the entire system. There are many advantages to using localized functions, and the arbitrary phase associated with the Bloch states (which can be scaled by $\mathrm{e}^{-\mathrm{i}\phi}$ with no change to the properties of the system) gives the freedom to do this. The most important route to understanding localization is via Wannier functions. Wannier functions have been used extensively in electronic structure theory [122–125]. They are formally defined for a periodic potential, with Bloch wavefunctions $|\psi_{n\boldsymbol{k}}\rangle$ with n labelling a band. Then, for the unit cell at $\boldsymbol{R}$, we can define the Wannier function as

$$|w_{\boldsymbol{R}n}\rangle = \frac{V}{2\pi} \int \mathrm{d}\boldsymbol{k} \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{R}} |\psi_{n\boldsymbol{k}}\rangle \tag{12}$$

with

$$\psi_{n\boldsymbol{k}}(\boldsymbol{r}) = \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{r}} u_{n\boldsymbol{k}}(\boldsymbol{r}) \tag{13}$$

and $u_{n\boldsymbol{k}}(\boldsymbol{r})$ the periodic part of the Bloch function. The inverse relationship allows us to write the wavefunction in terms of the Wannier functions:

$$|\psi_{n\boldsymbol{k}}\rangle = \sum_{\boldsymbol{R}} \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{R}} |w_{\boldsymbol{R}n}\rangle. \tag{14}$$

It is important to note the considerable freedom in the choice of Wannier functions that the arbitrary phase of the Bloch functions gives. The localization of Wannier functions is closely related to the range of the density matrix (considered below) [126]; it is also important in considering insulating behaviour against metallic behaviour in condensed matter systems [127]. The equivalent theory for non-periodic systems

was developed for localized molecular orbitals [128] which also can be linked to pseudopotential theory [129].

The study of localization of Wannier functions is far from trivial. The earliest results [123] showed that the decay was exponential with distance for a one-dimensional (1D) centrosymmetric crystal, using complex wave vectors. This was extended to three dimensions [125] for periodic solids with no degeneracy (with the decay rate related to the distance of the branch surface from the real axis), as well as general 1D [130] and three-dimensional (3D) [131, 132] solids (though only in the TB limit for the 3D case). Non-periodic systems in one dimension have been shown to exhibit exponential localization [133], and in the case of a localized perturbation, such as a defect, the Wannier functions converge to the periodic functions as the distance from the perturbation increases; this result has also been extended to three dimensions [134]. The general relation between eigenfunctions of a Hamiltonian and localized orbitals (leading to generalized Wannier functions) was investigated thoroughly [135]. The exact form of the decay of Wannier functions in 1D has been investigated in detail [136], and it was found that it can be written as a power law multiplying an exponential. If the rate at which the functions decay is given by $x^{-\alpha}\mathrm{e}^{-hx}$ for Wannier functions in 1D, a value of $\alpha = 0.75$ describes orthonormal Wannier functions, while non-orthonormal Wannier functions result in $\alpha = 0.5$ or, with careful construction, $\alpha = 1.5$. A more general, and formal, study of localization has shown that Wannier functions demonstrate exponential localization in insulators with a Chern number[7] of zero (i.e. which are time-reversal symmetric) in 2D and 3D [139]. A recent study of many-body (quantum Monte Carlo) Wannier functions and localization [140] has shown that localization (except for strongly correlated systems) is similar for one-electron and many-electron systems, with the difference related to the correlation hole. This builds on earlier work on natural Wannier functions in correlated systems (defined in terms of the natural orbitals) where similar localization properties were found [141]. An important study [142] of the number of iterations required to reach convergence in a given system, and how this number of iterations scales with system size, found that Wannier-representable insulators can be considered truly $\mathcal{O}(N)$, with the time to the ground state not dependent on system size (though some of the results of this paper have been shown to be pessimistic [143]).

It is not the intention of this review to cover all aspects of Wannier function theory and their use, though we summarize results relevant to linear-scaling methods below. There are excellent reviews on this subject, particularly as it relates to polarization [138, 144]. In quantum chemistry, the equivalent localization procedure (though without Bloch states) is known as Boys–Foster localization [145]. The modern approach to polarization relies on a definition in terms of the expectation values of the position operator in terms of Wannier functions of occupied bands [146]. Much of the modern theory of polarization [144], particularly relating to the Berry phase, is concerned with—and overlaps with—definitions of localization and the difference between insulators and metals; full details can be found elsewhere [147–150]. Recent important developments in the field also relate to topological insulators, and the wealth of physics contained therein.

Creating Wannier functions is a difficult problem, due to the many possible different definitions of functions themselves, and issues with composite bands. An early proposal showed that removing the orthogonality constraint created ultra-localized functions [151]. Kohn presented a variational method [124] and an efficient iterative method [152], both to find generalized Wannier functions. A key development in the use of Wannier functions was a reliable method to produce maximally localized Wannier functions (MLWFs) [137]. The freedom in phase of the Bloch states results in the Wannier transformation being underdetermined. The lack of determination allows extra restrictions to be placed on the Wannier functions, without loss of generality. The criterion used was to find the Wannier functions which minimized the functional Ω, defined as

$$\Omega = \sum_n \left[\langle r^2\rangle_n - \langle \boldsymbol{r}_n\rangle^2\right], \tag{15}$$

for a sum over the bands n in the simulation, with $\langle \boldsymbol{r}_n\rangle = \langle w_{0n}|\boldsymbol{r}|w_{0n}\rangle$ and $\langle r^2\rangle_n = \langle w_{0n}|r^2|w_{0n}\rangle$; w_{0n} is the nth Wannier function in the cell at the origin (cf equation (12)). Despite the real-space definitions, the transformations can be written in terms of the Bloch wavefunctions in reciprocal space. This technique has found widespread application throughout the first-principles community, and has been shown to be effective for disordered systems [153] and entangled bands [154, 155]; an efficient, iterative approach to forming MLWFs has been given [156]. Examples of MLWFs in silicon (used for linear-scaling evaluation of the exchange potential and energy [119]) are shown in figure 2(a), with clear sp^3 symmetry. The MLWFs from the entangled, narrow d bands in Cu, which overlap with a wide s band, are shown in figure 2(b) following disentanglement; the d-symmetry of these functions is clearly seen.

There have been many further developments in finding and using Wannier functions. Allowing some unoccupied bands to mix with the occupied bands [157] allows better localization, and in some cases a more intuitive picture of bonding (it is important to note that including unoccupied bands has been proposed before [154] for entangled bands; entangled bands overlap with other bands across the Brillouin zone, as opposed to, for instance, the isolated valence bands of an insulator). A linear-scaling technique has been demonstrated for the creation of Wannier functions [158], which uses projection as others have before [130, 131, 133, 135], while another technique derived separate eigenvalue equations for each Wannier state [159]; an alternative approach builds Wannier functions using perturbation theory to correct a simple initial approximation [160]. Another linear-scaling approach [161] starts from the ground-state density matrix and derives the MLWFs from there. A dynamic approach allowing on-the-fly localization during MD shows improvements in speed of simulation compared with normal band techniques [162]. Several of the linear-scaling methods for finding the ground

[7] The Chern number is related to the Berry connection; the Berry phase is an important part of the modern theory of polarization and is discussed extensively elsewhere [137, 138].

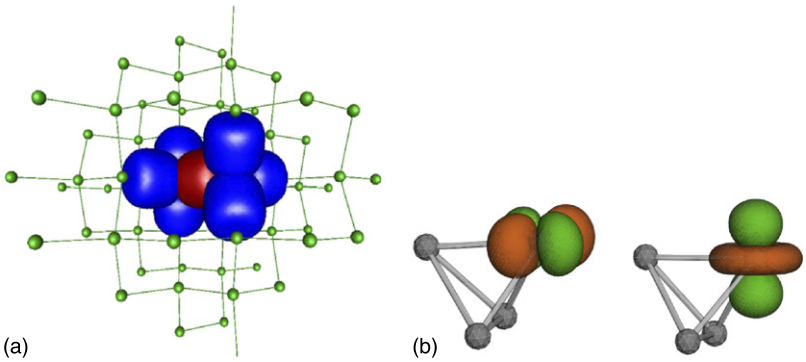

Figure 2. (*a*) MLWFs in silicon used in calculating exact exchange with linear-scaling effort (from [119]). (*b*) MLWFs in Cu (from [154]). Reprinted figures with permission from [119] and [154]. Copyright (2009) and (2001) by the American Physical Society.

state described in section 3.2.1 find the Wannier functions for the system (specifically those methods starting with the work of Mauri and co-workers [163, 164] and Ordejón *et al* [165, 166]), and this approach has been used for the calculation of polarization [167], while non-orthogonal Wannier functions are used for other problems, for instance for a self-consistent implementation of the DFT+U method [168]. The problem of finding non-orthogonal localized molecular orbitals [128, 169] can be simplified [170] by defining centroids based on single, double or triple bonds [171], though this requires some input and chemical intuition. A more recent study has investigated the general localization properties of bases for eigenvector problems [57]. An approach similar to Wannier functions to generate a minimal basis of quasi-atomic orbitals (QUAMBO) [172] for post-processing uses occupied and unoccupied states, and forces the orbitals to be as close as possible to free-atom orbitals; the method is applicable to metallic as well as insulating systems.

We also mention that localization is used in quantum chemistry (see section 3.2.7) to improve the scaling of perturbative and more accurate methods; for instance, a recipe to create localized orthonormal orbitals for fast MP2 calculations has been developed [173]. Wannier functions (and other localized orbitals) are becoming extremely powerful tools in extending the accuracy of DFT and Hartree–Fock methods, and we expect to see their use becoming widespread over the next few years.

The localization of the Wannier functions for a system is intimately related to the localization of its density matrix. This is easily seen as the density matrix can be written in terms of the bands, as in equation (10), and is unaffected by a unitary transformation of the bands; hence, the Wannier transformation allows the density matrix to be written in terms of the Wannier functions, and the localization properties follow. We can write [174]

$$\rho(\boldsymbol{r},\boldsymbol{r}') = \sum_n \sum_{\boldsymbol{R},\boldsymbol{R}'} W_n(\boldsymbol{r},\boldsymbol{R}) F_n(\boldsymbol{R}-\boldsymbol{R}') W_n^\star(\boldsymbol{r}',\boldsymbol{R}') \quad (16)$$

where the occupation matrix is defined by the Wannier transform of the occupancy $F_n(\boldsymbol{R}) = \frac{V}{2\pi}\int \mathrm{d}\boldsymbol{k} \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{R}} f_{n\boldsymbol{k}}$.

The general principle that the electronic structure of a system is localized is summarized by Kohn's principle of nearsightedness [175]. The principle is defined in terms of a typical de Broglie wavelength found in the wavefunction of the ground state, which in turn defines a local volume; the thermal de Broglie wavelength $\lambda = \sqrt{h^2/3m_\mathrm{e}k_\mathrm{B}T}$ might be a reasonable definition to start from [176]. Changes to distant parts of the system (such that the distant part is far from all points in the volume) have a negligible effect on the electronic structure in the local volume. The nearsightedness of electronic structure for systems with a non-zero gap can be expressed as

$$\rho(\boldsymbol{r}-\boldsymbol{r}') \sim \mathrm{e}^{-\gamma|\boldsymbol{r}-\boldsymbol{r}'|} \quad (17)$$

with $\gamma > 0$, though, as noted above, an algebraic prefactor can be defined. Much work in recent years has sought to relate γ to the properties of the system.

As with the Wannier functions, there has been considerable effort devoted to quantifying the localization (or, equivalently, the range) of the density matrix; an excellent overview is given by Goedecker [2]. The most elegant and appealing, as well as intuitive, results suggest that the decay rate should depend on the gap. Kohn [126] showed that the decay rate for Wannier functions is proportional to $\sqrt{m^\star\Delta}$ for a gap Δ (see also [123]); as $m^\star$ can be shown to depend on Δ [174], this gives an overall dependence on Δ. It was also shown that projection operators for specific bands localize exponentially at large distances in the one-electron approximation [130] and that in periodic solids with no degeneracies there is an exponential decay related to the distance of the branch surface from the real axis [125] (extending the earlier work on Wannier function ranges [123]); remembering that the density matrix is a projection operator for the whole system we see the importance of these results.

There is still no complete analytical understanding of the range of the density matrix. One study used Chebyshev polynomials [176] to explore the properties of the density matrix and the complexity of different linear-scaling methods. This suggested that the range was related to the gap (not exponentially necessarily, but $\propto 1/\sqrt{\Delta}$) for insulators; in metals, a finite electronic temperature, T, is required to localize the density matrix, and the range was found to be $\propto \sqrt{T}$. However, these results were subsequently shown to be incomplete. Using Fourier analysis, it was shown [177] that metals (again at finite electronic temperature T) show exponential localization proportional to $k_\mathrm{B}T/k_\mathrm{F}$ both

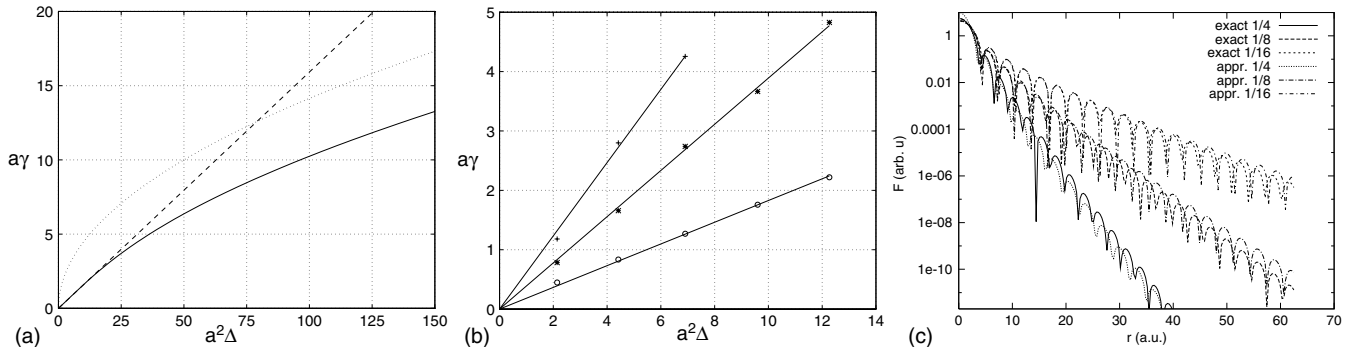

Figure 3. (*a*), (*b*) Density matrix range (γ) as a function of gap (Δ) for (*a*) a 1D insulating system and (*b*) a 3D insulating system (from [174]). (*c*) Density matrix range in a metal in a range of temperatures, given as a fraction of k_F (from [177]). Reprinted figures with permission from [174] and [177]. Copyright (1999) and (1998) by the American Physical Society.

in real space and also in Fourier space. These properties were studied further [178] using wavelets (see section 2.3), confirming the Fourier space nearsightedness. A careful analytic study of simple systems also found that decay is $\propto \Delta$ for semiconductors and $\propto T$ in metals at low T [174] (certainly in the weak binding limit); the study found that some materials show decay $\propto \sqrt{\Delta}$ in the TB limit, but the behaviour is complex (depending in detail on atomic potentials). At low temperatures, the decay is $\propto k_B T/k_F$ (in agreement with previous work [177]) with $\sqrt{T}$ behaviour at high T. It was also shown [136] for 1D systems that the density matrix decays exponentially (in the same way as Wannier functions) but with a prefactor of $\alpha = 0.5$. Overall, there is good evidence to support the exponential localization of the density matrix in insulators, and in metals at finite electronic temperature, though detailed behaviour depends strongly on the system.

Further work on the principle of nearsightedness [179] underlined the earlier results that the decay is proportional to Δ for 1D to 3D systems, and some interacting systems. There is a well-motivated suggestion that disorder increases the density matrix range for insulators and decreases it for metals. Numerical studies of the range of the density matrix for the Anderson model with varying levels of disorder [180] have shown that coherence is strongly affected by disorder, with the exponential localization depending inversely on disorder. Density matrix decay for both Hartree–Fock and DFT has been plotted for different systems [181]. It has also been shown that the correlation between fermionic operators is exponentially localized at non-zero temperatures [182]. Figures 3(*a*) and (*b*) show the behaviour of γ as a function of gap (Δ) for model insulating systems. Figure 3(*a*) plots the behaviour for a periodic 1D potential, period a; this is clearly linear for a weak potential (small gap), while stronger potentials are more complex. Figure 3(*b*) shows the behaviour for a simple cubic array of Gaussian potentials along different directions in the crystal; the linear dependence of density matrix decay on gap is clear in all directions. By contrast, figure 3(*c*) shows the spatial decay of the density matrix in a metal, comparing exact results with a simple model which simplifies to give decay proportional to $k_B T/k_F$ for different values of $k_B T$ as a fraction of k_F.

Further analytic studies have been performed on specific systems. A study of a simple cubic TB model [183] allowed the derivation of analytic results, showing both the exponential decay and a power-law prefactor for the density matrix. The decay length found depended on both the gap and the hopping integrals (specifically, the decay length $\lambda \propto \Delta/t$ for hopping integral t); this work was extended to TB models of metals [184], where a power law decay was found. Related earlier work using a simplified approach to DFT (a Sankey–Niklewski approach) for Si, C and Al [185] found algebraic decay for Al which was very close to the nearly free electron model, with anisotropic exponential decay in Si and C (with the directions along bonds showing the slowest decay). A one-body density matrix with TB model for insulators and decay found 1D and 2D analytic results with a universal power law times exponential but second energy scale emerges when hopping modulates, so that the decay is not entirely gap dependent [186, 187]. An extension of this work to 2D anisotropic hopping [188] showed exponential localization, but not with the same form as isotropic hopping; in particular, the correlation length does not vanish as the gap closes.

While much progress has been made in extending the analytic results for density matrix range, particularly based on Wannier functions, it is clear that there is still work to be done. The relationship of localization to the important solid-state problems of disorder and polarization (particularly as they extend into the new research areas of topological insulators) is fascinating. Furthermore, this area is relevant to problems in graphene, where defects can have an extremely strong effect on the electronic structure. Overall, the increasing complexity of a system changes the decay properties, and the whole area is far from simple.

3.2. Solving for the density matrix

Having established that the density matrix, and indeed the electronic structure, of matter is nearsighted, we can now turn to the matter of exploiting this nearsightedness in the search for the ground state. There have been developments in this field for over 30 years and, as might be expected, papers comparing the different methods [189–192] and general reviews of the subject have already been written [2, 70, 193–198]. This review will, naturally, build on these excellent surveys.

The fundamental quantity in linear-scaling techniques is the density matrix, and the fundamental property of the density

matrix is its sparsity. While there are methods which operate using, for instance, Wannier functions, they still rely on the density matrix as the fundamental quantity, and the short range of the functions to achieve linear scaling. To obtain a linear-scaling method, we must impose a range on the density matrix, which is a controllable approximation. An appropriate localized basis must used, which will make matrices sparse; however, they must also be stored and operated on as sparse matrices, which requires significant extra effort. Once these preparations have been made, the computational effort required to reach the ground state should scale linearly with system size. It is important to note that the choice of truncation is imposing an additional constraint on the system, and that there will therefore be an extra error (or energy difference) compared with an unconstrained problem.

The search for the ground state in terms of the density matrix cannot be made in terms of the original, six-dimensional object $\rho(\boldsymbol{r}, \boldsymbol{r}')$ (which is defined simply in terms of the bands of a system in equation (10) above). The most common approach is to work in terms of localized orbitals (also called *support functions*), and to assume that the density matrix is separable [199, 200]:

$$\rho(\boldsymbol{r}, \boldsymbol{r}') = \sum_{i\alpha, j\beta} \phi_{i\alpha}(\boldsymbol{r}) K_{i\alpha j\beta} \phi_{j\beta}(\boldsymbol{r}'), \tag{18}$$

where atoms are labelled with Roman letters i and j and orbitals on atoms with Greek letters α and β (we note that the density matrix is often notated only in terms of orbitals, as K_{ij}). The only approximation made by this assumption is that the original density matrix had a finite number of non-zero eigenvalues (which is at most the number of local orbitals used). However, this is not restrictive in the context of electronic structure calculations. Much of the rest of this section is devoted to discussing different methods for finding $\{\phi_{i\alpha}(\boldsymbol{r})\}$ and $K_{i\alpha j\beta}$. However, there are two further conditions which must be considered

(a) *Correct electron number*. The electron number is given by

$$N_\mathrm{e} = 2\mathrm{Tr}\,[KS] \tag{19}$$

where $S_{i\alpha j\beta} = \langle \phi_{i\alpha} | \phi_{j\beta} \rangle$ is the overlap matrix between the localized orbitals and we assume spin degeneracy, giving the factor of two. We will consider methods for imposing the correct electron number in section 4.2.

(b) *Idempotency*. The density matrix is a projector onto the occupied subspace, and must have eigenvalues of either zero or one. Another way of writing this is, in both operator and real-space notation,

$$\hat{\rho}^2 = \hat{\rho} \tag{20}$$

$$\rho(\boldsymbol{r}, \boldsymbol{r}') = \int \mathrm{d}\boldsymbol{r}'' \rho(\boldsymbol{r}, \boldsymbol{r}'') \rho(\boldsymbol{r}'', \boldsymbol{r}'). \tag{21}$$

This requirement is rather hard to impose exactly, and many approaches adopt a weaker restriction (often known as weak idempotency) [194, 201], where

$$0 \leqslant \lambda_\rho \leqslant 1 \tag{22}$$

for the eigenvalue of the density matrix λ_ρ. McWeeny [121] showed how an iterative scheme could be used to force an approximately idempotent matrix to exact idempotency; this will be discussed in section 3.2.1.

As alluded to above, there are often situations where the localized orbitals used to form the density matrix are not orthogonal. This leads to some complications in the formalism, which are described in section 4.1. Briefly, either the inverse of the overlap is required (which can be difficult to find for sparse matrices) or some form of orthogonalization must be applied. If the density matrix for only the *occupied* subspace is required, then it is equal to the inverse overlap matrix of the local orbitals (see, for instance, [58, 163, 165, 166]).

3.2.1. Direct and iterative approaches. In this section, we consider two major approaches to linear-scaling density matrix search, which turn out to share considerable theoretical background. We choose to group these methods to emphasize this shared background and stimulate further work on the development of effective techniques.

Galli and Parrinello [58] noted that, instead of writing the density matrix in terms of the occupied eigenfunctions (as seen, for example, in equation (10) above), it can equally well be written in terms of the same number of non-orthogonal orbitals ϕ_i:

$$\rho(\boldsymbol{r}, \boldsymbol{r}') = \sum_{ij} \phi_i^\star(\boldsymbol{r}') S_{ij}^{-1} \phi_j(\boldsymbol{r}). \tag{23}$$

Note that the matrix S_{ij}^{-1} is the inverse of the overlap matrix for the orbitals, $S_{ij} = \langle \phi_i | \phi_j \rangle$, which automatically makes the matrix $\rho(\boldsymbol{r}, \boldsymbol{r}')$ idempotent. They then impose localization constraints (in this case, using bucket-like potentials). The formulation allows the removal of any explicit orthogonalization between eigenfunctions (which leads to the asymptotic cubic scaling behaviour seen in normal methods). By taking advantage of the sparsity of the overlap and Hamiltonian matrices, linear scaling can also be achieved; the proposed methods for exploiting locality (involving local volumes where most operations are performed) have points of contact with both the divide-and-conquer method (section 3.2.2) and the FFT box method (see below under the ONETEP method in section 5.1). This method has been developed further to include unoccupied states, with a real-space grid as the basis (using finite differences, as described in section 2.1) [24], with an extension to allow the centres of the localization regions to adapt and move during MD [11, 37, 202]. These implementations are often not truly linear scaling, as the inverse overlap matrix is calculated exactly, though there are many methods to remove this final barrier; Galli and Parrinello proposed solving for the dual basis functions by an iterative application of $(I - S)$ (described in more detail in section 4.1).

A variational approach using the McWeeny purification transformation [121] has been used in a wide variety of approaches and methods [200, 201, 203–212]; we will refer to it as the density matrix minimization (DMM) method. The McWeeny transformation uses the function

$$f(x) = 3x^2 - 2x^3 \tag{24}$$

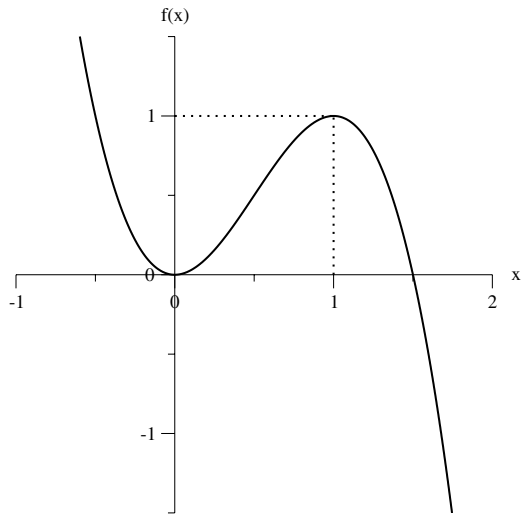

Figure 4. The McWeeny purification function, $f(x) = 3x^2 - 2x^3$.

which is plotted in figure 4. It has the property that for $-\frac{1}{2} \leqslant x \leqslant \frac{3}{2}, 0 \leqslant f(x) \leqslant 1$. This property is used to impose idempotency during a variational search: if an input matrix has eigenvalues bounded by $-\frac{1}{2}$ and $\frac{3}{2}$, then the output will be bounded by 0 and 1. This is known as *weak* idempotency. If an auxiliary density matrix is taken as $\sigma(\boldsymbol{r}, \boldsymbol{r}')$, then the true density matrix $\rho(\boldsymbol{r}, \boldsymbol{r}')$ is defined as

$$\rho = 3\sigma * \sigma - 2\sigma * \sigma * \sigma \tag{25}$$

where $\sigma * \sigma(\boldsymbol{r}, \boldsymbol{r}') = \int \mathrm{d}\boldsymbol{r}'' \sigma(\boldsymbol{r}, \boldsymbol{r}'')\sigma(\boldsymbol{r}'', \boldsymbol{r}')$. If the band energy (written as $2\mathrm{Tr}[\rho H]$) is varied with respect to the elements of the *auxiliary* density matrix, then the resulting minimization is variational and converges to idempotency at the ground state. It can easily be shown that

$$\frac{\delta E}{\delta \sigma} = 3(\sigma H + H\sigma) - 2(\sigma^2 H + \sigma H \sigma + H\sigma^2). \tag{26}$$

Provided some method is chosen to account for the chemical potential of the electrons (either adjusting it to keep the number of electrons fixed, or fixing the potential, as described in section 4.2), minimizing the total energy with respect to the auxiliary density matrix elements leads to the ground state in a variational manner. Imposing sparsity on the density matrix (using methods described in section 4.3) and the Hamiltonian results in linear scaling.

This approach was first proposed for orthogonal TB [201, 203] and subsequently extended to non-orthogonal bases [204] (discussed further in section 4.1) and DFT [200, 205] as well as finite electronic temperatures [213]; it is often referred to as the LNV method (after Li, Nunes and Vanderbilt). The utility of mixing an iterative McWeeny process to restore idempotency with the minimization has also been explored [208, 209, 214]. Implementations are numerous [200, 205–207, 211, 212, 215–217] and cover quantum chemistry approaches as well as DFT, and use minimization techniques such as conjugate gradients and direct inversion in the iterative subspace. Implementations of Car–Parrinello MD have also been described [218, 219], both for a simple fictitious electronic mass [218] and a variable fictitious mass to optimize convergence [219]. A closely related method to the DMM uses the sign matrix [220, 221], which is equivalent to expanding the Fermi matrix (as in the Fermi operator expansion (FOE) described in section 3.2.3); another approach including *implicit* purification for finite temperatures has been derived [222].

An alternative approach using similar ideas for purification starts by noting that the density matrix and the Hamiltonian should commute (they share eigenvectors, as the density matrix projects onto the occupied subspace). By starting with an initial density matrix which commutes with the Hamiltonian, and using iterative purification methods, it can be proved that the ground-state density matrix is the result [214]. This should not be surprising, as the density matrix can be written as a function of the Hamiltonian:

$$\rho = \theta(\mu \boldsymbol{I} - \boldsymbol{H}), \tag{27}$$

where μ is the Fermi level and $\theta(x)$ is the Heaviside step function (with $\theta(x) = 1, x > 0$ and $\theta(x) = 0, x < 1$). Approaches to expanding out the step function (or the Fermi function as it becomes at finite temperature) using polynomials and recursion are discussed in section 3.2.3.

The method first proposed [214] used the original McWeeny transform for grand canonical (i.e. fixed Fermi level) minimization, with the initial density matrix given by

$$\rho_0 = \frac{\lambda}{2}(\mu \boldsymbol{I} - \boldsymbol{H}) + \frac{1}{2}\boldsymbol{I} \tag{28}$$

$$\rho_{n+1} = 3\rho_n^2 - 2\rho_n^3 \tag{29}$$

with λ chosen to take the smaller value of $1/(H_{\max} - \mu)$ and $1/(\mu - H_{\min})$, so that the eigenvalues of ρ lie between 0 and 1. The same work also proposed a canonical method with an adapted version of McWeeny's purification function,

$$\rho_{n+1} = \begin{cases} \left[(1-2c_n)\rho_n + (1+c_n)\rho_n^2 - \rho_n^3\right]/(1-c_n) & \text{if } c_n \leqslant \frac{1}{2} \\ \left[(1+c_n)\rho_n^2 - \rho_n^3\right]/c_n & \text{if } c_n \geqslant \frac{1}{2} \end{cases} \tag{30}$$

$$c_n = \frac{\mathrm{Tr}[\rho_n^2 - \rho_n^3]}{\mathrm{Tr}[\rho_n - \rho_n^2]}. \tag{31}$$

This function allows the unstable fixed point of the McWeeny function at $x = c$ (where $f(x) = x$ and $f'(x) \geqslant 1$) to move away from $c = \frac{1}{2}$ to lie between $c = 0$ and $c = 1$. As a result, the electron number is conserved throughout the iteration. It is important to note [121, 214] that the grand canonical iteration is equivalent to a steepest descent minimization of the function $f(\rho) = \mathrm{tr}[\rho^2(I-\rho)^2]$, which links to the orbital minimization method (OMM) described below.

This basic idea has been further extended and elaborated in numerous ways. As noted above, this approach has been combined with variational minimization, both as an initialization for the density matrix [208] and to restore idempotency [209]. An iterative purification was introduced as a way of correcting density matrices following Car–Parrinello steps (rather than imposing orthogonality) [218]. A larger set of generic purifications was proposed [223], based on the equation $T_n(P) = I - (I - P)^n, n \geqslant 2$, and later

extended to systems with unoccupied states [224] (though it has been suggested that this extension is no more efficient than the original McWeeny transform [225]). Similar high-order polynomials have been derived elsewhere [49]; these methods can be formulated in terms of only a few matrix multiplications (e.g. four multiplies for a ninth-order polynomial), but are not in general use. The difference between methods which have the same number of filled orbitals as eigenstates, and therefore density matrix eigenvalues of 1 only, and those with more orbitals and hence eigenvalues of 0 and 1 is key in these expansions: an expansion which only has to consider filled states generally requires fewer matrix multiplications (and the form $2P - P^2$ proposed by Stechel [226] discussed below is a key example).

Building on this idea, Niklasson [227] suggested a trace-correcting approach, using different polynomial expansions for purification depending on whether the trace of the density matrix is below or above the correct number of electrons, N:

$$T_m(x) = \begin{cases} 1-(1-x)^m[1+mx], & N_e \geqslant N \\ x^m[1+m(1-x)], & N_e \leqslant N \end{cases} \tag{32}$$

If $m = 2$, then the original McWeeny expansion is recovered. When $m = 1$ the two polynomials are x^2 and $2x - x^2$; this has become known as the TC2 (trace-correcting second-order) method. It has been extended to spin-unrestricted methods [228] and non-orthogonal bases [229] (discussed more fully in section 4.1; the main change to the algorithm is to require the overlap matrix in density matrix products, so $P_\perp^2$ becomes PSP where $P_\perp$ is the density matrix in an orthogonal basis, and to require the inverse of the overlap for initialization), and compared with LNV [191], with some advantage found especially for high and low filling (though it is important to note that these iterative methods are not variational, and so forces cannot be calculated using the Hellmann–Feynman theorem).

Two closely related approaches use both the particle density matrix and the hole density matrix (defined as $Q = I - P$). Mazziotti [230, 231] recasts and combines the formulae of [224] in terms of the particle and hole matrices (with the method depending on the order—hole or particle first), and shows significant computational speed up compared with the McWeeny approach. A similar technique [232, 233] uses hole and particle density matrices, but includes an empirically adjusted parameter to optimize convergence.

Niklasson [234] also introduced a series of trace *resetting* algorithms. This method combines a purification polynomial, $F(x)$, and a reoccupation polynomial, $G(x)$, which between them purify the density matrix and keep its trace correct within a certain domain of applicability. If the trace falls outside the domain, then it is reset by application of the TC2 method. The following quartic polynomials have been empirically found to be effective:

$$F(x) = x^2(3x - 3x^2) \tag{33}$$

$$G(x) = x^2(1-x)^2, \tag{34}$$

leading to the TRS4 algorithm. Tests on both high- and low-filling problems (C_{60}, a zeolite, chlorophyll and water clusters) show that the approach is more efficient and more accurate than the original Palser–Manolopoulos method; at mid-filling the methods are of similar efficacy.

A comprehensive comparison of LNV-based minimization and iterative methods [192] finds considerable efficiency gains using purification, though it should be emphasized that these methods are not variational, and hence force calculations will be complicated. Studies of error propagation for the trace resetting method (TRS4) [234] with magnitude-based truncation (see section 4.3 for discussion of different methods of enforcing sparsity on matrices) as well as the TC2 method [235] show that, applying truncation at different stages, errors can be rigorously controlled. A method for controlling errors within the TC2 method has been proposed and demonstrated on water clusters [236]. These approaches show that it is possible to pursue high accuracy within linear-scaling methods.

Despite this plethora of methods, some fairly simple conclusions can be drawn. First, that McWeeny's original proposal has been remarkably robust, and has led to significant important physics. Second, that for variational methods the formulation of McWeeny's purification algorithm in terms of an auxiliary density matrix [201, 204] is the method of choice. It is probably the most commonly used method, and is ideal for distance-based truncation schemes. Third, for iterative approaches improvements over the original McWeeny formula can be made (e.g. the TC2 method) and error control can be introduced. The iterative methods are not variational, but are simple and are commonly used in conjunction with tolerance-based truncation.

Another class of methods, which is closely related to the minimization and iteration techniques just described, builds the Wannier functions of a system by direct minimization *without constraint*. The ideas were proposed independently by Mauri, Galli and Car [163] and Ordejón *et al* [165], though the method is generally referred to as MGC; given that the heart of the method is orbital minimization, we suggest that the method be known as the orbital minimization method (OMM), following [237]. In both cases, the density matrix is defined in terms of $N/2$ localized Wannier functions (LWFs) which tend to orthogonality as the minimization progresses. By introducing an approximation to the inverse overlap (which coincides with the density matrix if only occupied states are considered), we can write [163]

$$\rho(\boldsymbol{r}) = \sum_{ij} \phi_i(\boldsymbol{r}) Q_{ij} \phi_j(\boldsymbol{r}) \tag{35}$$

$$Q = \sum_{0}^{K} (I - S)^K \tag{36}$$

where K is an odd integer. When $K = 1$, then the approximation is $Q = 2I - S$. This same formula was derived by Ordejón *et al* [165] starting from a Lagrange multiplier approach to enforce orthogonality (adding a term $\sum_{ij} \Lambda_{ij}(S_{ij} - \delta_{ij})$ to the band energy) and substituting the value of the multipliers *at the minimum* ($\Lambda_{ij} = H_{ij}$) for all values of the orbitals. After a slight re-arrangement, this yields the same functional; the family of polynomials can also be derived

by rearranging the original equation to $\Lambda = H + (I - S)\Lambda$ and treating it as a recurrence relation [166]. (The Ordejón *et al* approach has some commonality with a method due to Wang and Teter described in section 3.2.4, though is more flexible.) It can be proved [163, 165, 166, 238] that the energy is a minimum at orthogonality provided that the Hamiltonian is negative definite (which can be enforced by applying a rigid shift to the Hamiltonian of an amount η). The functional often quoted for these methods is then

$$E[Q, \{\phi\}] = 2\left(\sum_{ij}^{N/2} Q_{ij}T_{ij} + F[\rho]\right) + \eta\left(N - \int \mathrm{d}\boldsymbol{r}\rho(\boldsymbol{r})\right), \tag{37}$$

where ρ is defined in equation (35) and Q is defined in equation (36). Given a suitable basis set (e.g. localized atomic orbitals, as described in section 2.3), the minimization expresses the orbitals $\{\phi_i\}$ in terms of this basis, and seeks the minimum energy in terms of the expansion coefficients, while also applying a localization criterion to the orbitals. At the minimum, the resulting orbitals will be orthogonal, by construction.

However, as described, the method has a serious problem: there are large numbers of local minima, leading to severe convergence problems [164, 166, 238]. While the method can be used, and has been implemented in this form, practical solutions require some input guess for the orbitals based on chemical intuition; a more general solution to this problem is required. Kim *et al* [164] showed that generalizing the original formulation so that more orbitals than just those spanning the occupied subspace could be used. This means that the orbitals are not orthogonal at the energy minimum, but has the advantage that local minima are avoided. Hierse and Stechel [199] also generalized the original OMM approach, using a different approximation for the trial density matrix, using equation (18). Yang [239] used a variational approach to derive a general functional of this class which can use unoccupied orbitals and which only requires a Hermitian matrix (as opposed to a positive definite matrix required in previous work [199]). It is interesting to note that Kohn suggested a method for building orthogonal Wannier functions [124] which may be seen as a precursor to these methods; he noted that, unless the starting functions were reasonably close to the final functions, there were multiple minima.

There are numerous implementations of these methods: the original papers [163, 165, 166, 238]; the generalized versions [164, 199, 240, 241]; in parallel [242, 243]; with ultra-soft pseudopotentials [36]. A real-space implementation of similar ideas [202] uses exact inversion of the overlap, thus not forming a strict linear-scaling method (though the cubic scaling part will have a small prefactor). A method for projecting localized functions onto the occupied subspace [244] using the FOE described in section 3.2.3 showed improvements in convergence, but does not solve the problem of initial functions.

More recently, Tsuchida [44, 237] has proposed an *augmented* OMM to overcome the convergence problems. The essence of the method is to define highly localized kernel functions (which contain the centre of a localization region, and are typically around $1a_0$ in radius); the localized orbitals are then forced to be orthogonal to the kernel functions. This change is sufficient to make the method stable and rapidly convergent. Indeed, the applications described in section 5.2 suggest that this is one of the most successful linear-scaling approaches.

At one level, the density matrix based methods appear quite different from the OMMs: in the first case, the elements of the density matrix are the variables, while in the second the orbitals themselves are the variables. However, we can make a close connection to the methods described above, by considering alternative ways of writing the density matrix; these demonstrations have been developed by a number of authors [167, 193, 194, 205]. If the OMM functional Q is written in terms of density matrices (compare equation (35)), then we find

$$\tilde{\rho}(\boldsymbol{r}, \boldsymbol{r}') = 2\rho(\boldsymbol{r}, \boldsymbol{r}') - \rho^2(\boldsymbol{r}, \boldsymbol{r}') \tag{38}$$

where $\rho(\boldsymbol{r}, \boldsymbol{r}') = \sum_i \chi_i(\boldsymbol{r})\chi_i(\boldsymbol{r}')$ is a trial density matrix, and $2\rho - \rho^2$ acts as a purification transformation. This can be generalized if the localized orbitals $\chi_i(\boldsymbol{r})$ are expanded in a basis:

$$\chi_i(\boldsymbol{r}) = \sum_\alpha b_{i\alpha}\phi_\alpha(\boldsymbol{r}). \tag{39}$$

We can then write the trial density matrix in exactly the form of equation (18), with $L_{\alpha\beta} = \sum_i b_{i\alpha}b_{i\beta}$. The purification transformation is then written $2L - LSL$, with $S_{\alpha\beta} = \langle\phi_\alpha|\phi_\beta\rangle$. Setting $L = I$ gives S as the density matrix, and recovers the original OMM form.

Returning to equation (38), we note that ρ is required to be positive semidefinite, and that the eigenvalues can thus be represented by [205]

$$\lambda_\rho = \kappa_\rho^2 \tag{40}$$

$$\lambda_{\tilde{\rho}} = \lambda_\rho(2 - \lambda_\rho) = \kappa_\rho^2(2 - \kappa_\rho^2) \tag{41}$$

where κ_ρ is real. The quartic function will lie in the range $[0, 1]$ when $-\sqrt{2} \leqslant \kappa_\rho \leqslant \sqrt{2}$, with turning points at $\lambda_{\tilde{\rho}} = 0$ and 1. This function is plotted along with the original McWeeny function in figure 5. It is clear that the two methods have a very similar form between zero and one, but the quartic potentially will introduce more local minima. The practical convergence rates for the methods (distinguishing between the MGC and Ordejón *et al* methods) have been examined [190], with the convergence rate in the simple systems (bulk Si and C) found to be dominated by the spectral properties of the Hamiltonian rather than the method.

The DMM and OMM methods share a fundamental connection, but the DMM is far more commonly implemented. This is in part due to simplicity and stability, and also to the single minimum present in the functional. The augmented OMM promises significant improvements and suggests that a new resurgence of OMM applications might be seen; either way, these two methods are the most commonly used and applied linear-scaling methods (see section 5.1 for a description of implementations in specific codes).

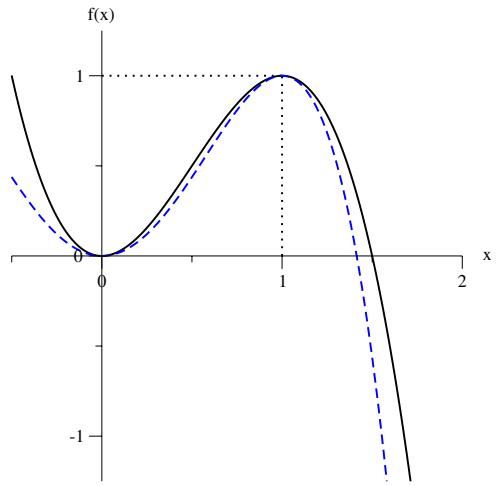

Figure 5. The McWeeny purification function (solid black line) plotted with the implicit purification function of the OMMs (dashed blue line). (Note that x is taken to mean κ_ρ in the latter case.)

3.2.2. Divide and conquer. The divide-and-conquer method [245–247] is conceptually the simplest of the linear-scaling approaches. Taking advantage of the nearsightedness of electronic structure, the system is divided into separate, small subsystems whose electronic structure is solved exactly, for example by diagonalization. (A similar method was proposed almost simultaneously [248], which combines the idea of dividing up the cell into subsystems with the ideas of orbital-free DFT (see section 3.2.5); however, this method has not proved significant.) A partition function p_α is defined for each subsystem, such that $\sum_\alpha p_\alpha(\boldsymbol{r}) = 1$ for all points in the system. If we rewrite the charge density, then we have

$$n(\boldsymbol{r}) = 2\langle \boldsymbol{r}|\theta(\epsilon_\mathrm{F} - \hat{H})|\boldsymbol{r}\rangle \tag{42}$$

$$= 2\sum_\alpha p_\alpha(\boldsymbol{r})\langle \boldsymbol{r}|\theta(\epsilon_\mathrm{F} - \hat{H})|\boldsymbol{r}\rangle = \sum_\alpha n_\alpha(\boldsymbol{r}) \tag{43}$$

$$n_\alpha(\boldsymbol{r}) = 2p_\alpha(\boldsymbol{r})\langle \boldsymbol{r}|\theta(\epsilon_\mathrm{F} - \hat{H})|\boldsymbol{r}\rangle \tag{44}$$

where $\theta(x)$ is the Heaviside step function. However, this requires some approximation of the system-wide KS Hamiltonian. Making a *local* approximation to the Hamiltonian requires some way to determine ϵ_F; Yang suggested writing $\tilde{n}_\alpha(\boldsymbol{r}) = n_\alpha(\boldsymbol{r})$, with

$$\tilde{n}_\alpha(\boldsymbol{r}) = 2p_\alpha(\boldsymbol{r})\langle \boldsymbol{r}|f(\epsilon_\mathrm{F} - \hat{H}_\alpha)|\boldsymbol{r}\rangle \tag{45}$$

where $\hat{H}_\alpha$ is the approximation of the KS Hamiltonian in the subsystem. Yang assumed non-orthogonal basis functions purely localized within the subsystem, $\{\phi_j^\alpha(\boldsymbol{r})\}$ and defined:

$$\hat{H}_\alpha = \sum_{ij} |\phi_i^\alpha\rangle H_\alpha^{ij} \langle\phi_j^\alpha| \tag{46}$$

where $H_\alpha^{ij} = \sum_{kl}(S_\alpha^{-1})^{ik}(H_\alpha)_{kl}(S_\alpha^{-1})^{lj}$, S^{-1} is an inverse matrix for the overlap $(S_\alpha)_{ij} = \langle\phi_i^\alpha|\phi_j^\alpha\rangle$ and $(H_\alpha)_{ij} = \langle\phi_i^\alpha|\hat{H}|\phi_j^\alpha\rangle$. The only system-wide constraint is the electron number, which gives a value for ϵ_F. A finite temperature is used to ensure that the electron number is monotonic and continuous as ϵ_F is varied. The approximate band energy is given by

$$\tilde{E} = 2\sum_\alpha \sum_i f(\epsilon_\mathrm{F} - \epsilon_i^\alpha)\epsilon_i^\alpha \langle\psi_i^\alpha|p^\alpha|\psi_i^\alpha\rangle. \tag{47}$$

Once a partitioning of the system, and a local basis set for each subsystem, has been chosen, the solution for each is found by direct diagonalization of H_α^{ij} and $(S_\alpha)_{ij}$. The system-wide Fermi level is fixed by the requirement on electron number, which is then used to construct a charge density, and fed back into the local Hamiltonians. A self-consistent solution is thus found for the entire system.

The formalism has been extended to solid-state systems [249]: here, the idea of buffer atoms is required. For each subsystem, a buffer zone of a certain width from the main system is retained to remove surface effects. Comparing the error as a function of the size of the buffer, the authors conclude that cohesive energy is converged to 0.1 eV when there are 40–50 atoms in buffer, but the electronic structure (in particular the density of states, or DOS) is much slower to converge with buffer size. Forces have been derived for the method [250], though the choice of applying the divide-and-conquer approach to the total system force (rather than differentiating the divide-and-conquer energy) may lead to slight discrepancies between the force and energy gradients. A more recent study of forces [251] showed that it is possible, though time-consuming, to derive an exact gradient, and proposed another approximate solution.

A further refinement based the partitioning and calculations on the system density matrix rather than the wavefunctions [252]: the charge density is built from the density matrix. A Mulliken-like partitioning is used for assigning the density matrix to subsystems; the partition *matrix* is defined as

$$p_{ij}^\alpha = \begin{cases} 1 & \text{if } i, j \in \alpha \\ 1/2 & \text{if } i \text{ or } j \in \alpha \\ 0 & \text{if } i, j \notin \alpha \end{cases} \tag{48}$$

and then we have

$$K_{ij} = \sum_\alpha p_{ij}^\alpha K_{ij} = \sum_\alpha K_{ij}^\alpha \tag{49}$$

$$K_{ij}^\alpha = 2p_{ij}^\alpha \sum f(\epsilon_\mathrm{F} - \epsilon_m^\alpha) C_{im}^\alpha C_{jm}^\alpha. \tag{50}$$

This method has the advantage that it removes the need to perform integrals between subsystem eigenstates and the projector function. An alternative partitioning [253] based on the number of subsystems a density matrix occupies has also been proposed.

The divide-and-conquer method has been implemented a number of times, for instance in the OpenMX code [254] (in this case in combination with a recursion method, as described in section 3.2.3), in the SIESTA code [255] and in an extremely large-scale approach [256–258] based on hierarchical real-space grids, which scales over multiple HPC centres. An implementation using the QUAMBO approach to localized orbitals also exists [259]. The divide-and-conquer approach has also been applied to quantum chemistry (up to the CCSD level) [260, 261], and implementation of exact exchange interactions has been described [262] (where the accuracy is very dependent on the size of the buffer region). An assessment of the computational time required for different systems [179] with divide and conquer has been made, and relates the time

needed to the required accuracy, as well as physical attributes of the system.

The main approximation in the divide-and-conquer method is in projecting the KS method onto small areas; this means that there is no variational principle, and the exact result will depend on the choice of partitioning unless extreme care is taken. Moreover, the convergence with subsystem size can be very slow, so that good convergence is hard. Formally exact partition theory [263, 264] suggests a route forward in putting these approaches on a more exact footing. The key idea is to transform a set of interacting fragments of a molecule or other system into a set of *non-interacting* fragments in an effective potential (in exact parallel of KS theory). However, as with KS theory, the approach does not have an analytic functional giving the effective potential; intriguingly, the obvious local approximations are strongly related to orbital-free DFT (discussed in section 3.2.5).

There are a number of methods related to divide and conquer. The 3D fragments [265, 266] builds on an earlier charge patching method [267, 268], and gives a good route to modelling large semiconductor systems, though it requires passivation of surfaces of the fragments. The key idea is that each fragment is part of several differently sized fragments. For simplicity, consider two dimensions and square fragments: then each fragment forms one small area (1×1), and is part of four rectangular areas (2×1 and 1×2) and four large square areas (all 2×2). These are combined (with differing signs) to yield the total energy. The method as formulated is variational, making the forces easy to calculate. The Mosaico method [269] combines aspects of divide-and-conquer and a Wannier function method (the method seeks local molecular orbitals within specific areas); it has been implemented and applied within SIESTA [270].

The fragment molecular orbital (FMO) method [271] uses a similar idea to the divide-and-conquer method (dividing a system up into small pieces) and applies it to create an approximate method for proteins and related biological molecules. The total energy of the whole molecule is calculated from the energy of fragments and pairs of fragments without solving the molecular orbitals of the whole molecule. In the calculation of the molecular orbitals of fragments, they introduce a special technique for treating the bonds at the boundary of the fragments. The method has found a wide variety of applications [272]. There is a closely related method (the molecular tailoring approach [273] which introduces a different method for defining the fragments. There are also proposals to use the result of an FMO-like calculation as the input density for a conventional or linear-scaling method [274].

Two recently proposed, related methods [275, 276] are related to the main idea of divide and conquer, but can yield exact results (though not in linear-scaling time). Domain decomposition in a quasi-1D system (i.e. one with two dimensions much smaller than the other) shows linear scaling and good efficiency [275], though the generalization to 3D is not given. Once the system is subdivided, a block-diagonal Hamiltonian can be built and either the KS eigenstates in each block can be used as a basis for creating a structured H or an LDL factorization can be used to generate the Green's function. A similar but more general method [276] uses contour integration of the Green's function to get the density matrix; after re-ordering the Hamiltonian to create a structured matrix, LDL^T factorization is used recursively to find the Green's function. The method, though not explicitly linear scaling, is efficient: it is $N(\log_2 N)^2$ in 1D, N^2 in 2D and $N^{7/3}$ in 3D, and is exact.

Other related approaches have been proposed for metallic systems. A KKR method where the system is divided into local interaction zones (LIZ) [277] then solves for the Green's function within the zone using multiple scattering theory. Improvements include embedding in a local environment [278, 279]. This contrasts with the screening in k-space found to lead to linear scaling in the number of layers for multi-layers and surfaces [280, 281]. Recent work suggests that combining the two approaches, so that real-space screened KKR is used for large imaginary energies and k-space screened KKR for small imaginary energies, reduces scaling to $\mathcal{O}(N^{1+\epsilon})$, where $\epsilon < 0.2$ [282, 283]. The method can be made truly linear scaling by using iterative minimization [284]; screened full-potential KKR has also been developed [285].

The simplicity of this type of approach makes it extremely attractive. It is easy to implement and has an obvious relation to the shortsightedness of electronic structure. However, the rate at which convergence to the exact result is achieved relative to subsystem size and the lack of a variational principle make these methods unsuitable for quantitative calculations with full DFT accuracy.

3.2.3. Recursive and stochastic approaches. The electronic structure of a system can be evaluated from the DOS as well as from the eigenvalues; for instance, the band energy can be written:

$$E_{\mathrm{band}} = \int_{-\infty}^{E_{\mathrm{f}}} \mathrm{d}E n(E) E \tag{51}$$

where $n(E)$ is the DOS and E_{f} is the Fermi level (practically, the lower bound on the integral is normally taken to be the bottom of the band of occupied states). The potential to use this in linear-scaling methods becomes clear when the DOS is written as a sum over *local* densities of states (LDOS). If a localized basis set such as PAOs is used, with $\{\phi_{i\alpha}(\boldsymbol{r})\}$, then we can write:

$$n(E) = \sum_{i\alpha} n_{i\alpha}(E). \tag{52}$$

As with any function, the LDOS can be described by its moments [286] (for instance: its mean, width, skew, etc, which correspond to the first, second, third, etc moments). But in a local orbital basis, these moments can be related to powers of the Hamiltonian; the pth moment of the LDOS $n_{i\alpha}$, $\mu_{i\alpha}^{(p)}$, is given by

$$\mu^{(p)} = \int \mathrm{d}E E^p n_{i\alpha}(E) = \langle i\alpha|\hat{H}^p|i\alpha\rangle \tag{53}$$

$$= \sum_{j_1\beta_1,\ldots,j_{p-1}\beta_{p-1}} H_{i\alpha j_1\beta_1} H_{j_1\beta_1 j_2\beta_2} \ldots H_{j_{p-1}\beta_{p-1} i\alpha} \tag{54}$$

where $H_{i\alpha j\beta} = \langle i\alpha|\hat{H}|j\beta\rangle$. In a simple TB picture, this corresponds to hopping around a lattice following closed loops of length p.

However, this is, in general, a rather unstable way of building the DOS for a system. One stable technique is the recursion method [287], which is a Green's function method. The LDOS is written as

$$n_{i\alpha}(E) = -\frac{1}{\pi}\lim_{\eta\to 0}\mathrm{Im}\left\{G_{i\alpha,i\alpha}(E+i\eta)\right\}. \tag{55}$$

This Green's function can be written in terms of a continued fraction, whose components are related to the elements of the tridiagonalized Hamiltonian of the system [287]. The element $G_{00}(Z)$ (where $G_{nm}(Z) = \langle U_n|\hat{G}(Z)|U_m\rangle$) can be found from:

$$G_{00}(Z) = \cfrac{1}{Z - a_0 - \cfrac{b_1^2}{Z-a_1-\cfrac{b_2^2}{Z-a_2-\cfrac{b_3^2}{\ddots}}}}, \tag{56}$$

with a_n the diagonal element and b_n the off-diagonal element. But most Hamiltonians are not tridiagonal.

The Lanczos recursion algorithm [288] is an efficient scheme for tridiagonalizing a matrix. Consider a matrix $\boldsymbol{H}$, which corresponds to the operator $\hat{H}$. Let the tridiagonal matrix be $\boldsymbol{H}'$, whose diagonal elements are given by a_n and whose off-diagonal elements are given by b_n. If the states which tridiagonalize the matrix are $|U_n\rangle$, then

$$H'_{mn} = \langle U_m|\hat{H}|U_n\rangle = \begin{cases} a_n, & \text{if } m = n, \\ b_n, & \text{if } m = n-1, \\ b_{n+1}, & \text{if } m = n+1, \\ 0, & \text{otherwise} \end{cases} \tag{57}$$

and the condition that the tridiagonalizing states are orthonormal ($\langle U_n|U_m\rangle = \delta_{n,m}$). The method can be extended to non-orthogonal basis vectors with a bi-orthogonal approach [289].

The overall procedure in a recursion method is therefore: choose an initial starting state; apply the Hamiltonian repeatedly to this (generating a Krylov subspace, which will be discussed in detail below); from the resulting tridiagonal Hamiltonian, construct the Green's function, DOS and density matrix as required. The application of recursion to TB has been described in detail elsewhere [216, 287], though it was used as the basis of one of the first DFT-based linear-scaling methods [290], with application to silicon using a finite difference approach.

The recursion method as described has a major drawback: the recursion for diagonal elements of the Green's function, which are required for the energy, are stable and easy to evaluate, but the *off-diagonal* elements, which are required for force calculations, are hard to evaluate with a tendency to numerical instability. The early techniques used to work around these problems (using the difference between bonding and anti-bonding orbitals for inter-site elements, and matrix recursion) did not adequately resolve convergence problems. The bond order potential (BOP) method [291–293] is a mathematically complex solution to this problem, which involves redefining the procedure in terms of a new basis, using sum rules for the Green's function. A block BOP [294, 295] uses a simpler basis to provide an efficient route to energies and forces via recursion. In all these cases, by limiting the range of the recursion to a cluster of specified size, and truncating the moments considered, the method becomes linear scaling. The block BOP method has been demonstrated with DFT [296], using a dual basis approach. However, BOP methods are more commonly used to bridge between TB and empirical methods [297]. Ozaki [254] has shown how the divide-and-conquer method (see section 3.2.2) can be combined with recursion to create a stable, easily extensible linear-scaling method. This will be discussed further below.

A large class of methods use Chebyshev polynomials, which can be defined recursively:

$$T_{j+1}(H) = 2HT_j(H) - T_{j-1}(H), \tag{58}$$

with $T_0(H) = I$ and $T_1(H) = H$. The FOE method [120, 298, 299] writes the finite temperature density matrix (or Fermi matrix) as

$$F_{\mu,T} = f\left(\frac{H-\mu I}{k_\mathrm{B}T}\right) \tag{59}$$

where $f(x) = 1/(1+\exp(-x/k_\mathrm{B}T))$ is the Fermi function, and then expands the Fermi function as a Chebyshev polynomial. Each column of the matrix can be found by a recursive procedure, starting with a localized orbital, and applying the Hamiltonian repeatedly. Without any truncation, this yields an $\mathcal{O}(N^2)$ method (which has been used for plane-wave DFT [300]); with truncation of the elements in a column of the Fermi matrix, the method scales linearly with the atoms in the system. It can be shown that the polynomial expansion should have a degree of the order of $n \simeq (\epsilon_\mathrm{max} - \epsilon_\mathrm{min})k_\mathrm{B}T$, to represent the DOS adequately. The early applications were to TB, but the method can be extended to DFT [299]. The forces as originally calculated [298] are not exact derivatives of the energy; an exact expression for the forces has been derived [301, 302] but it involves significant computational cost.

There are a number of methods which are closely related to FOE. The kernel polynomial method [301–303] differs only in the way that the DOS is reconstructed: Gibbs factors are used to weight each term in the polynomial, reducing oscillations which result from the truncation of the polynomial at a finite level. (Note that this is related to the use of a terminator in a continued fraction in the recursion method.) Liang *et al* [304] give various improvements to the original method: they demonstrate fast resumming of polynomials; they explore different approximations to the Fermi function, concluding that the complementary error function is the best to use; and they show how the inverse temperature, β, can be related to the accuracy required and the system properties (in particular band width and gap). The use of Chebyshev polynomials to select part of the spectrum has been extended to $\mathcal{O}(N)$ methods [305], where the key scaling issue is inversion of the overlap matrix (which is performed iteratively; overlap matrices are discussed in greater detail in section 4.1). A method using maximum entropy techniques to extract the moments of DOS by applying Chebyshev polynomials of

the Hamiltonian to random vectors [306] uses importance sampling to select vectors which give an idempotent density matrix. This idea was extended [307] to project the random vectors on the occupied subspace as a starting point for the recursion method. Wang [308] similarly defines moments of the DOS in terms of Chebyshev polynomials, and then calculates these using random wavefunctions and maximum entropy methods with a plane-wave basis set to generate the DOS and optical absorption spectra.

Random vectors have also been used in a stochastic approach to invert the exponential of a Hamiltonian [309]. Noting that the correct thermodynamic potential for spinless fermions is given by

$$\Omega = -\frac{1}{\beta}\log\det\left(I + \mathrm{e}^{\beta(\mu - H)}\right) \tag{60}$$

the exponential is then rewritten as

$$\Omega = -\frac{1}{\beta}\sum_{l=1} P\log\det\left(M(l)\right) \tag{61}$$

where $M(l) = 1 + \exp(\mathrm{i}\pi(2l-1)/P)\exp(\beta/P)(\mu - H)$, with the exponential written as a Chebyshev polynomial or a Trotter expansion. It can be shown [309] that expectation values can be written in terms of a sum over the inverses of $M(l)$, so that the main computational effort is in inverting these sparse matrices. The inverse is found by selecting a random vector ψ, solving the linear equation $M(l)\phi = \psi$ and calculating $M(l)^{-1} = \langle\phi\psi^{\dagger}\rangle$, with the expectation value taken over the stochastic process. The original application was to a very simple system, but it was developed to TB applications in 1D systems [310, 311], using Langevin dynamics and noisy forces. The decomposition of the grand potential has innate scaling $\mathcal{O}(N)^{3-2/d}$ for a d-dimensional system. However, by rewriting the partition function in terms of both electrons and ions, and using a careful sampling of the Boltzmann distribution an efficient $\mathcal{O}(N)$ method can be developed which is valid for metals [312], though only demonstrated so far on TB systems (in this case, metallic carbon nanotubes (CNTs)). The method has been further developed with an exploration of efficient decompositions of the Fermi operator [313].

The energy renormalization group (ERG) method is an approach which takes a different view of the density matrix. Conventionally, the density matrix is viewed as the zero-temperature limit of the Fermi matrix defined above, and as the temperature is lowered, the correlations become longer ranged (particularly in small gap systems and metals) so that the matrix becomes non-sparse. The ERG method instead writes the density matrix as a telescopic sum of terms, with the temperature in each term decreasing by a factor $q > 1$:

$$\hat{\rho} = \hat{F}_{\beta_0} + (\hat{F}_{\beta_1} - \hat{F}_{\beta_0}) + (\hat{F}_{\beta_2} - \hat{F}_{\beta_1}) + \cdots \tag{62}$$

$$= \hat{F}_{\beta_0} + \hat{\Delta}_1 + \hat{\Delta}_2 + \cdots. \tag{63}$$

The first term in the series is a high-temperature Fermi matrix which is short-ranged. Each successive term gradually corrects for lower temperatures, so that $\beta_n = q^n\beta_0$. The Fermi matrix is written as a Chebyshev polynomial of the Hamiltonian, just as in the FOE method described above. The expansion is substituted into expressions for expectation values (rather than calculating a long-range density matrix); a recursive approach is given to coarse-grain the Hamiltonian in each successive space, using the Chebyshev polynomials. It is developed elsewhere [314], but has not shown true linear scaling beyond 1D systems and has not been applied beyond TB implementations. Recursive bisection of the density matrix [315] gives a method which scales linearly for 1D systems without truncation, but is highly efficient; given density matrix truncation, it would scale linearly with a small prefactor. A recursive bisection of real space [316] allows the truncation of KS wavefunctions to be controlled, and suggests a possible route to localized orbitals without the need to specify centres and extents *a priori*.

The essence of a subspace method is the repeated application of the Hamiltonian operator to a vector to form a new set of vectors, which span a space; this is often referred to as a Krylov subspace method, with the Krylov subspace spanned by the set $\{|\phi\rangle, H|\phi\rangle, H^2|\phi\rangle, H^3|\phi\rangle, \ldots, H^{n-1}|\phi\rangle\}$. For instance, a Lanczos procedure generates basis vectors in a Krylov subspace which are orthogonal. It has been proposed [317] that diagonalizing within the subspace (similar, in effect, to the method of Ozaki [254] described above) will give an efficient linear-scaling method; they find that around 30 vectors is sufficient. A more efficient variant of this method [318] solves linear equations:

$$(z - H)|x_j\rangle = |j\rangle \tag{64}$$

for given basis vectors $|j\rangle$ to generate vectors $|x_j\rangle$, from which the Green's function elements $G_{ij} = \langle i|x_j\rangle$ can be found; the basis vectors are the Krylov subspace vectors. The technique used is the conjugate–orthogonal conjugate gradient method, which generates a set of residual vectors *independent* of the energy shift z. Once the vectors $|x_j\rangle$ have been generated for one energy, further energies are almost trivial, and certainly scale as $\mathcal{O}(1)$. An overview and summary of applications using TB have been given [319, 320], and the method has been extended to non-orthogonal orbitals [321].

The recursion methods were the first set of linear-scaling methods proposed, and Lanczos approaches are widely used in many areas of physics and mathematics. The most successful *ab initio* recursion methods are based around FOE-like methods, mainly because of the simplicity of Chebyshev polynomials and the overall approach. The convergence and general lack of variational properties have made these methods less successful than other approaches.

3.2.4. Penalty functionals. An alternative approach to imposing idempotency was suggested by Kohn [175]: add a term to the energy functional which penalizes non-idempotent density matrices. This way, as the ground state is sought, idempotency is automatically included. The original method defines Hermitian trial functions $\tilde{n}(\boldsymbol{r}, \boldsymbol{r}')$, in terms of which the density matrix $\bar{n}$ is written as $\bar{n} = \tilde{n}^2$. The ground-state search is written in terms of a functional:

$$Q\left[\tilde{n}(\boldsymbol{r}, \boldsymbol{r}')\right] = E\left[\bar{n}\right] - \mu N\left[\bar{n}\right] + \alpha P\left[\tilde{n}\right] \tag{65}$$

where $E[\tilde{n}]$ is the usual KS energy functional, μ is a chemical potential, $N[\tilde{n}]$ is the number of electrons and $P[\tilde{n}]$ is a penalty functional,

$$P\left[\tilde{n}\right] = \left[\int \mathrm{d}\boldsymbol{r}\tilde{n}^2\left(1-\tilde{n}\right)^2\right]^{1/2}. \tag{66}$$

For idempotent $\tilde{n}$, $P[\tilde{n}] = 0$, so that the penalty functional has no effect. The number of electrons is set by choosing μ, leaving only α as a parameter. It can be shown [175] that, for a given μ and α larger than a critical value α_{C}, the correct, idempotent ground-state density matrix is found. However, α_{C} cannot be predicted exactly; too small a value of α will yield local minima while too large a value will slow convergence.

A lower bound on α_{C} can be derived [322] for a slightly different functional ($Q[\tilde{n}(\boldsymbol{r},\boldsymbol{r}')] = E[\tilde{n}] - \mu N[\tilde{n}] + \alpha P[\tilde{n}]$ where the trial function is used as the density matrix):

$$\alpha_{\mathrm{C}} > 2\max_i |\epsilon_i - \mu|. \tag{67}$$

However, it was also shown [322] that, due to the square-root form of the penalty functional, there is a branch point at the minimum which prevents standard minimization approaches such as conjugate gradients from being effective. A corrected functional was proposed [323] which removes these problems:

$$Q\left[\tilde{n}\right] = E\left[\tilde{n}\right] - \mu N\left[\tilde{n}\right] + \alpha P\left[\tilde{n}\right]^2 \tag{68}$$

where $\tilde{n}$ is now taken as the density matrix throughout. This approach does not impose idempotency exactly (nor weak idempotency), with occupancy errors which can be written as $\delta f_i = -(\epsilon_i - \lambda)/\alpha$ where ϵ_i is a KS eigenvalue at the minimum; the occupancies are such that occupied bands have more than one electron and unoccupied bands have negative occupancies, which also gives an α-dependent error in the total energy. A correction to the energy can be applied as $\mathrm{Tr}[\rho(1-\rho)^2(1-2\rho)]$ for occupied bands, which can be evaluated in $\mathcal{O}(N)$ steps; the unoccupied bands require a more complex correction. Following correction, the method gives a total energy independent of α. An alternative approach uses this functional as the heart of an iterative method via an augmented Lagrangian method [324].

The penalty functional method is used within the ONETEP code [212], as part of a cascading sequence of methods (from canonical purification, through penalty functional and finally LNV). We note that there was an early proposal which used a penalty functional with linear scaling [325]; the method sought highly LWFs in a basis of atomic TB orbitals without explicit orthogonalization, but a penalty applied for non-orthonormality (using a simple sum over all atoms i and their neighbours j of the form $\lambda\sum_i\sum_j|\langle\psi_i|\psi_j\rangle|^2$). Overall, penalty functional methods have not been widely taken up.

3.2.5. Orbital-free DFT. Orbital-free DFT [326, 327] returns to the original spirit of DFT, and seeks only the ground-state charge density and not the associated orbitals, giving a huge advantage in terms of simplicity and speed. However, to find the energy and hence the ground-state density, a functional for the energy in terms of the charge density is required; this follows standard DFT methods, except for the KE, and much of the difficulty in orbital-free DFT lies in finding a KE functional for the charge density. An overview of different approaches can be found in a recent review paper [328].

For a uniform electron gas, the Thomas–Fermi functional gives

$$T_{\mathrm{TF}} = c_k\int \mathrm{d}\boldsymbol{r}n(\boldsymbol{r})^{5/3} \tag{69}$$

where $c_k = \frac{3}{10}(3\pi^2)^{2/3}$. However, this is not sufficient for systems where the charge may vary. Rapidly varying perturbations can be represented by the von Weizsäcker functional:

$$T_{\mathrm{vW}} = \frac{1}{8}\int \mathrm{d}\boldsymbol{r}|\nabla n|^2/n. \tag{70}$$

These two functionals represent limits of behaviour, and early work on orbital-free DFT combined the two to form a single functional. While the functionals form the main topic of research, there is a further problem regarding pseudopotentials, which are in common use throughout physics. As the charge density is a local quantity (a function only of one position, $\boldsymbol{r}$) the pseudopotentials used must be local only. This introduces a significant restriction, as much of the transferability and accuracy of modern pseudopotentials comes from the angular freedom given by non-local potentials (where dependence on two positions $\boldsymbol{r}$ and $\boldsymbol{r}'$ allows different potentials to be used for different angular momenta). Progress has been made in creating local pseudopotentials suitable for bulk use [329], though this still causes problems, and restricts the transferability of the method. The most successful potentials are for metallic systems, and in particular those closest to the nearly free electron model.

Most of the work beyond this conceptual starting point is in creating new functionals [330–332], with extra freedom found by including density dependence in the kernel [333] and non-local functionals [334, 335]. An approach to allow the kernel to be calculated in non-periodic systems [336] has enabled development of methods which combine FE modelling and coarse graining (with the properties of each element based on a single representative atom modelled with OFDFT) [337]. Extension to covalent systems and semiconductors have been introduced [338, 339]; the key step forward was the use of a non-local KE functional, involving two parameters which are transferrable within environments with similar coordination numbers. There is a freely available implementation of orbital-free DFT, which has been demonstrated for a calculation on one million atoms of Al [340, 341]. Very recently, a new theoretical approach has been suggested [342] which writes the density as a functional of the potential (reversing the standard approach); this appears to give a well-defined route to KE functionals, though its impact on the field is yet to be seen.

While this family of methods gives a good way to model large metallic systems, there is still the concern that the KE functional is not exact, and the pseudopotentials are limited, despite significant effort. The recursion methods (section 3.2.3) are a viable alternative, though not widely used, and a detailed investigation of the relative accuracies of the different approaches, and their convergence, would be a valuable contribution.

3.2.6. Expansion of the density matrix and tensorial approaches. A recent class of approaches has emerged which use a change of variables (generally an exponential parametrization) to impose the idempotency of the density matrix, or the orthogonalization of the KS orbitals implicitly. These methods offer efficient routes to optimization of conventional methods, as well as potential for linear scaling.

The first proposal within linear scaling was an exponential parametrization of the density matrix [343], $\boldsymbol{D}$. We write

$$\boldsymbol{D}(\boldsymbol{X}) = \exp(-\boldsymbol{X}\boldsymbol{S})\boldsymbol{D}\exp(\boldsymbol{S}\boldsymbol{X}) \tag{71}$$

where $\boldsymbol{X}$ is an anti-symmetric matrix, and $\boldsymbol{S}$ is the overlap matrix between basis functions, as usual. This transformation preserves both the idempotency and trace of the density matrix, and gives a set of variables which allow a search for the ground state; note that the starting density matrix is key. The exponential can also be written in terms of a Baker–Campbell–Hausdorff (BCH) expansion:

$$D(X) = D + [D, X]_S + \tfrac{1}{2}[[D, X]_S, X]_S + \cdots, \tag{72}$$

$$[A, X]_S = ASX - XSA \tag{73}$$

where the commutator in a non-orthogonal basis is notated $[,]_S$. From this formalism, both Hartree–Fock and DFT can be made linear scaling, and an implementation has been given, with details on preconditioning of the minimization [344]. A similar idea has been independently derived from geometric considerations of the minimization [345], though without linear-scaling application. A related, orthogonal basis method using curvy steps has demonstrated linear scaling [346], and its generalization to non-orthogonal bases is discussed in detail below. The method has also been shown to be effective as an extended Lagrangian approach to Car–Parrinello MD [347]. This type of approach, using an exponential transformation as a unitary transform, has been used a number of times in the past in conventional methods (e.g. for Car–Parrinello MD [348]).

Another example of the unitary transformation leads to an approach to optimization of the density matrix with curvy steps [210]. Here, the unitary transformation is used to preserve idempotency and electron number, with the unitary operator written as an exponential, via a BCH expansion. Much care is devoted to tensorial correctness (as discussed in section 4.1). The transformed density matrix is written as a power series following the exponential expansion, which is again written in terms of an anti-Hermitian operator, $\hat{\Delta}$, with $\hat{U} = \mathrm{e}^{\hat{\Delta}}$. Using tensor notation (explained in section 4.1), they write

$$\tilde{P}^{\mu}_{\ \nu} = \mathrm{e}^{-\Delta^{\mu}_{\ \lambda}} P^{\lambda}_{\ \sigma} \mathrm{e}^{-\Delta^{\sigma}_{\ \nu}} \tag{74}$$

$$= \sum_{j=0} \frac{1}{j!} \left(P^{[j]}\right)^{\mu}_{\ \nu} \tag{75}$$

$$\left(P^{[j+1]}\right)^{\mu}_{\ \nu} = \left[P^{[j]}, \Delta\right]^{\mu}_{\ \nu} = \left(P^{[j]}\right)^{\mu}_{\ \alpha} \Delta^{\alpha}_{\ \nu} - \Delta^{\mu}_{\ \alpha} \left(P^{[j]}\right)^{\alpha}_{\ \nu}. \tag{76}$$

While the matrices in these equations are mixed, fully contravariant matrices, such as the density matrix, are simpler to work with, giving symmetric commutators. The gradient can be found as $S^{-1}HK - KHS^{-1}$ in the contravariant notation; taking a step in this direction and expanding out the energy in terms of the density matrix gives a polynomial search to some order, rather than the usual linear assumption. If the expansion for the expansion is truncated at linear order, then the gradients for the LNV method are recovered. As the expansion is truncated, idempotency is not exactly conserved, and must be re-applied using McWeeny iterations. The main advantage of the method is that it allows longer steps to be taken than would be possible with a linear method.

A method for the direct minimization of the wavefunction coefficients while remaining on the Grassmann manifold [349] has also been developed (the Grassmann manifold arises when the energy of the system depends only on the space spanned by the orbitals, and not on the orbitals themselves: the transformation from eigenstates to Wannier functions remains on the Grassmann manifold, for instance). The basis set coefficients are written as a matrix Φ, which is defined as a transformation of the bands Ψ. Then the overlap of the basis functions is incorporated into a covariant matrix $\bar{\Phi}^{\dagger} = (\Phi^{\dagger}\boldsymbol{S}\Phi)^{-1}\Phi^{\dagger}$, with the application of $(\Phi^{\dagger}S\Phi)^{-1}$ as a metric. Writing the energy as $E = \mathrm{Tr}[\bar{\Phi}^{\dagger}\boldsymbol{H}\Phi]$ and minimizing with respect to Φ gives a method which automatically includes the constraint of idempotency, and stays on the Grassmann manifold to first order. The inverse overlap matrix, S^{-1}, must be applied for tensorial correctness; the authors suggest finding this via inversion of a sub-matrix of S made only from the orbitals within the localization region [226]. The resulting method is closely related to Wannier function methods described above in section 3.2.1, particularly [226].

Another application of the parametrized unitary transform is an approach to sparsity [350]. The ℓ_1-norm (defined for a vector $\boldsymbol{x}$ as $\ell_1 = \sum_j |x_j|$, compared with the more usual $\ell_2 = \sqrt{\sum_j x_j^2}$) is used as a sparsity measure for the wavefunction coefficients. The key idea of the method is to perform unitary transformations on the orbital coefficients so as to maximize the sparsity of individual columns, using the ℓ_1-norm. The resulting method has a slight restriction, in that only gradients for steepest descent have been defined, but it shows promise, and is intended for linear scaling.

Unitary transformations and sparsity have also been used within the CP2K code [351, 352]. The normal constrained problem (with the constraint being orthogonality of eigenstates) is transformed to a locally unconstrained one; it is suggested that linear-scaling behaviour should result for sparse problems [352]. The method is based on an orbital transformation [351]: the wavefunction coefficients, C, are transformed to new variables x:

$$\boldsymbol{c}(\boldsymbol{x}) = \boldsymbol{c}_0 \cos(\boldsymbol{U}) + \boldsymbol{x}\boldsymbol{U}^{-1}\sin(\boldsymbol{U}) \tag{77}$$

$$\boldsymbol{U} = (\boldsymbol{x}^{\mathrm{T}}\boldsymbol{S}\boldsymbol{x})^{1/2} \tag{78}$$

$$\boldsymbol{c}_0^{\mathrm{T}}\boldsymbol{S}\boldsymbol{c}_0 = \boldsymbol{I} \tag{79}$$

$$\boldsymbol{c}(\boldsymbol{x})^{\mathrm{T}}\boldsymbol{S}\boldsymbol{c}(\boldsymbol{x}) = \boldsymbol{I}\ \forall\ \boldsymbol{x}. \tag{80}$$

Then $\boldsymbol{x}$ can be used in a minimization, and remain linear (compared with previous ideas). The constraint applied is that $\boldsymbol{x}^{\mathrm{T}}\boldsymbol{S}\boldsymbol{c}_0 = 0$. In the first-proposed form, a diagonalization

of $\boldsymbol{x}^{\mathrm{T}}\boldsymbol{S}\boldsymbol{x}$ is required to get eigenvectors. The linear-scaling method [352] uses iterative refinement for the orbital transformation (OT/IR)—a function is defined such that $f_n(Z) \sim f(Z)$ such that $Z^{\mathrm{T}}SZ = 1$. The authors propose to use the method by Niklasson [353], with fourth order found to be particularly efficient:

$$f_4(Z) = \frac{1}{128} Z(315 - 420Y + 378Y^2 - 180Y^3 + 35Y^4) \quad (81)$$

where $Y = Z^{\mathrm{T}}SZ$. The iterative refinement uses $f_n(\ldots f_n(Z)\ldots)$. To achieve linear scaling, a Taylor expansion is made for matrix functions rather than diagonalizing; after a conjugate step, iterative refinement is used to reimpose the constraint. The main drawback of the method is that the preconditioners required are based on dense algebra, giving $\mathcal{O}(N^3)$ scaling, but this may be lifted.

These methods are not yet widely used, but sit at an interesting junction between standard approaches and linear-scaling ones. The mathematical identities underlying the methods may well find wider use in linear-scaling applications, if they can be translated efficiently.

3.2.7. Quantum chemistry. The general area of wavefunction-based methods (starting, for instance, with Hartree–Fock wavefunctions and adding correlations via perturbative methods such as Møller–Plesset (MP) methods or coupled-cluster approaches) tends to be known as quantum chemistry. Many of the approaches in this area scale prohibitively with system size: the simplest, MP2, scales asymptotically with N^5; coupled-cluster single double (CCSD) methods scale as N^6 and CCSD with perturbative triples, CCSD(T), scales as N^7. Local correlation methods [354] can reduce scaling, and various other methods have been proposed which build on the ideas of locality for correlation. This is an area of sufficient complexity which has been reviewed elsewhere [354, 355]; below, we will briefly summarize some key ideas.

The MP2 approach in its canonical form involves pairs of integrals between occupied and virtual molecular orbitals. Local approaches to correlation reduce the scaling, but it is possible to produce linear-scaling methods. The key step in one approach was the development of Laplace MP2, where the exact energy is written in terms of non-canonical orbitals, via a Laplace transform of the original energy denominator. This can be extended to use atomic orbitals, which gives asymptotic N^2 scaling; by defining spherical interaction domains centred on atoms, efficient linear scaling has been demonstrated [356]. Another implementation of the same approach uses multipole-based integral estimates for screening [357]. An alternative approach builds on the local correlation methods [358], and uses density fitting and explicitly correlated wavefunctions (which depend on inter-electron distance, and improve the convergence of the method) [359]. It is also possible [360] to work in terms of atomic orbitals, and truncate the excitations based on the number of atoms involved in excitations (yielding a hierarchy of methods which can lead to linear scaling). This has been extended with a method for forming localized orbitals [173], with related work leading to a coupled-cluster algorithm which scales nearly linearly [361]. The implementation of some of these methods in a standard code has been discussed [100].

These ideas have also been extended beyond MP2: linear-scaling CCSD has been demonstrated using non-orthogonal localized molecular orbitals confined to fragments [362] and by expanding the coupled-cluster wavefunction in a local basis formed from a divide-and-conquer approach (see section 3.2.2); the localization is adapted dynamically to ensure error control [261]. MRSD-CI has also demonstrated linear scaling by using local correlation and integral screening [363].

The divide-and-conquer method itself (section 3.2.2) has also been extended to MP2 [364] and CCSD [260]. In both cases, the *full* HF orbitals from the subsystem are used (as opposed to other quantum chemistry approaches which typically localize the molecular orbitals or use atomic orbitals). Quantum chemistry is showing great promise in the area of increasing system size; at present, calculations on many tens of atoms are possible, and this trend should continue even for the most expensive methods.

3.2.8. Extensions. While much of the work on linear-scaling methods has been devoted to finding the ground state efficiently, there is also effort being put into going beyond the ground state to model the response of large systems to perturbations, in particular excitations. In this section, we briefly survey this work, though it is not a comprehensive list; other reviews cover parts of the ground in more detail [70, 355].

One obvious route for extending DFT is to perform real-time propagation of the density matrix (instead of the wavefunctions) within the framework of time-dependent DFT (TDDFT). It can be easily shown that the time variation is given by

$$\mathrm{i}\hbar \frac{\partial \rho(t)}{\partial t} = [H(t), \rho(t)] \quad (82)$$

though there has yet to be any comprehensive investigation of the effect of truncation on the accuracy of propagation. This has been implemented [365] following earlier work on time-dependent Hartree–Fock [366–368]. It has been extended to density functional TB [369] and applied to calculating spectral properties [370] using Chebyshev polynomials to expand the exponential of the Hamiltonian. However, the authors note that the amount of effort required is still high (and propagation for around 35 fs is needed for 0.1 eV energy resolution in medium to large systems, though the size is not quantified). Another approach to real-time propagation TDDFT [371, 372] uses random vectors and a projection method (as described in section 3.2.3) to calculate response functions. Using an empirical pseudopotential, it has been applied to calculation of optical properties of silicon nanostructures [373].

Standard implementations of TDDFT normally avoid real-time propagation, and instead search for solutions at the linear response level. The time propagation, equation (82), can be recast in terms of a Liouvillean superoperator, $\mathcal{L}$, whose eigenvalues represent vertical excitation energies. If the full Liouvillean is used, this is the random phase approximation (RPA), while if certain off-diagonal terms are neglected,

the Tamm–Dancoff approximation (TDA) results. By using projectors onto the occupied subspace (the density matrix, $\boldsymbol{P}$) and its complement ($\boldsymbol{Q} = \boldsymbol{I} - \boldsymbol{P}$), the eigenvalues of the Liouvillean can be found [374]; while Krylov subspace approaches (closely related to recursion methods described in section 3.2.3) are efficient, a direct variational method based on Rayleigh quotients [375] is also tested, and has been extended [376] and applied to CNTs and polymers. A different approach to speeding up TDDFT also uses recursion [377]. Again, the method avoids explicit representation of the virtual orbitals (a common theme in methods to improve speed and convergence of TDDFT and many-body perturbation theory), though it is not linear scaling in its present form. Another efficient, though not yet linear scaling, TDDFT method uses Lagrange functions (see section 2.3) and domain decomposition (see section 3.2.2) [378].

Density matrix perturbation theory [379] (which has been extended to non-orthogonal basis functions [229]) uses the trace-correcting TC2 method to generate a sequence of density matrices ($X_n^{(0)}$ for the unperturbed Hamiltonian). The expansion $X_n = X_n^{(0)} + \Delta_n$ allows a recursive expression for Δ_{n+1} to be derived in terms of Δ_n and $X_n^{(0)}$. The method is easily extended to different levels of perturbation theory. The first application [380] was to calculate polarizability of water clusters, going to 150 water molecules with a 6-31G** basis set. Other groups have built on this method for calculating polarization. One approach [381] uses only the non-diagonal part of perturbation to find polarizability and the Born effective charge in solids. A different approach to the coupled-perturbed equations [382] allowed linear-scaling calculation of the derivative of the density matrix with respect to a parameter (e.g. atomic position or electric field) using the McWeeny expansion; this method has been made numerically more stable and applied, among other things, to calculation of NMR [383]. It has recently been reformulated in terms of a Laplace transform [384] (in the same way as MP2 calculations, described in section 3.2.7, were reformulated) for greater efficiency. A less complex version of this approach has also been given, based on the projection properties of the density matrix; it has been shown to be competitive with other linear-scaling methods [385]. A post-hoc method for calculating polarizability with linear scaling [386] uses another variant of perturbation theory. A further approach to molecular response [387] uses an exponential parametrization (as described in section 3.2.6) to find excitation energies and polarizabilities.

Finally, it is clear that band edges can be found from linear-scaling methods (for instance, see the discussion in the work by Stechel [226]). Recent developments [388] have made the search more efficient. The maximum eigenvalue of ρH is sought *after* finding the ground state, using the density matrix as a projection operator; equivalently, the minimum eigenvector of $(I - \rho)H$ is sought for the LUMO. The simplest solution uses the Lanzcos algorithm for an extreme eigenvalue. The method has been applied to the case of a doped semiconductor where there are mid-gap or band-edge states. Band-edge states can also be found efficiently using iterative purification methods [389].

4. Technical details and parallelization

4.1. Non-orthogonal basis functions

In general, the localized orbitals used as basis functions in most linear-scaling methods are non-orthogonal; it has been shown that these functions are more contracted [169] and give computational advantages in various systems (e.g. silicon [390] and organic molecules [170]). This introduces complications in maintaining the correctness of tensors, and in defining different types of operator representations. The clearest notation uses covariant and contravariant notation (lower and upper indices), which was introduced for recursion methods with an excellent overview [289]. There is also a general explanation of the notation, and an application to second quantization [391]. The key implication for linear-scaling methods is the need for a good approximation to the inverse of the overlap matrix, or its decomposition, which can be found in linear-scaling time. We summarize the basics below, and urge the interested reader to find more in the references given.

If a non-orthogonal basis is defined, $\{|e_\alpha\rangle\}$, then the overlap matrix has elements defined by

$$S_{\alpha\beta} = \langle e_\alpha | e_\beta \rangle = S^\star_{\beta\alpha} \tag{83}$$

where we have assumed that the basis is real. Any matrix represented in terms of the original basis and notated with two lower indices is called *covariant*; it is actually a tensor. There also exists a *dual* basis, $\{|e^\alpha\rangle\}$, which satisfies the relation:

$$\langle e^\alpha | e_\beta \rangle = \delta^\alpha_{\ \beta} = \langle e_\alpha | e^\beta \rangle = \delta_\alpha^{\ \beta}. \tag{84}$$

A matrix represented in terms of the dual basis and notated with two upper indices is called *contravariant*; it is also a tensor. It is easy to show that, for a complete basis, we can write

$$\sum_\alpha |e_\alpha\rangle\langle e^\alpha| = \sum_\alpha |e^\alpha\rangle\langle e_\alpha| = \mathcal{I} \tag{85}$$

where $\mathcal{I}$ is the identity. The overlap between elements of the dual basis also forms the *inverse* of the overlap for the overlap matrix in the original basis; confusingly, different authors notate this in different ways. If we write $T^{\alpha\beta} = \langle e^\alpha | e^\beta \rangle$, then

$$T^{\alpha\gamma} S_{\gamma\beta} = \delta^\alpha_{\ \beta} \tag{86}$$

$$T^{\alpha\beta} |e_\beta\rangle = |e^\alpha\rangle \tag{87}$$

$$S_{\alpha\beta} |e^\beta\rangle = |e_\alpha\rangle \tag{88}$$

$$\sum_{\alpha\beta} |e_\alpha\rangle T^{\alpha\beta} \langle e_\beta| = \sum_{\alpha\beta} |e^\alpha\rangle S_{\alpha\beta} \langle e^\beta| = \mathcal{I} \tag{89}$$

where we have used the Einstein summation convention in equations (86)–(88). Some authors write $T^{\alpha\beta} = (S^{-1})^{\alpha\beta}$, while others write $T^{\alpha\beta} = S^{\alpha\beta}$. Some care has to be taken when considering different representations: as well as fully covariant and contravariant forms, there are mixed forms, such as $\langle e^\alpha | \hat{A} | e_\beta \rangle = A^\alpha_{\ \beta}$. A Hermitian operator is represented by a Hermitian matrix in the co- or contravariant forms, but not in the mixed form:

$$H_{\alpha\beta} = \left(H_{\beta\alpha}\right)^\star \tag{90}$$

$$H^{\alpha\beta} = \left(H^{\beta\alpha}\right)^{\star} \tag{91}$$

$$H_{\alpha}^{\ \beta} = \left(H^{\beta}_{\ \alpha}\right)^{\star}. \tag{92}$$

This necessitates careful notation, with the position of the indices (i.e. whether the upper or lower index comes first) being significant.

It is when considering differentials and the difference between classes of tensor that the notation and choice of metric becomes important. If the Hamiltonian is represented in terms of the original basis then it forms a covariant tensor, and the density matrix is a contravariant tensor. But the differential of the energy with respect to the density matrix (as used in methods such as the LNV technique described in section 3.2.1) is covariant, and should be scaled by an appropriate metric before it can be added to the density matrix [392]; this metric is $T^{\alpha\beta}$ (we note that it is possible to proceed with a different metric, which is equivalent to the original formulation [204]). If we write the auxiliary matrix $\sigma(\boldsymbol{r}, \boldsymbol{r}') = \sum_{\alpha\beta} \phi_{\alpha}(\boldsymbol{r}) L^{\alpha\beta} \phi_{\beta}(\boldsymbol{r}')$ and search for a minimum energy with respect to L matrix elements, for instance, then one search step in the minimization might be written:

$$\frac{\partial E}{\partial L^{\alpha\beta}} = \sigma_{\alpha\beta} \tag{93}$$

$$\sigma^{\alpha\beta} = T^{\alpha\gamma} \sigma_{\gamma\delta} T^{\delta\beta} \tag{94}$$

$$L^{\alpha\beta} = L_0^{\alpha\beta} + \lambda \sigma^{\alpha\beta} \tag{95}$$

where λ is a varying step length. Of course, the problem can be formulated in terms of a mixed representation, as suggested by Stephan [196], where all matrices are mixed. This requires that the Hamiltonian be premultiplied on the left by $T^{\alpha\beta}$, which connects closely with inverse preconditioning approaches to minimization, such as the AINV method described below; we also note that this goes back to observations on non-Hermiticity for localized molecular orbitals [129]. The search step is now written:

$$\frac{\partial E}{\partial L^{\alpha}_{\ \beta}} = \sigma^{\alpha}_{\ \beta} \tag{96}$$

$$L^{\alpha}_{\ \beta} = \left(L^{\alpha}_{\ \beta}\right)_0 + \lambda \sigma^{\alpha}_{\ \beta}. \tag{97}$$

An alternative approach is to effectively orthogonalize the basis, which can be achieved by transforming the matrices to an orthogonal representation. Some decomposition of the overlap matrix is required for this; either a symmetric one (using $S^{-1/2}$, the Löwdin transformation) or a non-symmetric one (e.g. a Cholesky decomposition [207], where $S = UU^{\dagger}$ and transforming with U^{-1}). Either way, we require either the inverse overlap matrix or the inverse of a decomposed matrix; naturally, these must be found using a linear-scaling method[8]. Cholesky decomposition can be made to scale linearly with system size for sparse matrices [207, 393], though these techniques are extremely hard to parallelize efficiently [394].

There is therefore a need to find the inverse or decomposed inverse of the overlap, for sparse matrices with linear-scaling time (and ideally good parallelization). An excellent overview of approaches from a computational science stance [395] makes the important point that the sparsity pattern of the inverse of a matrix may not be the same as that of the original matrix. This raises the problem of how sparsity is imposed, which is discussed fully below in section 4.3. We note that some of the methods in section 3.2.1 effectively converge on the inverse overlap matrix. The range and sparsity pattern of S^{-1} are extremely important, as is the condition number of S. It can be shown that, for a localized S, S^{-1} is exponentially localized [204], though the range will depend on the spread of eigenvalues of S. The condition number of the overlap matrix (the ratio of the largest to the smallest eigenvalue) will determine how easily an inverse can be found; the condition number will depend on the basis and the number of support functions/localized orbitals [396].

An iterative approach, known as Hotelling's method or Schultz's method [35], is extremely effective:

$$X_{n+1} = 2X_n - X_n A X_n \tag{98}$$

will converge quadratically on the inverse of A, so that $X_{\infty} = A^{-1}$. The iteration must be started from a suitable initial guess (which can be shown to be $B_0 = \epsilon A$, with ϵ a small number). This formula appears in a number of places: this iterative approach to updating an inverse from a previous step (or close to convergence) was suggested [226]. The close relation to iterative purification methods (section 3.2.1) should not be surprising, as the inverse overlap coincides with the density matrix when only occupied states are considered. The main drawback of this approach is that the iteration stalls when the truncation error is similar to the change in inverse at a single step.

A divide-and-conquer-like algorithm was first suggested in the context of recursion methods [397]: in order to form the matrix elements $M^{\alpha}_{\beta} = (S^{-1})^{\alpha\gamma} H_{\gamma\beta}$ for neighbours β of an atom α, a series of linear equations in clusters centred on each atom α in the system are solved: $H = MS$. A similar idea was put forward using conjugate gradients [226]: the problem can be written as $\sum_j S_{ij} D_{jk} = \delta_{ik}$, and conjugate gradients is then used to solve a least-squares problem with $(\sum_j \lambda_j^{(k)} \delta_{jk} - [SD]^2_{jk}$ for each localized orbital k, with $\lambda_j^{(k)} = S^2_{jk}$, or set to one only if within range). Enforcing symmetry and idempotency ($D = DSD$) for the diagonal elements was found to improve stability. An initial guess for the inverse is taken to be δ_{jk} or $2\delta_{jk} - S_{jk}$. Linear scaling follows from enforcing sparsity on the matrices.

Defining the generalized inverse of the overlap matrix, S^{-}, by $SS^{-}S = S$ [239, 398], a truncated approximation can be found by minimizing:

$$\mathrm{Tr}[BS^{-}] = \min_{X} \mathrm{Tr}[B(2X - XSX)] \tag{99}$$

where X is constrained to be Hermitian and B is *any* positive definite matrix. It was shown that a variational expression for the energy can be written for any number of localized orbitals (i.e. including unoccupied states) if the density matrix is taken to be $K = 2X - XSX$ and the energy is minimized with respect

[8] We note that various authors use cubic scaling methods to find the inverse overlap, on the grounds that the prefactor for this operation is rather small; while a pragmatic approach, it is not a scalable one.

to X. (A similar expression was found before [226], though without the variational derivation.)

The AINV method [399] has been used by Challacombe [209]) as a route to form $S^{-1}H$. The approximate sparse inverse is constructed from a sparse incomplete factorization of the overlap; while this is an effective method, it is hard to parallelize. If the overlap is assumed to be decomposed as $S = LL^{\mathrm{T}}$, traditional methods find L and then create $Z = L^{-1}$ by an incomplete linear solution. This can introduce inaccuracies, and the AINV method avoids these by solving directly for $Z = L^{-1}$. Once Z has been found, it is possible to create $Z(Z^{\mathrm{T}}A) \equiv S^{-1}A$ without ever creating the inverse [209], which may be dense or have unusual sparsity patterns.

Ozaki has proposed using the recursion method, normally applied to finding the density matrix, to solve for the inverse overlap [400]. He applies the block BOP method (described above in section 3.2.3) to the resolvent:

$$R(Z) = (S - ZI)^{-1} \tag{100}$$

for Z a complex number. It is clear that $S^{-1} = \mathrm{Re}[R(0)]$; however, for a method implemented with finite ranges, the matrix must be symmetrized to ensure stability. The method is also compared with three other approaches described above: the divide-and-conquer-like method [397]; Hotelling's method; and a Taylor expansion approach (as suggested in unconstrained minimization methods [163, 166] described above in section 3.2.1). All methods are effective for diamond, though the recursion method is slightly faster, while the Hotelling method yields larger errors for fcc Al. The choice of Taylor expansion (simply using a polynomial) precluded testing on the systems, as it did not converge. A dual-basis recursion method has also been suggested [296], with the dual basis calculated simultaneously with the basis (once a starting state has been defined).

A method for *improving* an approximate factorization of an inverse [353] (or an inverse itself, by decomposing $S^{-1} = IS^{-1}$) iterates $Z_{n+1} = Z_n(\sum_{k=0}^{m} a_k X_n^k)$ with $X_n = Z_n^{\mathrm{T}} S Z_n$; this method is related to another iterative approach [401]. The obvious drawback is the lack of an initial factorization, but this has been solved [402] by starting with a recursive decomposition of the S matrix, which allows a convergent calculation of the inverse. An extremely similar method has been derived for the symmetric square root [403]; both discussions also point out that convergence can be improved by scaling the overlap so that the eigenvalues lie within a convergent radius.

In all these approaches, there is the problem of whether the inverse (or the decomposed inverse) should be sparse at all, or share the same sparsity patterns as the overlap; in this case, a sparsity algorithm based on drop tolerances may give some advantages. VandeVondele [69] deliberately optimized basis sets to yield overlap matrices with small condition numbers, and show revealing data on the sparsity of inverse overlaps with basis size: the sparsity decreases as basis increases. Plots of number of non-zero matrix elements for different truncation thresholds and different systems [181] also show that the inverse is less sparse than the overlap. It is not yet unambiguously clear whether convergence can be achieved for ill-conditioned S matrices.

4.2. Preserving electron number

When varying the density matrix or localized orbitals to find the ground state, as well as maintaining idempotency (which has effectively occupied most of the methods discussed so far), the correct number of electrons must be maintained, as mentioned above. It is also possible to work at a fixed Fermi level (chemical potential for electrons, as suggested in one of the earliest methods [201]) though this is often a less physically reasonable approach. A grand potential is often defined:

$$\Omega = E_{\mathrm{Tot}} - \mu N_{\mathrm{e}} \tag{101}$$

$$= \mathrm{Tr}\left[K\left(H - \mu I\right)\right] \text{ if } \langle\phi_\alpha|\phi_\beta\rangle = \delta_{\alpha\beta} \tag{102}$$

$$= \mathrm{Tr}\left[K\left(H - \mu S\right)\right] \text{ if } \langle\phi_\alpha|\phi_\beta\rangle = S_{\alpha\beta}. \tag{103}$$

An early TB approach [404] included the derivative of the electron number in the gradient of the energy with respect to L matrix element, with the chemical potential from the previous step used for μ. This is similar to the approach used in the CONQUEST code [205], though instead of using the previous value of μ, the CONQUEST code projects out the direction of electron change ($\partial N_{\mathrm{e}}/\partial L_{\alpha\beta}$) from the search direction. After the line minimization for energy, a separate search for the correct electron number should be performed [205, 404].

An alternative is to treat the auxiliary density matrix L as the real density matrix, and use the McWeeny purification to alter the gradient [207]. In this case, the electron number becomes $N_{\mathrm{e}} = \mathrm{Tr}[LS]$ in a non-orthogonal basis, and it can be shown that a traceless gradient can be found by defining a slightly different grand potential, and treating μ as a parameter:

$$\Omega(L) = \mathrm{Tr}\left[KH\right] + \mu\left(\mathrm{Tr}\left[LS\right] - N_{\mathrm{e}}\right) \tag{104}$$

$$\mu = -\mathrm{Tr}\left[3(HLS + SLH)\right. \tag{105}$$

$$\left. - 2(HLSLS + SLHLS + SLSLH)\right]/N_{\mathrm{e}}.$$

The key advantage of working with this functional is that the search for the ground state does not perturb the electron number, so that given a starting point with the correct electron number, only μ needs to be updated. However, the density matrix thus defined will be less idempotent than K will be, imposing an additional approximation.

The final approach taken to maintain correct electron number is simply to rescale the density matrix, either after each line minimization [211] or continuously [212]. If the auxiliary density matrix is used as the density matrix, then a scaling factor $L \rightarrow N_{\mathrm{e}}L/\mathrm{Tr}[LS]$ is applied at each step. Alternatively, the purification transformation can be adapted to form a purification and normalization transformation [212]:

$$K = N_{\mathrm{e}}\frac{3LSL - 2LSLSL}{\mathrm{Tr}[(3LSL - 2LSLSL)S]} \tag{106}$$

where N_{e} is the number of electrons in the system. This transformation potentially introduces multiple minima, though is reported [212] not to adversely affect convergence.

4.3. Parallelization and sparse matrices

In this section we consider two important technical problems: parallelization of linear-scaling codes, and the implementation of sparse matrix methods. Sparse matrices are key to

linear-scaling computational time and storage, while efficient parallelization is needed for access to systems of more than 10 000 atoms or so. We will consider sparse matrices first, and then turn to strategies taken for parallelization.

We must consider both the technology of sparse matrices and how sparsity is imposed, and what errors different approaches will impose; the two main approaches to sparsity (or truncation of a matrix) are as follows: first, to consider a distance-based criterion (appealing to the results of section 3.1, so that an element is set to zero and neglected if the distance between atoms or localization centres is greater than some cutoff); and second, to drop an element if its value falls below some threshold. The first approach tends to give a clear, well-defined sparsity pattern while the second requires application of the tolerance after each operation. It is, however, easier to estimate and control an overall error due to sparsity with the second method. A recent analysis of matrix sparsity and how the function of a sparse matrix decays [405] provides an excellent, in-depth understanding of sparsity. It is highly recommended for all those developing methods in this area.

An early implementation of a TB OMM [165, 242] used distance-based truncation, and stored matrices in a compressed row format, following a parallelization described below. The communication pattern used for matrix multiplication followed a synchronous, Cartesian-based technique (for a processor at the centre of a cube, communicating with 26 other processors, communicate with neighbouring processors in the sequence: (1, 0, 0); (1, 1, 0); (0, 1, 0); (−1, 1, 0); etc). Similar approaches occur in other methods based on the orbital minimization family [243, 406].

The first approaches to drop tolerance used a criterion based on the individual matrix elements [206, 207]. Challacombe [209] introduced a sparse matrix algebra based on atom blocks which are dropped if the Frobenius norm of the *block* (defined as $A_{\mathrm{F}} = \sqrt{A_{IJ}^{\mathrm{T}} A_{IJ}}$ for atoms I and J) is smaller than a specified tolerance. The tolerance is re-applied after each matrix operation (e.g. addition or multiplication). The reasoning for using blocks rather than elements is that a change of basis is less likely to change sparsity patterns and associated errors. As the number of electrons on an atom can be related to the trace of the on-site block, this seems reasonable. A sparse-approximate matrix multiply (SPAMM) [209, 407] has also been introduced: the contraction over two atom blocks in a multiply ($C_{IJ} = A_{IK} B_{KJ}$) is neglected if B_{F} is smaller than the threshold divided by the number of blocks and A_{F}. This leads to small elements being treated more approximately than large ones, and potential computational savings. The method has been generalized [407] to a recursive approach on successively smaller sub-matrices; this bears some resemblance to an interesting approach to sparsity, that of hierarchical or $\mathcal{H}$-matrices [408], though these have not been used in linear-scaling methods to our knowledge.

A sparsity analysis [396] of matrices in electronic structure methods (pointed towards quantum chemistry methods) concentrated on linear alkanes, and applied drop tolerances. For the Hartree–Fock method, provided that a well-conditioned set of localized *occupied* orbitals exist, they prove both that the density matrix is localized, and that the overlap has a localized inverse. However, it does not follow that a localized density matrix can always be found (nor that localized orbitals can be found); a demonstration of the inverse of S was shown before [204], and its significance is discussed above in section 4.1.

CONQUEST [216, 410] truncates following distance-based criteria, and uses atom blocks and compressed row storage, with each on-site block stored in the place of a matrix element. An intermediate level of organization between atoms and the unit cell (called partitions, and described in detail below) is used to distribute storage and communications for multiplies. Special matrix-multiplication routines [410] have been developed, with a matrix kernel isolated to allow optimization. There are two multiply routines, depending on the radius of the final matrix compared with the initial matrices (*extension* when the final matrix has a larger radius than the other two, and *reduction* when the final matrix has a smaller radius). A similar approach to division of the unit cell and communications is used for the grid points (where the charge density is found on a real-space grid). Matrix multiplication efficiency has been shown for a number of systems, including large bulk silicon cells of up to 16 384 atoms [410] and, recently, perfect linear scaling and efficiency has been shown for systems containing millions of atoms [411].

ONETEP [412] uses atom-blocking, with data for columns of matrices stored on the atom-responsible processor. Hand-coded multiplies for block sizes relating to numbers of valence states (1,4,5,9,10) are used for efficiency. The developers suggest that sparsity of 90% or more is needed to benefit from routines (especially when going to product matrices without truncation); their analysis shows that the matrix KSK is 73% dense even for an 8000 atom system. More recent work [409] allows combined dense/sparse operations. Matrix columns are again divided into columns, and then further into process segments after assigning columns to processors. Each segment is designated dense (stored in full) or sparse based on fraction of non-zero elements and a threshold. Plots of density matrix structure for different systems are shown in figure 6, illustrating the different sparsities found in varying simulations. They show an analysis of performance versus sparsity threshold (which is typically set around 0.3 fraction), and find up to factor of two improvement over the original approach. Calculations have been demonstrated on systems with up to 36 000 atoms of silicon (as well as 16 000 atoms of DNA and 3200 atoms of alumina).

A multi-atom blocked approach to sparse matrix multiplication [413] has been developed for quantum chemistry based approaches. The method divides the cell along Cartesian dimensions in a binary way if the cell dimension is larger than some cutoff, R_{c}. This gives the multi-atom blocks, which necessarily include some zeros, but allow use of efficient BLAS calls. The block size depends on the basis, from about 30 for a minimal basis to about 70–80 for larger bases; the overall method is significantly faster than element-by-element sparse multiplies.

A study of truncation methods [235] suggested that distance-based truncation does not allow error control, and proposed a drop-tolerances method based on sub-matrix magnitude using a norm, as well as implementing sparsity

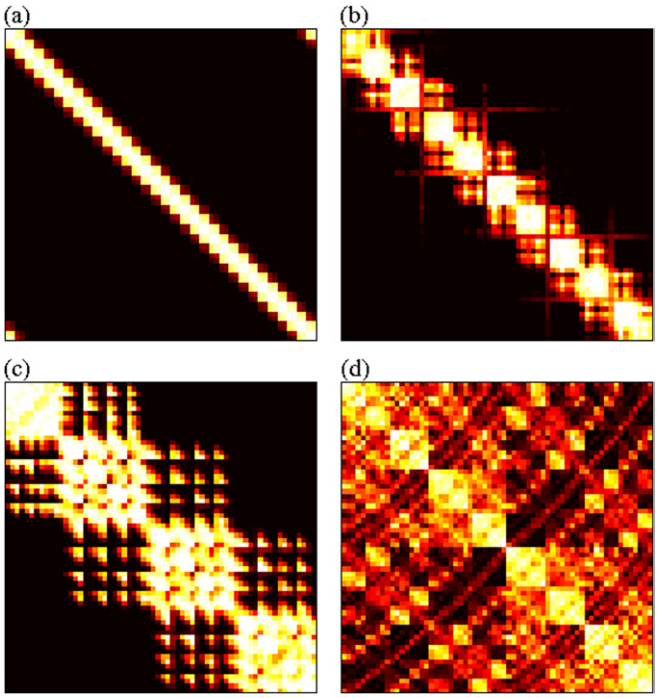

Figure 6. Segment-by-segment filling factors of sparse matrices in typical large systems divided over $P = 64$ processes. Matrices of the sparsity pattern KS (the product of the density kernel and overlap matrices) are shown for (*a*) a (10,0) zigzag nanotube (4000 atoms), (*b*) a 64 base-pair sequence of B-DNA (4182 atoms), (*c*) a H-terminated wurtzite-structure GaAs nanorod (4296 atoms) and (*d*) 8 (8 × 8) supercells of eight-atom cubic unit cells of crystalline Si (4096 atoms). Each pixel represents a segment, whose colour shows the fraction of matrix elements in that segment which are non-zero: black segments contain no non-zero elements, through red, then yellow, to white segments containing all non-zero elements. The non-zero elements are seen to be clustered near the diagonal of the matrix (though less so with increasing periodicity and complexity of the structure). The space-filling curve ensures that in a given column, there are non-zero overlaps only with rows of atoms on nearby processes, so the non-zero elements form a broad band centred on the diagonal. This is clearest for the simple structure of the nanotube, but even for the crystalline solid, there are segments in which there are a few or no non-zero elements. Reprinted with permission from [409]. Copyright 2009, American Institute of Physics.

control *during* matrix multiplication for efficiency. In this original paper, the 1-norm was proposed. A hierarchical approach to matrix storage evolved out of this work [414], which subdivides a matrix into sub-matrices, each of which can be further divided into more sub-matrices. At the lowest level, a matrix consists of real numbers; they found that five levels was enough for 36 000 × 36 000 matrices. The method permits blocks which are not related to atoms for performance reasons, and specifies uniform block sizes (32 × 32 in the examples given) at the lowest level. The developers use their own algorithms for symmetric square and inverse Cholesky based on symmetry, and show good performance and reduced storage relative to optimized libraries. The sparsity of the density matrix was studied for different molecules [181]. The effects of truncation were analysed, and possible ways of truncating matrices to a given tolerance were examined (so that the error introduced by truncation is controlled). A Euclidean norm-based method is accurate but computationally expensive, while calculation of Frobenius norms scales poorly with system size; they suggest a mixed norm based on the Euclidean norm of blocks, where the Frobenius norm of the block is found. A related idea [415] uses a tolerance based on either the number of atoms within a localization region, or a dynamically updated number of atoms based on the residual of the localized functions. We note, however, that a drop-tolerance approach can be less scalable than a distance-based approach, as the sparsity pattern changes with each iteration. The distance-based truncation is variational, while the drop-tolerance truncation may not be (and is often used in non-variational methods such as the purification methods).

The parallelization of linear-scaling techniques requires considerable care: the balance between communications and calculation will affect efficiency, and significant numbers of operations or variables which require storage or work on all atoms in the system on all processors will lead to N^2 scaling, as well as to memory problems for large systems. Nevertheless, owing to the spatial locality inherent in the methods, they are natural candidates for parallelization, and impressive scaling results can be achieved. We recollect the two types of scaling in common use: weak scaling, where the load per processor is kept fixed and the system size and number of processors are increased simultaneously; and strong scaling, where the system size is fixed, and the number of processors is increased. Efficiency for strong scaling is harder to achieve than for weak scaling, though weak scaling may well be the mode of operation chosen for linear-scaling codes.

An early TB method [242] divided space into a number of parallelepipeds equal to the number of processors, and assigned all atoms in a parallelepiped to a processor. The processor is then responsible for calculating forces and positions of atoms in the parallelepiped, and for storing matrix rows corresponding to those atoms. As this method also uses LWFs with well-defined centres not necessarily associated with an atom, the processor is also responsible for all LWF whose centre lies within its parallelepiped, and any matrices indexed purely by LWF. Scaling on up to 512 processors and 85 000 atoms was demonstrated, and an extensive analysis of scaling was made, noting that as the volume assigned to each processor decreases relative to a boundary area (due to localization radii) the amount of communication will change from depending on the number of processors (as $N_{\mathrm{proc}}^{-1/3}$) to depending on the volume of the boundary. The same approach has been used for an implementation of orbital minimization [164] within an *ab initio* TB method [243], though MPI and OpenMP parallelization are shared; the resulting code was demonstrated on up to 1024 processors and 6000 atoms.

Another implementation of orbital minimization [406] spent considerable time and effort on sparse matrices and parallelization to allow efficient MD. Matrices were stored as orbital blocks and neighbours of atoms (described as four-dimensional storage). The merits of particle versus spatial distribution of atoms between processors was discussed; in particular, list calculation (using the link-cell method) relies on locality for efficiency. They chose particle/orbital division for

simplicity (and as the method was communications limited). The problem of symmetric matrices and distribution is also considered: if symmetry is used to reduce storage, even distribution becomes more complex. Load balancing is achieved dynamically, by subdividing the system into 3D blocks. The blocks are ordered by x, y, z, with the atoms in each block ordered by z, y, x. The assignment to processors is balanced to even work (the time per site is calculated roughly). The resulting method showed good weak scaling, and reasonable strong scaling (the normal problem when going to a few atoms per processor, resulting in communications dominating).

The approach in the MondoSCF code (now called FreeOn) [416] is to order the atomic coordinates using a space-filling curve, so that atoms which are close in space are close on the curve. An overlap of communications and computation is achieved by posting a series of non-blocking receives (`MPI_Irecv`) and using blocking sends (`MPI_Send`). It is suggested that this arrangement (rather than a series of non-blocking sends followed by blocking receives) is more efficient, and less likely to lead to communication imbalance. Load balancing in the code is achieved by minimizing the imbalance of a characteristic matrix (typically the Fock or density matrix) based on the distribution of numbers of non-zero elements between processors. The resulting tests show reasonable parallelization up to 16 processors and sustained performance and efficiency increases up to 95 processors. Significant effort has been put into the linear-scaling calculation of the Fock matrix [103, 417]. The exchange-correlation matrix is calculated by hierarchical cubature [108]: a Cartesian grid is divided into blocks, with load balancing done dynamically to balance times. The key assumption is that the computational time is proportional to the total charge in box. Good scaling is found, and the approach is linear scaling. A similar approach is taken to the Coulomb problem (as described in section 2.5).

Conquest [216, 418, 410] subdivides the unit cell into small orthorhombic partitions of atoms and blocks of grid points (which are not necessarily the same shape and size), which are then distributed between processors using either space-filling curves [418] or an optimization procedure which considers both communication and computation time. Communication is performed by small group (leading to a compromise between local storage of unneeded elements and communications efficiency) [216]. The indexing and searching rely on different sets [410]: the primary set (the atoms or blocks for which a processor is responsible); the halo (all groups within range of a primary set); the covering set (an orthorhombic set of groups within range of a primary set for searching over). Processors are responsible for group of atoms (a bundle) and grid points (a domain) which should ideally overlap to reduce communications overhead. Scaling up to 4096 processors and 2 000 000 atoms has been demonstrated [411], as illustrated in figure 7. The code shows perfect weak scaling from 8 to 4096 cores, and reasonable strong scaling (hut clusters of Ge on Si(0 0 1) with either 11 620 or 22 746 atoms showed excellent speed-up from 16 to 128 cores and still 20 times faster going from 16 to 512 cores) [411].

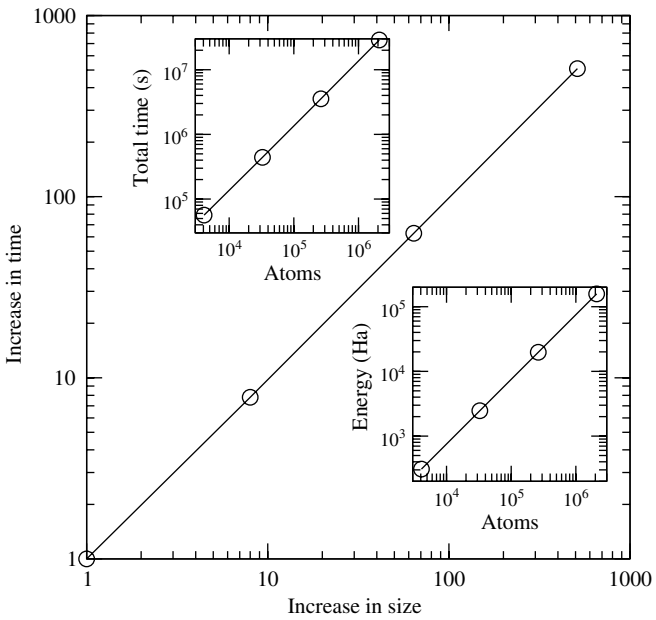

Figure 7. Linear and parallel scaling for bulk silicon on 512–4096 cores. The insets show total time and total energy (made positive to enable log plot) while main graph shows increase in time with system size; from [411].

There have been significant developments recently in the parallelization of onetep [419]. The first scheme used for parallelization of matrix multiplication [412] started on the diagonal and looped over processors using the modulus of the number of processors. This approach was much more efficient than the original communications which used an inefficient `MPI_Bcast` call. The localized orbitals (non-orthogonal generalized Wannier functions or NGWFs) are distributed by sending lists of parallelepipeds and psinc coefficients on the grid points in each parallelepiped. A batch system is used on columns (to get as many NGWFs as will fit in memory), with the most intensive operation being the Hamiltonian acting on the NGWFs. Using this scheme, scaling up to 27 000 atoms and 256 cores was demonstrated, though the speed-up when increasing number of cores by a factor of four was only 2.5. A new method [409] uses segments in each matrix to identify blocks, and each processor sends only contributing blocks. Rather than a round-robin approach (where core N sends to $N + 1$, $N + 2$, etc) the code now has an on-demand communication pattern. Scaling has been demonstrated again up to 256 cores, with a three times speed-up for increasing the core count by four times. onetep switches between sparse and dense matrix algebra depending on the filling of the matrices in question.

siesta can be run in parallel, though in its normal implementation only on small numbers of processors (scaling up to 32 or so). However, recent work [420] has changed the parallelization. The code originally divided up the cell by columns on the integration grid; the new implementation uses recursive bisection and weighting based on the number of neighbours. It also schedules communications to avoid unnecessary all-to-all communication. The resulting code shows an improvement for scaling 262 water molecules on

8–128 processors from 24% speed-up (old scheme) to 52% (new scheme—for an inhomogeneous distribution).

The divide-and-conquer method has been parallelized [421] following the obvious route: the individual subsystems are assigned to processors, with the limiting parallelization given by subsystem size. In the implementation described, the system data were copied to all processors, with the intended aim being small numbers of processors. The method has also been included in a massively multiscale MD approach with divide-and-conquer as the embedded quantum method [256, 258]. This approach has scaled to systems with 1.2 million quantum atoms and billions of classical atoms [257].

Orbital-free DFT has also been parallelized [422], by dividing up the real-space grid evenly between processors. Since the real-space data are purely real, only the positive half-sphere of reciprocal space is needed; these points are also divided evenly between processors, but the sparse grid involves a map. Load balancing requires consideration of a compromise: if it is performed by points, this gives ideal balance, but a communications penalty; by line gives better communications, but worse balance; by plane gives good communications by potentially poor load balance.

4.4. Structure relaxation

Before considering details of a subset of the applications which use linear-scaling methods, we note that the most common task required of a code is to relax the atoms to their lowest energy configuration. A point becoming more widely discussed is the scaling of relaxations: just because the electronic ground state can be found in linear-scaling time does not mean that a relaxed atomic structure can be found with the same scaling. Indeed, the difficulty of relaxing a system with low- and high-frequency phonon modes (or equivalently with very different length scales in the curvature of the energy) will depend on the ratio of the largest and smallest eigenvalues of the Hessian. It seems likely that, in general, the number of iterations to find a relaxed structure will increase with system size if relaxation is implemented naively; simple arguments indicate that, for conjugate gradient relaxations, the number of iterations will increase with the largest linear dimension in the system [423].

There have been a number of proposals to alleviate this problem. Preconditioning or relaxing low-frequency relaxations using insights from elasticity theory has been suggested [423, 424]. One approach [423] transforms the atomic forces onto a discrete grid (by smearing them) and uses multigrid approaches to solve for the Hessian without diagonalizing. The resulting method shows that the number of force evaluations (and hence time per atom required to find relaxed structure) is independent of system size on scaling from 500 to 800 000 atoms of silicon with an interatomic potential. An alternative approach [424] uses a model Hessian either to precondition conjugate gradients or as the input Hessian for a quasi-Newton method (in both following exact diagonalization or inversion, though this could presumably be made linear scaling if needed). The model Hessian is effective in improving convergence, though is not tested on increasing system size.

The obvious alternative approach, particularly for large molecules, is to use internal coordinates (bond lengths, angles and dihedral angles) instead of Cartesian coordinates. A linear-scaling transformation to and from internal coordinates has been proposed [425], which uses Cholesky decomposition on a sparse transformation matrix (which is shifted to avoid problematic zeros) and iterated. This method has been harnessed to allow a linear-scaling relaxation method [426] where the overall problem is decomposed into a set of $3N$ independent relaxation problems (using a weighted fit). The method is applied to a protein (with 263 atoms) and is efficient, though scaling is not tested explicitly.

Relaxation can be made more efficient, and less dependent on system size, if each step requires a shorter time to relaxation. A method to allow extrapolation of the density matrix from a previous step in a relaxation to the current one [427] uses the trace-correcting formalism to extrapolate. The method can be applied very efficiently to a local perturbation, with only the overlap matrix affected (and the density matrix extrapolated); it also improves convergence for standard geometry optimization.

The final approach extends the FMO method (section 3.2.2) to geometry optimization. Similarly to the perturbation approach just described, the system is divided into frozen, polarizable and active domains: the frozen domain is calculated just once, at the start, and not updated; the electronic structure of the polarizable domain is updated in response to changes in the active domain, but the atomic positions are frozen; while both the atomic and electronic degrees of freedom for the active domain are updated at each step (note that this is a form of embedding [428]). The method has been applied to prostaglandin synthase in complex with ibuprofen, with over 19 000 atoms.

5. Implementation and applications

5.1. Implementations

Implementations of linear-scaling methods can be split into two camps: first, those that build on existing codes and methods and simply add a new solver to the self-consistent loop (these often include quantum chemistry based codes and methods, such as the fragment methods described in section 3.2.2); and second, those that create an entirely new implementation, often involving careful parallelization (as described in section 4.3). At the start of this section, we will briefly survey the second class, though this is not an all-inclusive list, and more details can be found in papers referenced above describing approaches to linear scaling.

The main codes that we are aware of with linear-scaling functionality are, in alphabetical order: CONQUEST [429]; ERGOSCF [430]; FEMTECK [43, 237]; FREEON [431]; ONETEP [432]; OPENMX [433]; PROFESS [434]; and SIESTA [435]. We will discuss the basis sets used to represent the localized orbitals (or support functions or Wannier functions) as well as the linear-scaling kernel.

SIESTA [59, 76, 193, 436] was one of the earliest linear-scaling codes made widely available, and is still in widespread use today, though most of the applications use exact diagonalization to find the ground state. The code uses PAOs

(section 2.3) as a basis, and the method for calculating forces and stresses is clear [59]. The linear-scaling kernel used is the Kim, Mauri and Galli functional (section 3.2.1) which ameliorates the convergence problems of the OMM functional, and recently [255] a divide-and-conquer kernel has been implemented. Two different linear-scaling implementations of the vdW-DF functional for van der Waals corrections [437, 438] have been developed within SIESTA. The code is freely available for academic use.

The ONETEP code [212, 412, 419, 439–443] represents the density matrix in terms of NGWFs [444], which are in turn represented by periodic sinc functions (which are periodic, bandwidth-limited delta functions) on a fixed, real-space grid. The KE (and preconditioning) is calculated using local Fourier transforms, known as the FFT box method [445]. The linear-scaling kernel [212] is the LNV method [201, 204], and the method used for calculating forces has recently been published [446] (it is worth noting that the contribution from the local pseudopotential scales with N^2, but with a small enough prefactor that it only starts to become significant around 15 000 atoms in the published scaling tests). Empirical van der Waals interactions have been implemented [447] and extensive tests comparing with plane-waves have been carried out [96]. The code is commercial, though can be obtained for a modest fee by academics.

The OPENMX code [254, 295] uses a carefully constructed set of PAOs [63, 64] as a basis for the density matrix, and, similarly to SIESTA, concentrates on exact diagonalization and other efficient routes to the ground state [276], though recursion-based linear-scaling functionality is available [254]. In particular, the code uses an ingenious combination of divide-and-conquer and recursion methods. An extensive set of extensions has been implemented, including non-collinear spin, constrained DFT, LDA + U, Wannier function construction and polarization calculations. The code is available freely.

The FREEON (formerly MondoSCF) code [98, 108, 111, 113, 115, 209, 416] comes from the quantum chemistry community, using standard basis sets to represent the density matrix, with linear scaling coming from trace-correcting methods [234]. Forces are calculated within the standard framework [448, 449]. Recent extensions have considered perturbation theory and time-dependent DFT. The code is freely available.

The ERGOSCF code [104] also comes from the quantum chemistry community, using Gaussian basis sets to represent the density matrix and trace-correcting methods to find the ground state. The code is parallelized for shared memory machines, and is freely available under the GNU Public Licence.

The CONQUEST code [66, 67, 216, 217, 410, 411, 418, 450–453] can use either PAOs or systematically improvable blip (B-spline) functions to represent the density matrix, and the LNV approach to linear scaling, following initialization with McWeeny iteration; exact diagonalization via SCALAPACK has also been implemented. The forces are calculated as exact derivatives of the energy with linear-scaling time [454]. The code implements a number of standard features, including GGA with non-self-consistent forces [110], and recently spin and exact exchange. It also includes constrained DFT, which has been shown to converge in a linear-scaling manner [455]. The code is in late-stage beta release, and will be freely available in 2012.

The FEMTECK code [43, 237] uses FEs to represent the Wannier functions, and the augmented OMM method [237] to find the ground state in linear-scaling time. The code has been applied to liquid ethanol [44] and a fast ion conductor [456] with stability and efficiency.

The PROFESS code [332, 340, 341, 422] is an orbital-free DFT method (section 3.2.5) which represents the electron density directly on the grid, and has implemented a number of different KE functionals. As with all OFDFT codes, only local pseudopotentials can be used, which, along with the limitations on functionals, restricts its use to simple metals and some properties of semiconductors. Nevertheless, extremely large systems can be addressed: recent developments in ion–ion and electron–ion calculations have allowed the efficient parallelization of OFDFT, with a benchmark calculation on a million atom sample of perfect bulk Al demonstrated [341, 434]. The code is freely available.

5.2. Applications

In this section, we survey the applications of the linear-scaling methods explained in previous sections. There are many research areas where large-scale DFT calculations are expected to play an important role, and recently calculations for actual scientific research have been emerging. However, the applications of linear-scaling methods are still rather limited. It is necessary, at this stage in the development of the methods, to investigate the availability, accuracy and efficiency of the techniques employed in each study. It is also not obvious what kind of quantities can be calculated by such large-scale DFT studies. So far, different methods have been used depending on the system or the phenomena of the research area. From the examples of the applications introduced here, we would like to summarize what has been found so far, such as the size of the systems, which can be treated, robustness of the calculations, and the accuracy of the calculation method, including the quality of the basis sets used in the calculations.

5.2.1. $\mathcal{O}(N)$ calculations on biological systems. One of the most promising targets for $\mathcal{O}(N)$ DFT study are biological systems. In spite of the complex structures of large biomolecular systems, they provide atomically controlled systems for realizing surprisingly sophisticated functions. With the rapid increase in experimental information, the demand for accurate modelling of large biological systems is also growing. It is a challenge to understand the mechanism of such phenomena from the atomic scale, especially with quantum mechanics.

So far, most theoretical studies on biological systems from atomic scale have been done using classical force fields. Although these methods are powerful tools to investigate various phenomena in biological systems, they have a serious problem that the calculated results sometimes depend strongly

on the parameters used for interatomic potentials. Different force fields or even different version of the same force field can show qualitatively different results. In addition, it is quite difficult for the methods to treat the phenomena of bond forming or breaking properly. Thus, hybrid approaches like ONIOM or QM/MM (quantum mechanics/molecular mechanics) methods are often used for the study of chemical reactions, like enzyme reactions in biology. With these methods, the important region where chemical reactions take place is treated by a method based on quantum mechanics, and the dynamics or mechanics of the atoms in the surrounding region is calculated using a classical force field. However, it is sometimes uncertain how to define these two (or more) regions and it is not clear how accurate the method is, especially when the QM region is not large enough. Obviously DFT calculations on the entire or the sufficiently large region of complex biological systems are of great importance and linear-scaling DFT methods are expected to answer this request.

With such demands, there was already a report in 2000, showing the $\mathcal{O}(N)$ DFT calculations on a dry DNA model system containing 715 atoms [457]. The system is a periodic double helix DNA consisting of 11 guanine–cytosine (G–C) base pairs in the unit cell. For this system, $\mathcal{O}(N)$ calculations were performed using the Siesta code to employ the structure relaxation mainly with a double-ζ basis sets. Using the relaxed structure, a conventional diagonalization technique was also performed to calculate the KS energy and orbitals, and to confirm the accuracy of the forces with the $\mathcal{O}(N)$ method. The results for the simple polyG–polyC system show that the topmost valence bands are made by the 11 highest occupied molecular orbitals of guanine bases, and they are connected along the direction of the DNA chain; the 11 highest occupied molecular orbitals and the 11 lowest unoccupied molecular orbitals are illustrated in figure 8. The effects of a defect structure were further investigated by introducing the mutation of one G–C base pair, aiming to mimic the electronic structure of λ-DNA, which has a random sequence of the DNA bases. Following the same procedure, the orbitals of this system were calculated and they found that the orbitals showed cleavage at the point where the swap was introduced. This suggests that the resistivity of λ DNA should be very high, consistent with the measurements of the electron transport of the system. Although some of the results relied on a conventional method, this pioneering work clearly shows that $\mathcal{O}(N)$ DFT study will be powerful in the study of biological systems. Similar hybrid works, combining $\mathcal{O}(N)$ and conventional calculations on DNA, have also been performed [276, 458].

Biomolecular systems usually have large HOMO–LUMO gaps and thus the electronic structure is expected to be well localized. In this respect, the systems should be generally suitable for $\mathcal{O}(N)$ DFT studies. This aspect was clearly demonstrated in a theoretical study on a test DNA system using the CONQUEST and the DMM method [67]. The system investigated in their study is a B-DNA decamer (5′-d(CCATTAATGG)$_2$-3′), with 932 water molecules and 9 Mg^{2+} counter ions. Its structure is shown in figure 9(*a*), including 3439 atoms in total. The atomic positions were prepared from a snapshot taken from an MD simulation with the AMBER

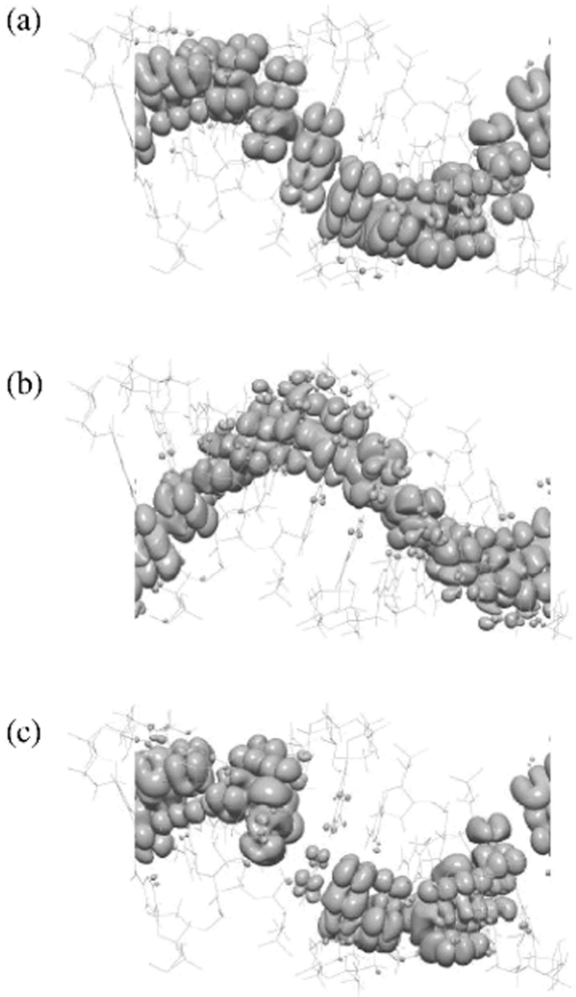

Figure 8. Isosurfaces of the electron density of dry DNA calculated using SIESTA [457] for: (*a*) the 11 highest occupied molecular orbitals; (*b*) the 11 lowest unoccupied molecular orbitals and (*c*) the 11 highest occupied orbitals following a single mutation. Reprinted with permission from [457]. Copyright 2000 by the American Physical Society.

classical force field. They showed that the DFT method can be applied to a system of this size, and the accuracy of the DMM method is reported to be surprisingly good. Figure 10 shows the dependence of total energy on the cutoff length R_L, calculated in a non-self-consistent way with a minimal basis set of PAOs. Here, both dry and hydrated DNA systems are calculated and the dry system is made by removing all water molecules from the system shown in figure 9(*a*). Since the dry system includes only 643 atoms (634 atoms for DNA and 9 Mg atoms), it was possible to also employ an exact diagonalization method to test the convergence of the linear-scaling method. The results are shown in figure 10. The energy obtained by diagonalization is shown by a horizontal dotted line, and the total energy calculated with various cutoffs using the DMM method is plotted by a red line with circles. These results clearly illustrate that the total energy by the DMM method converges very rapidly. The error at $R_L = 8.47$ Å is already 0.046 eV (7.2×10^{-5} eV/atom) and, if we increase R_L to 9.53 Å, the error becomes only 0.0078 eV (1.2×10^{-5} eV/atom). This rapidly convergent behaviour can

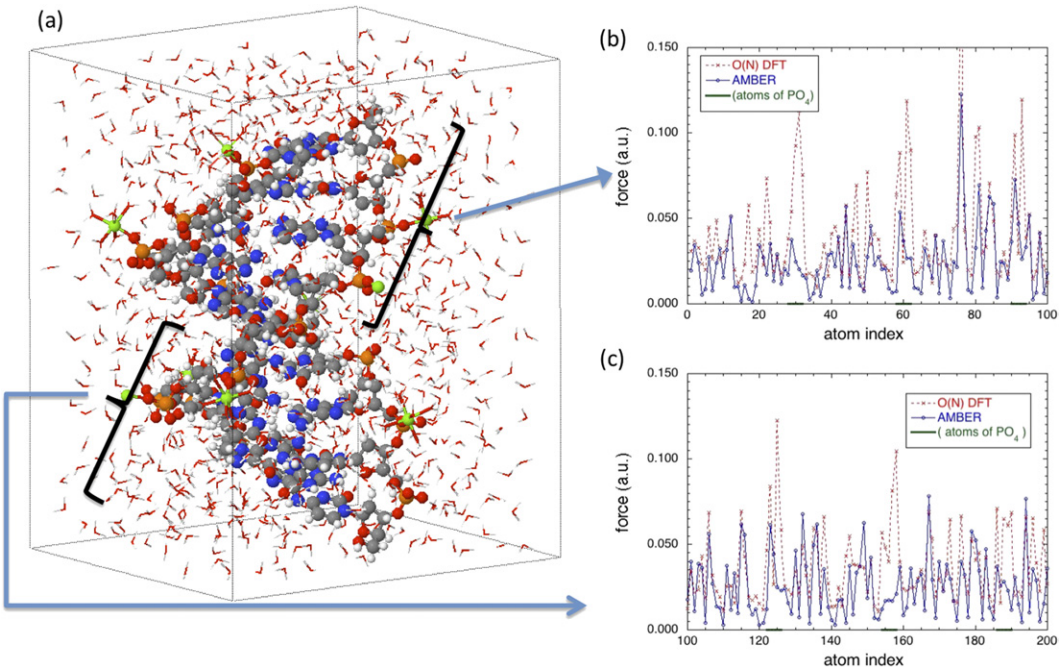

Figure 9. (*a*) Snapshot of the structure of hydrated ten-mer of DNA, and the calculated atomic forces on (*b*) 1st–100th and (*c*) 101st–200th atoms of DNA by O(N) DFT and AMBER force field calculations. In (*b*) and (*c*), green bars on the horizontal axis show the atoms of phosphoric acids.

also be observed in the calculation of a DNA system hydrated with many water molecules. The total energy of this system calculated as a function of R_L is plotted by a blue line with triangles in figure 10. Note that the energy scale for this system (right in figure 10) is the same as the one for the dry system (left in the figure), though shifted. The convergence of this system is also very rapid and the total energy at $R_L = 13.23$ Å can be considered as a well converged value. Then, the error at $R_L = 8.47$ Å is 0.094 eV, which corresponds to 2.7×10^{-5} eV/atom. If we use $R_L = 9.53$ Å the error becomes 0.017 eV for 3439 atoms (4.9×10^{-6} eV/atom). With such accuracy, it is possible to discuss the difference of the total energy induced by a local reaction in a system containing several thousand atoms. There are reports showing that similar accuracy can be obtained also in other linear-scaling methods [276]. These results suggest that $\mathcal{O}(N)$ DFT methods will be able to provide quite accurate results in molecular biology.

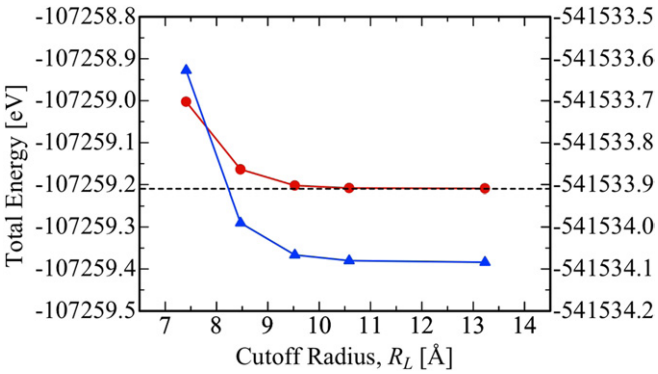

Figure 10. Dependence of total energy on density matrix cutoff, R_L, for DNA without water molecules (red line and circles, left axis) and full system (blue line and triangles, right axis). The dashed line shows the exact diagonalization result for DNA without water; from [67].

The FMO method (described in section 3.2.2) is efficient and accurate for biomolecular systems, and already has many published examples [272]. One of the largest problems tackled with the approach was a calculation of active sites on influenza A viral haemagglutinin [459]. The calculation used a polarizable continuum model and about 24 000 atoms of protein. It was found that the binding of a class of chemicals known as sialosides is not regulated by allosteric effects (in other words the number of ligands bound does not affect binding affinity); this result can be used in the design of drugs for coping with influenza pandemics. One of the advantages of the FMO method is that DFT can be used, as well as Hartree–Fock and post-Hartree–Fock methods, like the MP2 method. Another advantage is that the method can easily calculate the quantities by which we can discuss whether the interactions between two fragments are attractive or repulsive. This is called pair interaction energy decomposition analysis, and can give useful information for the ligand–protein interactions. There has been an attempt to apply the FMO method to silicon systems [460], but we note here that an investigation of the MTA fragment method for 2D π-conjugated systems [461] found that large fragments were needed for accuracy in these types of system.

In the actual research on biological systems using DFT, it may be a serious problem that DFT functionals such as LDA or GGA cannot express van der Waals interactions correctly. However, there have been many attempts to solve this problem. A new class of exchange-correlation functional which includes the vdW interactions has been proposed by Dion *et al*, called the vdW-DF method. There are already many examples using this method, and a revised version (vdW-DF2) has recently

been proposed [462]. There is also a more efficient method, called the DFT + D method, which has been applied to many biological or organic systems [463, 464]. This is a simple method where total energy and forces are calculated by adding empirically parametrized interatomic potentials to the total energy calculated by DFT. In the early version of the method, the parameters used in the method depended only on the species of atoms and did not vary in different environments. Recently, methods to improve the transferability, by changing the parameters using the charge density calculated by DFT method or from the analysis of the local coordination numbers, have been proposed [464, 465]. The results obtained by the new methods reported so far are encouraging. These methods, vdW-DF or DFT-D, can be easily applied to the $\mathcal{O}(N)$ methods and there are already some reports for their implementation to the $\mathcal{O}(N)$ codes [437, 447]. The vdW interactions are usually very weak and seem to be more effective for larger systems, thus more important in $\mathcal{O}(N)$ DFT calculations.

Another serious problem in the DFT study on biological systems is that the simulation time of MD is very short. One candidate to overcome this problem at present is to use the 'force matching' method [466], which aims to refine the classical interatomic potentials using the DFT results. To employ this method in complex biological systems, we must be able to calculate the total energy and atomic forces for very large biological systems using DFT.

As mentioned previously, the DMM method can calculate the atomic forces easily and accurately. For the system in figure 9(*a*), the total energy and atomic forces are recently calculated with DZP basis sets and compared with those by the AMBER force field [467, 468]. The calculated forces acting on the 1st to the 200th atoms, from a part of DNA, are shown in figures 9(*b*) and (*c*). In the figure, the green bars on the horizontal axis show the indices of the atoms forming the phosphoric acids. The result shows that the atomic forces calculated by these two methods agree well for most of the atoms. However, we can see that the agreement for the atoms in the phosphoric acids is much worse, compared with the forces on atoms of DNA bases. The deviation in the forces on the phosphoric acid part seems to depend on the position of the Mg counter-ion close to this part. We can expect that such DFT results would be useful to revise the accuracy of classical force fields. All of these results suggest that $\mathcal{O}(N)$ DFT (or other quantum mechanics) methods will soon be able to play a significant role in the study of complex biological systems.

5.2.2. Order-N DFT study on nanoscale structures of Ge islands on Si(0 0 1). Another class of targets which can greatly benefit from $\mathcal{O}(N)$ DFT calculations are nanomaterials or nanoscience. Here we would like to introduce the study on the energetics of a nano-structured system, Ge 3D islands grown on Si(0 0 1). The Ge/Si(0 0 1) system has been extensively studied because it is a prototypical example of hetero-epitaxial Stranski–Krastanov growth. It is also technologically important because of the formation of organized quantum dots. Many experimental studies have been reported so far, and they confirm the following results. When Ge atoms are deposited on Si(0 0 1), growth initially occurs layer by layer, up to a critical thickness of about three monolayers (ML). Strain due to lattice mismatch is relieved by the formation of regularly spaced rows of dimer vacancies in this two-dimensional (2D) structure, resulting in the $2 \times N$ structure. Deposition of further Ge atoms leads to another strain-relief structure, 3D pyramid-like structures called hut clusters [469], whose four facets are well established to be {1 0 5} surfaces. The typical side length of the hut cluster is about 150 Å. Recently, all-atom DFT calculations on the hut cluster including a substrate were performed using an $\mathcal{O}(N)$ technique to study the transition from the 2D to 3D structures [470, 471]. Here, we introduce this $\mathcal{O}(N)$ DFT study.

The stability of the 3D structures grown on surfaces is usually governed by the competition between the release of the strain energy from the formation of a 3D structure and the energy increase due to the larger surface area of the facets on the surface. So far, theoretical approach on the Stranski–Krastanov growth mode has usually used continuum elasticity theory for the first part, and employed DFT calculations for the latter. However, a unique situation exists in the Ge/Si(0 0 1) system. The surface structure of the strained Ge(1 0 5) was clarified by a combination of STM and DFT studies and the DFT calculations show that the strained Ge(1 0 5) surface is much more stable than the strained Ge(0 0 1) surface. Therefore, even though the surface area increases, the contribution from the surface energy is extremely small or may be lowered by the formation of facets. If this is the case, we need other contributions to reproduce the 2D–3D transition. It is important to consider the energy contributions from the wetting layer, as well as the edges where the facets meet each other. In addition, as the area of the facets of the experimentally observed Ge hut cluster is not large, the evaluation of the surface energy using conventional DFT is also doubtful. For these reasons, the validity of previous theoretical approaches is uncertain, especially for small hut clusters. To overcome these problems, it is necessary to employ all-atom DFT simulations on this system, including the entire Ge hut cluster, wetting layer and Si substrate. Since the number of atoms exceeds a few thousands even for small hut clusters, we need a linear-scaling technique to employ DFT calculations on such large systems and it has recently been shown that the calculations are possible with the CONQUEST code.

Before the work on full systems, the accuracy of the computational methods was thoroughly examined by the calculations on the strained Ge(1 0 5) surfaces [470]. It should be noted that even the semi-empirical TB method, though it is based on quantum mechanics, is found not to be accurate enough for the energetics of the surfaces in this system. Using the results, $\mathcal{O}(N)$ DFT calculations on the 3D Ge hut clusters have been employed [471]. At the non-self-consistent level, structural optimization on systems of different sizes were employed using a standard CG method. The largest system calculated in this work, shown in figure 11, contains ~23 000 atoms. As we can see in figure 12, the structure optimization is robust and accurate enough even for such large systems. For the 3D hut clusters, three structural models having different facet and edge structures are investigated. Furthermore, the structure of 2D $2 \times N$ reconstructions ($N = 4$, 6 and 8) and

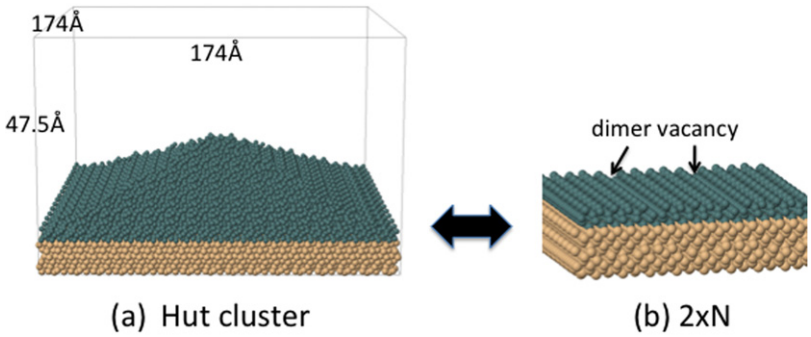

Figure 11. (*a*) Structure of Ge hut clusters on Si(0 0 1). (*b*) Structure of the wetting layer of Ge on Si(0 0 1), showing the $2 \times N$ reconstruction.

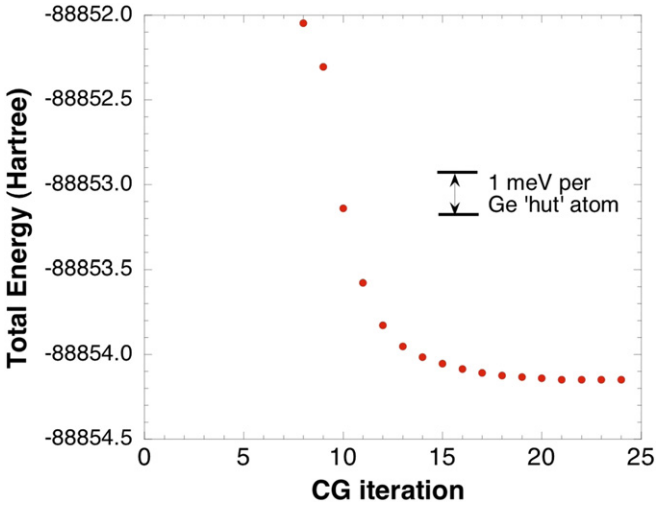

Figure 12. Convergence of total energy during structural relaxation of Ge hut cluster on Si(0 0 1), which contains over 20 000 atoms; from [471].

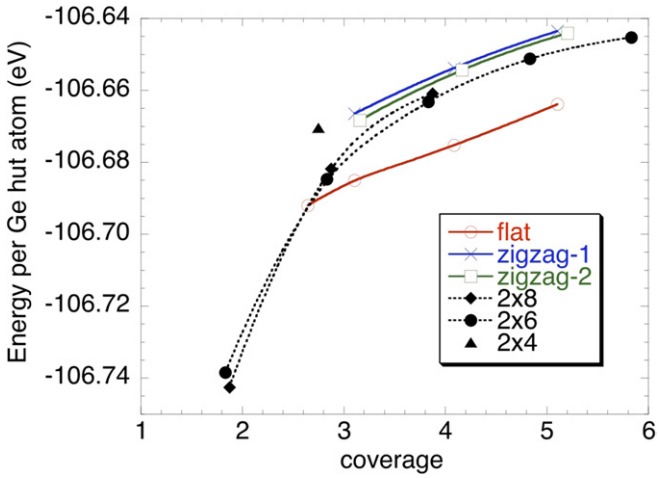

Figure 13. Comparison of energetics for different 2D (dotted lines with filled symbols) and 3D structures (solid lines with open symbols or crosses) for Ge on Si(0 0 1); from [471].

its total energy are calculated for comparison using the same calculation condition. The energy of these structural models as a function of coverage is illustrated in figure 13. It shows that the 2D structure is more stable for small coverage of Ge atoms, but the 3D hut structure becomes stable when the coverage exceeds 2.7 ML. Interestingly, this coverage agrees with the experimental value showing the 2D–3D transition. This $\mathcal{O}(N)$ DFT study has succeeded in clarifying the energetics of the 3D hut cluster systems, but the kinetic aspects are also important to simulate actual growth. In this respect, since the $\mathcal{O}(N)$ DFT study can treat the entire system, it is also possible to work on the dynamical aspects by putting additional Ge atoms on the optimized structures. Such works have been unavailable so far but we expect fruitful information could be obtained by $\mathcal{O}(N)$ DFT studies in the near future.

5.2.3. Other examples. Here, we survey other applications. Of course, it is impossible to show all examples, but we will try to show various areas of applications using different $\mathcal{O}(N)$ methods.

First of all, an excellent problem for linear-scaling methods is a 1D system; there is also considerable interest in the transport properties of molecules [472, 473]. One approach to calculating transport for large systems [474] divides the system into layers with local coupling between them, and then uses normal DFT calculations to find the electronic structure of the layers. The resulting method is linear scaling (and related to divide-and-conquer techniques, section 3.2.2). This approach was applied to defects in CNTs, and the localization lengths associated with them. By examining systems with up to 25 defects it is predicted that Anderson localization will be observed in CNT at room temperature. The method allows systems hundreds of nanometres long to be simulated.

Divide-and-conquer-like algorithms have found a relatively wide application. The ONIOM method [475] was used, along with successive geometry updates on different fragments, to relax the 150 000 atom photosystem I trimer [476]. The resulting relaxed structure allowed quantitative identification of the location of key hydrogen atoms, as well as the insertion of missing atoms and correction of misaligned features from x-ray diffraction data. A similar approach, with divide-and-conquer embedded within a multiscale modelling framework [477], allows for massively parallel deployment. Tests have been carried out on [258]: alumina up to 11 796 480 atoms on 131 072 Blue Gene/L processors (though with a rather coarse grid of $0.4a_0$); MD for 432 atoms of Rb, which yields good agreement with x-ray pair distribution data; MD on 512

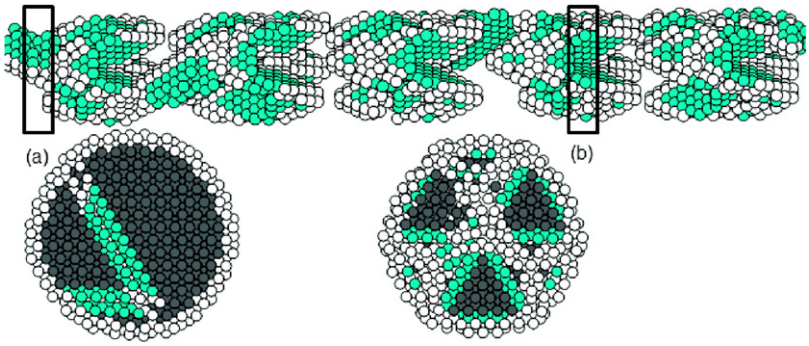

Figure 14. Slip planes of a 4 nm diameter Al nanowire formed upon loading one step past the elastic limit. The top image shows the non-fcc interior atoms for the entire 20 nm long nanowire ([1 1 1] axis extending from left to right). The bottom two images show all the atoms in three-layer cross sections of the nanowire at (*a*) 4–10 Å and (*b*) 164–174 Å along the wire length. Light blue atoms have hcp structure; grey have fcc structure; and white have unknown structure. Reprinted with permission from [480]. Copyright 2011 American Chemical Society.

atoms of graphene to look at vibrational spectra; and calculations on the electron affinity of a CdSe nanorod with 432 atoms. The same method has been used to simulate a thermite reaction with 1152 atoms (Al/Fe_2O_3), performing MD at 2000 K for 5 ps. The key result was a metal–oxygen flip mechanism that enhances diffusivity [478].

Orbital-free DFT (section 3.2.5) is generally used for metallic systems. The PROFESS code has been used extensively to model Al, such as the energetics and mobility of vacancies [479], Al nanowires [479, 480] and crack tips in Al [481]. In the first example, vacancy formation and migration energies are calculated using cells up to 500 atoms for tri-vacancies. The authors find that while nearest-neighbour vacancy pairs are unstable, next-nearest-neighbour vacancy pairs are stable, and predict that vacancy clusters grow preferentially through next-nearest-neighbour vacancies. For the second example, Al nanowires (up to 16 770 atoms) with 1–8 nm diameter and up to 20 nm long are stretched to examine elastic and plastic behaviours. The elastic deformation is found to be qualitatively similar, but quantitatively different with respect to the diameter; thinner nanowires are more compressed relative to the bulk fcc structure. On the other hand, clear size-dependent behaviour is observed in the plastic region. Partial slip as mechanism for plastic deformation is only seen for 4 nm wires and above, while amorphous deformation is seen in narrower wires. These are illustrated in figure 14. In the third example, the system is also calculated with the embedded atom model (EAM) method. With the OFDFT a system of up to 7800 atoms is treated. Qualitative differences of the OFDFT result to EAM for one orientation are found, in particular in how emission of twinning partial dislocations changes crack length; the difference is likely to be down to the surface energies being incorrect for EAM. There is also a quantitative difference in the onset of emission of partial dislocations with load.

For the study of vacancies in aluminium, there is another OFDFT study which combined FE modelling and coarse graining with OFDFT [482]. In this method, far from the region of interest, they use the energy of distorted bulk Al based on the local environment to coarse-grain. The authors find that 10^3–10^4 atoms are needed for convergence of monovacancy formation energy. Interestingly, they find a change of sign in interaction energy for the di-vacancies, when they increase the size of the system from 32 atoms to larger. However, this result does not seem to be consistent with the previous result by PROFESS. In particular, the results for monovacancy formation energy, while giving almost the same value as the PROFESS result [479], show different behaviour with cell size. There may be some effects from the coarse graining approximation; however, much of the difference comes from the boundary conditions used, particularly for the electrostatics, and not from the OFDFT component (which seems to be consistent between the two implementations) [483].

The OMM (section 3.2.1) has been applied to molecular dynamics (MD) calculations of liquids. It has been used to test the effect of cell size on diffusivity of liquid water [484]: the diffusivity of liquid water calculated from DFT is too low over long timescales. Using MD simulations with up to 128 water molecules, the study found no appreciable size effects (between 32 and 128 molecules), nor any effects due to parameters chosen (including the localization radius and basis chosen). The final, detailed study used exact diagonalization with a DZP basis optimized for water, with basis set superposition errors corrected.

The augmented OMM has been used to perform challenging and scientifically relevant simulations with an $\mathcal{O}(N)$ code. It has been applied to MD simulations of liquid ethanol [44]; this system requires large system sizes (of the order of 10^3 atoms to ensure the structure of hydrogen bonded chains is correct). Using seven localization regions for each molecule (centred on the bond centres with one on the oxygen) with radii of $12a_0$ energy is conserved extremely well, and the computation cost is reduced by a factor of 4.6 compared with a conventional method. The comparison with experimental data is impressive: the radial distribution function shown in figure 15 compares well; the self-diffusion coefficient is close ($8.2 \times 10^{-6}\,cm^2\,s^{-1}$ for simulation compared with $1.1 \times 10^{-5}\,cm^2\,s^{-1}$ from experiment); and the red shift in O–H vibration mode due to hydrogen bonds in the liquid agrees as well ($\sim 350\,cm^{-1}$ in simulation compared with $320 \pm 10\,cm^{-1}$ in experiment). The same method has been applied to calculations of Li^+ conductivity in $LiBH_4$ [456]. Using 100 unit cells (each with two $LiBH_4$ units, for a total of

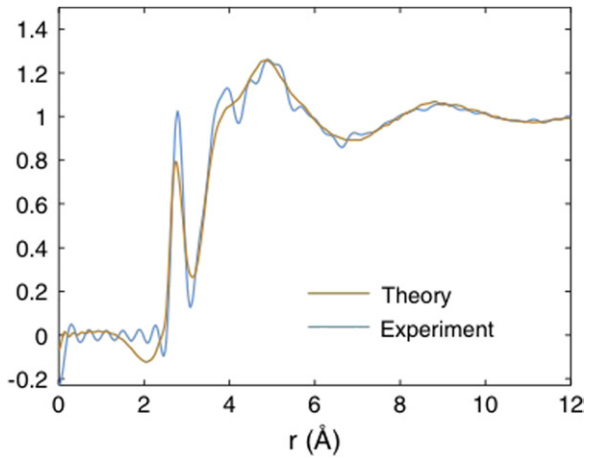

Figure 15. X-ray weighted radial distribution functions of liquid ethanol calculated from *ab initio* O(N) MD simulations (brown line) and obtained from experiments (blue line); from [44].

1200 atoms) and localization regions of $10a_0$ for Li and $14a_0$ for B, the CPU time was reduced by a factor of 2.5 with no reduction in accuracy. The authors found that the high ionic conductivity results from a metastable interstitial site generated by a splitting of the units in the *c* direction, which is a new mechanism for ion conductivity.

The LNV method is implemented in both CONQUEST and ONETEP; applications of CONQUEST have already been described above in sections 5.2.1 and 5.2.2. A study of the formation energy of vacancies in alumina with ONETEP [409] (α-Al_2O_3) found, as would be expected, significant simulation cell size dependence. Using cells from 120 atoms up to 3240 atoms, a linear correlation between defect formation energy and Madelung energy was found, and used to extrapolate an infinite cell size result. The simulation used a density matrix range of $24a_0$ and radii of $8a_0$ for the Wannier functions, though the convergence of the energy with these values is not shown. The same approach has been used to calculate the electronic structure of silicon nanorods with hydrogen passivation on the surface [485]. The calculations used a kernel radius of $24a_0$ and a Wannier function radius of $7a_0$. Nanorods containing up to 1648 atoms were modelled, though the geometry optimization was performed using an *ab initio* TB code. The DOS and states around the band gap are found using exact diagonalization.

Krylov subspace methods (section 3.2.3) have been applied with TB Hamiltonians (rather than *ab initio*) to study cleavage in silicon. The cleavage of nano-crystalline silicon [486] (for systems up to 100 000 atoms) showed formation of (1 0 0) facets with a partial reconstruction, and explored the effects of size difference and competition between bulk and surface energy. Cleavage of bulk Si along the (1 1 1) plane [487] used 11 096 atoms (periodic in one direction only). The formation of a (2×1) reconstruction as well as steps during cleavage was observed, with a similar interplay between surface and bulk terms.

As we have mentioned above, there are still not many scientific applications using $\mathcal{O}(N)$ DFT methods. However, as we can see from the examples surveyed in this section, there are now real simulations using $\mathcal{O}(N)$ DFT methods and the number of applications is growing rapidly. As the codes develop in robustness they will be more widely used, and more experience will emerge in important areas such as convergence and basis sets. This will in turn encourage confidence in results and more applications.

6. Conclusions

In recent years, the trend in computing has been a dramatic increase in the numbers of cores on a processor, and to massive numbers of processors in high-performance computing centres; the recent emergence of hundreds or thousands of processors on graphics processing units has taken this trend even further. As computational science is driven by the hardware base available, we have seen a strong movement towards real-space basis methods as a route to exploit hardware efficiently, though the eigenvalue solvers have retained traditional scaling rather than shifting towards linear scaling.

Since the first DFT linear-scaling methods were proposed 15–20 years ago, it can legitimately be said that development in $\mathcal{O}(N)$ methods has been slow. However, in recent years, real progress towards applications has been made (for instance, see the proceedings of a CECAM workshop held in 2007 [488]). The reasons for this slow development are easily understood: linear scaling introduces more parameters and sources of instability; standard methods have developed rapidly in efficiency and robustness; parallelizing linear-scaling methods is complex.

Nevertheless, we have reached a point where approximate linear-scaling methods (such as divide-and-conquer and orbital-free DFT, which are hard to pursue to high accuracy) are producing real applications, as discussed in section 5.2. Methods which have the capacity for exact behaviour are now at the point where they are more efficient than conventional methods for systems over about a thousand atoms, and are starting to demonstrate real applications with predictive capability.

Linear-scaling codes can seem more complex to use than standard codes; at the moment there is certainly less expertize in their use and appropriate convergence criteria compared with standard methods. One of the key advantages of a plane-wave basis set is the simplicity of convergence: the cutoff energy offers a single, simple variational parameter; linear-scaling methods which use a variational basis set (such as blips or psincs) have a directly equivalent grid spacing. Atomic-orbital- or Gaussian-based codes, whether conventional or $\mathcal{O}(N)$ methods, by contrast, have basis sets which are hard to converge in a variational manner. Linear-scaling methods do have a cutoff both on the density matrix and on the localized orbitals, though again cutoffs on local orbitals are not unique to linear-scaling codes. We note that these codes do *not* in general include an integral in reciprocal space over the first Brillouin zone, and therefore *remove* the need to converge the *k*-point mesh. Linear-scaling methods do introduce an extra convergence criterion during minimization (one for the solution of the density matrix and another for the localized orbitals). So we see that, while apparently more complex, there

is no reason why linear-scaling methods cannot become as widely used as conventional methods. With more researchers working on practical calculations using these methods, the field will rapidly improve, in the same way as the field of *ab initio* calculations changed in the years after the Car–Parrinello method was first proposed.

Just as in conventional approaches, there is no consensus on appropriate basis sets (in particular whether atomic-like or variational is better). However, this should not hold the field back: a variational basis can be characterized in the same way as plane-waves or real-space grids, and the same cutoffs can be used. There is a widely held assumption that double zeta (or double valence) plus polarization basis sets are required for atomic-like representations, though without the consensus achieved in quantum chemistry on different basis set qualities and likely errors. This should come with maturity of the field. Even within variational basis sets, however, it is not yet clear how many orbitals are required when we impose localization on the basis set functions: the smallest size, the same as the number of bands, carries potential convergence problems; a number equivalent to a minimal basis gives a good compromise between computational effort and variational freedom; some calculations have found that larger numbers (e.g. equivalent to the valence orbitals plus polarization orbitals) are needed for accuracy, while other calculations achieve accuracy with the same number of orbitals as valence orbitals. As more calculations are performed, a deeper understanding will emerge.

We now consider the challenges faced by the linear-scaling community and future routes for development. Naturally these are personal choices, but they certainly represent important problems in the area. The first challenge is that of accuracy: how accurately can linear-scaling methods reproduce exact methods, and with what accuracy can important quantities be calculated? The question of accuracy (and convergence) becomes more important when considering energy differences in large systems, which is a natural area for applications of linear-scaling codes: tight convergence will be required, for instance, when comparing different structures in biomolecules. There have been some investigations already in this area (showing energy difference convergence for Ge nanostructures on silicon [471] and for solvated DNA [67], comparing absolute energy convergence between exact and linear-scaling methods for bulk silicon [96], and comparing performance for purification and DMM methods for a linear alkane [192], as well as work on error control within linear-scaling methods [235]) but more of these studies are needed. It is becoming clear that good accuracy can be achieved, though it naturally increases the computational time required; this accuracy is important for the future of these methods.

The second challenge is that of metallic systems: there is no clear route to linear-scaling solution for systems with low or zero gaps and extended electronic structure. While methods such as orbital-free DFT offer some hope, and certainly allow large system sizes to be addressed, they do not at present give sufficient accuracy for quantitative calculations. There are approaches with reduced scaling over standard approaches, though they will slow down at some point. It may be that these methods offer the best route forward either until fully linear-scaling methods are developed or until better orbital-free functionals are found.

The third challenge, which is faced by codes in many other areas, is to make efficient use of new computing architectures, particularly given the shift towards petascale computing. Real-space methods in general are well placed to adapt to multi-core and GPU-based computing, but the communications patterns developed for previous generations of high-performance computing are not necessarily best suited for novel architectures. As linear-scaling methods often use specific approaches developed for the problem, it may be harder to adapt than for other approaches using standard packages. However, this is an area where linear-scaling codes can be extremely successful, and the effort should be made.

The fourth challenge is to improve functionality while maintaining linear-scaling behaviour. Recent years have seen DFT improved by adding features such as exact exchange and dispersion forces (also known as van der Waals interactions). Methods have been proposed to implement these with linear scaling, though, as always, adapting them to the approach used (and parallelizing while maintaining linear scaling and efficiency) poses problems. Time-dependent DFT is certainly possible with linear scaling, as are certain parts of quantum Monte Carlo and MP2 calculations, but approaches such as GW cannot be adapted in their present form (though the influence and portability of Wannier functions is making many interesting approaches viable). Embedding of more accurate methods into DFT (or approximate DFT methods) may become important, and linear-scaling approaches are ideal for this, starting from the locality of electronic structure.

The final, and in many ways most important, challenge is that of applications to large systems, as already highlighted. Long timescales pose a challenge to all atomistic simulation methods, and larger systems give longer length scales which typically are associated with slower response times. However, this is a generic problem. Weak forces such as hydrogen bonding and van der Waals interactions are also important in large systems, particularly biological systems. While there are semi-empirical and *ab initio* methods for calculating these forces, accuracy and testing are vitally important. For the problem of preparing input for, and analysing output from, calculations with millions of atoms, this community can learn valuable lessons from the molecular dynamics community, and use existing tools from that area. Finally, as this is a new field, it will take time to understand which physical properties can be calculated reliably and efficiently.

This survey of recent developments in linear-scaling approaches has shown that we stand at a fascinating point. In 1999, in a previous review of linear-scaling methods, Goedecker wrote, 'Even with $\mathcal{O}(N)$ algorithms it will not be possible in the foreseeable future to treat systems containing millions of atoms at a highly accurate density-functional level using large basis sets, as would be necessary for certain materials science applications' [2]. However, we now have the first, true DFT calculations on millions of atoms [411], and fully converged, highly accurate results on systems of this size will be obtained in the very near future. With this capability,

there is a rich variety of systems and phenomena which can be tackled with accurate, linear-scaling DFT techniques.

Acknowledgments

The authors gratefully acknowledge many years of support and encouragement from Professor Mike Gillan, and the hard work and contributions of the CONQUEST developers: Michiaki Arita, Veronika Brázdová, Ian Bush, Rathin Choudhury, Chris Goringe, Eduardo Hernández, Conn O'Rourke, Takao Otsuka, Alex Sena, Umberto Terranova, Milica Todorović, Lianheng Tong, Antonio Torralba and Lionel Truflandier. They are also grateful for many helpful suggestions and useful data from colleagues in the field, including Emilio Artacho, Emily Carter, Vikram Gavini, Stefan Goedecker, Peter Haynes, Nick Hine, Linda Hung, Sohrab Ismail-Beigi, Nicola Marzari, Anders Niklasson, Mark Rayson, Emanuel Rubensson, Elias Rudberg, Annabella Selloni, Chris-Kriton Skylaris, Jose Soler, David Vanderbilt, Lin-Wang Wang and Weitao Yang. This work was partly supported by the Royal Society, a Grant-in-Aid for Scientific Research on Innovative Areas (No. 22104005), the Strategic Programs for Innovative Research (SPIRE), MEXT, and the Computational Materials Science Initiative (CMSI), Japan.

References

[1] Car R and Parrinello M 1985 *Phys. Rev. Lett.* **55** 2471
[2] Goedecker S 1999 *Rev. Mod. Phys.* **71** 1085
[3] Kohn W and Sham L J 1965 *Phys. Rev.* **140** A1133
[4] Hamann D R, Schlüter M and Chiang C 1979 *Phys. Rev. Lett.* **43** 1494
[5] Bachelet G B, Hamann D R and Schlüter M 1982 *Phys. Rev.* B **26** 4199
[6] Rappe A M, Rabe K M, Kaxiras E and Joannopoulos J D 1990 *Phys. Rev.* B **41** 1227
[7] Troullier N and Martins J L 1991 *Phys. Rev.* B **43** 1993
[8] Vanderbilt D 1990 *Phys. Rev.* B **41** 7892
[9] Blöchl P E 1994 *Phys. Rev.* B **50** 17953
[10] Gervasio F L, Carloni P and Parrinello M 2002 *Phys. Rev. Lett.* **89** 108102
[11] Fattebert J L and Gygi F 2004 *Comput. Phys. Commun.* **162** 24
[12] Iwata J I, Takahashi D, Oshiyama A, Boku T, Shiraishi K, Okada S and Yabana K 2010 *J. Comput. Phys.* **229** 2339
[13] Hutter J and Curioni A 2005 *ChemPhysChem* **6** 1788
[14] Gygi F 2006 *J. Phys.: Conf. Ser.* **46** 268
[15] Goringe C M, Bowler D R and Hernández E 1997 *Rep. Prog. Phys.* **60** 1447
[16] Galli G and Mauri F 1994 *Phys. Rev. Lett.* **73** 3471
[17] Beck T L 2000 *Rev. Mod. Phys.* **72** 1041
[18] Torsti T *et al* 2006 *Phys. Status Solidi* b **243** 1016
[19] Saad Y, Chelikowsky J R and Shontz S M 2010 *SIAM Rev.* **52** 3
[20] Bernstein N, Kermode J R and Csányi G 2009 *Rep. Prog. Phys.* **72** 026501
[21] Senn H M and Thiel W 2009 *Angew. Chem. Int. Edn* **48** 1198
[22] Pask J E and Sterne P A 2005 *Modelling Simul. Mater. Sci. Eng.* **13** R71
[23] Briggs E L, Sullivan D J and Bernholc J 1996 *Phys. Rev.* B **54** 14362
[24] Fattebert J L and Bernholc J 2000 *Phys. Rev.* B **62** 1713
[25] Maragakis P, Soler J and Kaxiras E 2001 *Phys. Rev.* B **64** 193101
[26] Skylaris C K, Diéguez O, Haynes P D and Payne M C 2002 *Phys. Rev.* B **66** 073103
[27] Bernholc J, Hodak M and Lu W 2008 *J. Phys.: Condens. Matter* **20** 294205
[28] Chelikowsky J R, Troullier N and Saad Y 1994 *Phys. Rev. Lett.* **72** 1240
[29] Zhou Y, Saad Y, Tiago M L and Chelikowsky J R 2006 *Phys. Rev.* E **74** 066704
[30] Alemany M, Jain M, Tiago M L, Zhou Y, Saad Y and Chelikowsky J R 2007 *Comput. Phys. Commun.* **177** 339
[31] Iwata J I, Oshiyama A and Shiraishi K 2006 *Physica* B **376–377** 196
[32] Fujimoto Y and Oshiyama A 2010 *Phys. Rev.* B **81** 205309
[33] Heiskanen M, Torsti T, Puska M J and Nieminen R M 2001 *Phys. Rev.* B **63** 245106
[34] Mortensen J J, Hansen L B and Jacobsen K W 2005 *Phys. Rev.* B **71** 035109
[35] Press W H, Flannery B P, Teukolsky S A and Vetterling W T 1992 *Numerical Recipes in FORTRAN: The Art of Scientific Computing* 2nd edn (Cambridge: Cambridge University Press)
[36] Hoshi T and Fujiwara T 1997 *J. Phys. Soc. Japan* **66** 3710
[37] Fattebert J L 2008 *J. Phys.: Condens. Matter* **20** 294210
[38] White S R, Wilkins J W and Teter M P 1989 *Phys. Rev.* B **39** 5819
[39] Tsuchida E and Tsukada M 1995 *Phys. Rev.* B **52** 5573
[40] Pask J E, Klein B M, Fong C Y and Sterne P A 1999 *Phys. Rev.* B **59** 12352
[41] Hernández E, Gillan M J and Goringe C M 1997 *Phys. Rev.* B **55** 13485
[42] Tsuchida E and Tsukada M 1996 *Phys. Rev.* B **54** 7602
[43] Tsuchida E and Tsukada M 1998 *J. Phys. Soc. Japan* **67** 3844
[44] Tsuchida E 2008 *J. Phys.: Condens. Matter* **20** 294212
[45] Wijesekera N R, Feng G and Beck T L 2007 *Phys. Rev.* B **75** 115101
[46] Feng G and Beck T L 2006 *Phys. Status Solidi* b **243** 1054
[47] Genovese L *et al* 2008 *J. Chem. Phys.* **129** 014109
[48] Arias T A 1999 *Rev. Mod. Phys.* **71** 267
[49] Niklasson A M N, Tymczak C J and Röder H 2002 *Phys. Rev.* B **66** 155120
[50] Liu Y, Yarne D A and Tuckerman M E 2003 *Phys. Rev.* B **68** 125110
[51] Lee H S and Tuckerman M E 2006 *J. Phys. Chem.* A **110** 5549
[52] Mostofi A A, Skylaris C K, Haynes P D and Payne M C 2002 *Comput. Phys. Commun.* **147** 788
[53] Hoshi T, Arai M and Fujiwara T 1995 *Phys. Rev.* B **52** R5459
[54] Varga K, Zhang Z and Pantelides S T 2004 *Phys. Rev. Lett.* **93** 176403
[55] Varga K and Pantelides S T 2006 *Phys. Status Solidi* b **243** 1110
[56] Alfè D and Gillan M J 2004 *J. Phys.: Condens. Matter* **16** L305
[57] E W, Li T and Lu J 2010 *Proc. Natl Acad. Sci.* **107** 1273
[58] Galli G and Parrinello M 1992 *Phys. Rev. Lett.* **69** 3547
[59] Soler J M, Artacho E, Gale J D, García A, Junquera J, Ordejón P and Sánchez-Portal D 2002 *J. Phys.: Condens. Matter* **14** 2745
[60] Blum V, Gehrke R, Hanke F, Havu P, Havu V, Ren X, Reuter K and Scheffler M 2009 *Comput. Phys. Commun.* **180** 2175
[61] VandeVondele J, Krack M, Mohamed F, Parrinello M, Chassaing T and Hutter J 2005 *Comput. Phys. Commun.* **167** 103
[62] Ozaki T 2003 *Phys. Rev.* B **67** 155108
[63] Ozaki T and Kino H 2004 *J. Chem. Phys.* **121** 10879
[64] Ozaki T and Kino H 2004 *Phys. Rev.* B **69** 195113
[65] Kenny S D, Horsfield A P and Fujitani H 2000 *Phys. Rev.* B **62** 4899

[66] Torralba A S, Todorovic M, Brázdová V, Choudhury R, Miyazaki T, Gillan M J and Bowler D R 2008 *J. Phys.: Condens. Matter* **20** 294206
[67] Otsuka T, Miyazaki T, Ohno T, Bowler D R and Gillan M J 2008 *J. Phys.: Condens. Matter* **20** 294201
[68] Briddon P and Jones R 2000 *Phys. Status Solidi* b **217** 131
[69] VandeVondele J and Hutter J 2007 *J. Chem. Phys.* **127** 114105
[70] Shang H, Xiang H, Li Z and Yang J 2010 *Int. Rev. Phys. Chem.* **29** 665
[71] Sankey O F and Niklewski D J 1989 *Phys. Rev.* B **40** 3979
[72] Junquera J, Paz O, Sánchez-Portal D and Artacho E 2001 *Phys. Rev.* B **64** 235111
[73] Horsfield A P and Bratkovsky A M 2000 *J. Phys.: Condens. Matter* **12** R1
[74] Porezag D, Frauenheim T, Köhler T, Seifert G and Kaschner R 1995 *Phys. Rev.* B **51** 12947
[75] Horsfield A P 1997 *Phys. Rev.* B **56** 6594
[76] Artacho E, Sánchez-Portal D, Ordejon P, Garcia A and Soler J M 1999 *Phys. Status Solidi* b **215** 809
[77] Anglada E, M. Soler J, Junquera J and Artacho E 2002 *Phys. Rev.* B **66** 205101
[78] Esfarjani K, Farajian A A and Kawazoe Y 1999 *Comput. Mater. Sci.* **15** 351
[79] Sánchez-Portal D, Artacho E and Soler J M 1995 *Solid State Commun.* **95** 685
[80] Sánchez-Portal D, Artacho E and Soler J M 1996 *J. Phys.: Condens. Matter.* **8** 3859
[81] Chen M, Guo G C and He L 2010 *J. Phys.: Condens. Matter* **22** 445501
[82] Talman J D 2000 *Phys. Rev. Lett.* **84** 855
[83] Haynes P D and Payne M C 1997 *Comput. Phys. Commun.* **102** 17
[84] Monserrat B and Haynes P D 2010 *J. Phys. A: Math. Gen.* **43** 465205
[85] Porezag D and Pederson M R 1999 *Phys. Rev.* A **60** 2840
[86] Lee M S and Head-Gordon M 1997 *J. Chem. Phys.* **107** 9085
[87] Lee M S and Head-Gordon M 2000 *Int. J. Quantum Chem.* **76** 169
[88] Berghold G, Parrinello M and Hutter J 2002 *J. Chem. Phys.* **116** 1800
[89] Aquilante F, Pedersen T B, de Merás A S and Koch H 2006 *J. Chem. Phys.* **125** 174101
[90] Basanta M, Dappe Y, Jelínek P and Ortega J 2007 *Comput. Mater. Sci.* **39** 759
[91] Gusso M 2008 *J. Chem. Phys.* **128** 044102
[92] Rayson M J and Briddon P R 2009 *Phys. Rev.* B **80** 205104
[93] Rayson M 2010 *Comput. Phys. Commun.* **181** 1051
[94] Boys S and Bernardi F 1970 *Mol. Phys.* **19** 553
[95] Artacho E, Machado M, Sánchez-Portal D, Ordejón P and Soler J M 2003 *Mol. Phys.* **101** 1587
[96] Skylaris C K and Haynes P D 2007 *J. Chem. Phys.* **127** 164712
[97] Haynes P, Skylaris C K, Mostofi A and Payne M 2006 *Chem. Phys. Lett.* **422** 345
[98] Challacombe M and Schwegler E 1997 *J. Chem. Phys.* **106** 5526
[99] Watson M A, Sałek P, Macak P and Helgaker T 2004 *J. Chem. Phys.* **121** 2915
[100] Shao Y *et al* 2006 *Phys. Chem. Chem. Phys.* **8** 3172
[101] Ozaki T and Kino H 2005 *Phys. Rev.* B **72** 045121
[102] Challacombe M, Schwegler E and Almlöf J 1996 *J. Chem. Phys.* **104** 4685
[103] Gan C K, Tymczak C J and Challacombe M 2004 *J. Chem. Phys.* **121** 6608
[104] Rudberg E, Rubensson E H and Sałek P 2011 *J. Chem. Theory Comput.* **7** 340
[105] Sodt A, Subotnik J E and Head-Gordon M 2006 *J. Chem. Phys.* **125** 194109
[106] Genovese L, Deutsch T and Goedecker S 2007 *J. Chem. Phys.* **127** 054704
[107] Watson M A, Kurashige Y, Nakajima T and Hirao K 2008 *J. Chem. Phys.* **128** 054105
[108] Challacombe M 2000 *J. Chem. Phys.* **113** 10037
[109] Balbás L C, Martins J L and Soler J M 2001 *Phys. Rev.* B **64** 165110
[110] Torralba A S, Bowler D R, Miyazaki T and Gillan M J 2009 *J. Chem. Theory Comput.* **5** 1499
[111] Schwegler E and Challacombe M 1996 *J. Chem. Phys.* **105** 2726
[112] Burant J C, Scuseria G E and Frisch M J 1996 *J. Chem. Phys.* **105** 8969
[113] Schwegler E, Challacombe M and Head-Gordon M 1997 *J. Chem. Phys.* **106** 9708
[114] Ochsenfeld C, White C A and Head-Gordon M 1998 *J. Chem. Phys.* **109** 1663
[115] Schwegler E and Challacombe M 1999 *J. Chem. Phys.* **111** 6223
[116] Lambrecht D S and Ochsenfeld C 2005 *J. Chem. Phys.* **123** 184101
[117] Weber V and Challacombe M 2006 *J. Chem. Phys.* **125** 104110
[118] Rudberg E, Rubensson E H and Sałek P 2008 *J. Chem. Phys.* **128** 184106
[119] Wu X, Selloni A and Car R 2009 *Phys. Rev.* B **79** 085102
[120] Goedecker S and Colombo L 1994 *Phys. Rev. Lett.* **73** 122
[121] McWeeny R 1960 *Rev. Mod. Phys.* **32** 335
[122] Wannier G H 1937 *Phys. Rev.* **52** 191
[123] Kohn W 1959 *Phys. Rev.* **115** 809
[124] Kohn W 1973 *Phys. Rev.* B **7** 4388
[125] Blount E 1962 *Solid State Phys.* **13** 305
[126] Kohn W 1995 *Int. J. Quantum Chem.* **56** 229
[127] Kohn W 1964 *Phys. Rev.* **133** A171
[128] Adams W H 1971 *Chem. Phys. Lett.* **11** 71
[129] Weeks J D, Anderson P W and Davidson A G H 1973 *J. Chem. Phys.* **58** 1388
[130] des Cloizeaux J 1964 *Phys. Rev.* **135** A685
[131] des Cloizeaux J 1964 *Phys. Rev.* **135** A698
[132] Nenciu G 1983 *Commun. Math. Phys.* **91** 81
[133] Kohn W and Onffroy J R 1973 *Phys. Rev.* B **8** 2485
[134] Nenciu A and Nenciu G 1993 *Phys. Rev.* B **47** 10112
[135] des Cloizeaux J 1963 *Phys. Rev.* **129** 554
[136] He L and Vanderbilt D 2001 *Phys. Rev. Lett.* **86** 5341
[137] Marzari N and Vanderbilt D 1997 *Phys. Rev.* B **56** 12847
[138] Resta R 2002 *J. Phys.: Condens. Matter* **14** R625
[139] Brouder C, Panati G, Calandra M, Mourougane C and Marzari N 2007 *Phys. Rev. Lett.* **98** 046402
[140] Hine N D M and Foulkes W M C 2007 *J. Phys.: Condens. Matter* **19** 506212
[141] Koch E and Goedecker S 2001 *Solid State Commun.* **119** 105
[142] Annett J F 1995 *Comput. Mater. Sci.* **4** 23
[143] Kresse G and Furthmüller J 1996 *Phys. Rev.* B **54** 11169
[144] Resta R 1994 *Rev. Mod. Phys.* **66** 899
[145] Foster J M and Boys S F 1960 *Rev. Mod. Phys.* **32** 300
[146] Vanderbilt D and King-Smith R D 1993 *Phys. Rev.* B **48** 4442
[147] Stephan U, Martin R M and Drabold D A 2000 *Phys. Rev.* B **62** 6885
[148] Resta R and Sorella S 1999 *Phys. Rev. Lett.* **82** 370
[149] Souza I, Wilkens T and Martin R M 2000 *Phys. Rev.* B **62** 1666
[150] Kramer B and MacKinnon A 1993 *Rep. Prog. Phys.* **56** 1469
[151] Anderson P W 1968 *Phys. Rev. Lett.* **21** 13
[152] Kohn W 1993 *Chem. Phys. Lett.* **208** 167
[153] Silvestrelli P L, Marzari N, Vanderbilt D and Parrinello M 1998 *Solid State Commun.* **107** 7
[154] Souza I, Marzari N and Vanderbilt D 2001 *Phys. Rev.* B **65** 035109
[155] Birkenheuer U and Izotov D 2005 *Phys. Rev.* B **71** 125116

[156] Berghold G, Mundy C J, Romero A H, Hutter J and Parrinello M 2000 *Phys. Rev.* B **61** 10040
[157] Thygesen K S, Hansen L B and Jacobsen K W 2005 *Phys. Rev. Lett.* **94** 026405
[158] Stephan U and Drabold D A 1998 *Phys. Rev.* B **57** 6391
[159] Hoshi T and Fujiwara T 2000 *J. Phys. Soc. Japan* **69** 3773
[160] Geshi M, Hoshi T and Fujiwara T 2003 *J. Phys. Soc. Japan* **72** 2880
[161] Xiang H J, Li Z, Liang W Z, Yang J, Hou J G and Zhu Q 2006 *J. Chem. Phys.* **124** 234108
[162] Iftimie R, Thomas J W and Tuckerman M E 2004 *J. Chem. Phys.* **120** 2169
[163] Mauri F, Galli G and Car R 1993 *Phys. Rev.* B **47** 9973
[164] Kim J, Mauri F and Galli G 1995 *Phys. Rev.* B **52** 1640
[165] Ordejón P, Drabold D A, Grumbach M P and Martin R M 1993 *Phys. Rev.* B **48** 14646
[166] Ordejón P, Drabold D A, Martin R M and Grumbach M P 1995 *Phys. Rev.* B **51** 1456
[167] Nunes R W and Vanderbilt D 1994 *Phys. Rev. Lett.* **73** 712
[168] O'Regan D D, Hine N D M, Payne M C and Mostofi A A 2010 *Phys. Rev.* B **82** 081102
[169] Liu S, Pérez-Jordá J M and Yang W 2000 *J. Chem. Phys.* **112** 1634
[170] Feng H, Bian J, Li L and Yang W 2004 *J. Chem. Phys.* **120** 9458
[171] Cui G, Fang W and Yang W 2010 *J. Phys. Chem.* A **114** 8878
[172] Lu W C, Wang C Z, Chan T L, Ruedenberg K and Ho K M 2004 *Phys. Rev.* B **70** 041101
[173] Subotnik J E and Head-Gordon M 2005 *J. Chem. Phys.* **122** 034109
[174] Ismail-Beigi S and Arias T A 1999 *Phys. Rev. Lett.* **82** 2127
[175] Kohn W 1996 *Phys. Rev. Lett.* **76** 3168
[176] Baer R and Head-Gordon M 1997 *Phys. Rev. Lett.* **79** 3962
[177] Goedecker S 1998 *Phys. Rev.* B **58** 3501
[178] Goedecker S and Ivanov O V 1999 *Phys. Rev.* B **59** 7270
[179] Prodan E and Kohn W 2005 *Proc. Natl Acad. Sci. USA* **102** 11635
[180] Sacksteder V E 2005 *Numer. Linear Algebra Appl.* **12** 827
[181] Rubensson E H and Rudberg E 2011 *J. Comput. Chem.* **32** 1411
[182] Hastings M B 2004 *Phys. Rev. Lett.* **93** 126402
[183] Taraskin S N, Drabold D A and Elliott S R 2002 *Phys. Rev. Lett.* **88** 196405
[184] Taraskin S N, Fry P A, Zhang X, Drabold D A and Elliott S R 2002 *Phys. Rev.* B **66** 233101
[185] Zhang X and Drabold D A 2001 *Phys. Rev.* B **63** 233109
[186] Jędrzejewski J and Krokhmalskii T 2004 *Phys. Rev.* B **70** 153102
[187] Jędrzejewski J and Krokhmalskii T 2005 *Phys. Rev.* B **72** 233106
[188] Jędrzejewski J and Krokhmalskii T 2007 *Europhys. Lett.* **78** 37002
[189] Bowler D R, Aoki M, Goringe C M, Horsfield A P and Pettifor D G 1997 *Modelling Simul. Mater. Sci. Eng.* **5** 199
[190] Lippert R A and Sears M P 2000 *Phys. Rev.* B **61** 12772
[191] Jordan D K and Mazziotti D A 2005 *J. Chem. Phys.* **122** 084114
[192] Rudberg E and Rubensson E H 2011 *J. Phys.: Condens. Matter* **23** 075502
[193] Ordejón P 1998 *Comput. Mater. Sci.* **12** 157
[194] Galli G 1996 *Curr. Opin. Solid State Mater. Sci.* **1** 864
[195] Galli G 2000 *Phys. Status Solidi* b **217** 231
[196] Stephan U 2000 *Phys. Rev.* B **62** 16412
[197] Wu S Y and Jayanthi C S 2002 *Phys. Rep.* **358** 1
[198] Scuseria G E 1999 *J. Phys. Chem.* A **103** 4782
[199] Hierse W and Stechel E B 1994 *Phys. Rev.* B **50** 17811
[200] Hernández E and Gillan M J 1995 *Phys. Rev.* B **51** 10157
[201] Li X P, Nunes R W and Vanderbilt D 1993 *Phys. Rev.* B **47** 10891
[202] Fattebert J and Gygi F 2006 *Phys. Rev.* B **73** 115124
[203] Daw M S 1993 *Phys. Rev.* B **47** 10895
[204] Nunes R W and Vanderbilt D 1994 *Phys. Rev.* B **50** 17611
[205] Hernández E, Gillan M J and Goringe C M 1996 *Phys. Rev.* B **53** 7147
[206] Daniels A D, Millam J M and Scuseria G E 1997 *J. Chem. Phys.* **107** 425
[207] Millam J M and Scuseria G E 1997 *J. Chem. Phys.* **106** 5569
[208] Bowler D R and Gillan M J 1999 *Comput. Phys. Commun.* **120** 95
[209] Challacombe M 1999 *J. Chem. Phys.* **110** 2332
[210] Head-Gordon M, Shao Y, Saravanan C and White C 2003 *Mol. Phys.* **101** 37
[211] Li X, Millam J M, Scuseria G E, Frisch M J and Schlegel H B 2003 *J. Chem. Phys.* **119** 7651
[212] Haynes P D, Skylaris C K, Mostofi A A and Payne M C 2008 *J. Phys.: Condens. Matter* **20** 294207
[213] Corkill J L and Ho K M 1996 *Phys. Rev.* B **54** 5340
[214] Palser A H R and Manolopoulos D E 1998 *Phys. Rev.* B **58** 12704
[215] Carlsson A E 1995 *Phys. Rev.* B **51** 13935
[216] Goringe C M, Hernández E, Gillan M J and Bush I J 1997 *Comput. Phys. Commun.* **102** 1
[217] Bowler D R, Miyazaki T and Gillan M J 2002 *J. Phys.: Condens. Matter* **14** 2781
[218] Schlegel H B, Millam J M, Iyengar S S, Voth G A, Daniels A D, Scuseria G E and Frisch M J 2001 *J. Chem. Phys.* **114** 9758
[219] Li X, Moss C, Liang W and Feng Y 2009 *J. Chem. Phys.* **130** 234115
[220] Beylkin G, Coult N and Mohlenkamp M J 1999 *J. Comput. Phys.* **152** 32
[221] Nemeth K and Scuseria G 2000 *J. Chem. Phys.* **113** 6035
[222] Niklasson A M N 2003 *Phys. Rev.* B **68** 233104
[223] Kryachko E S 2000 *Chem. Phys. Lett.* **318** 210
[224] Holas A 2001 *Chem. Phys. Lett.* **340** 552
[225] Habershon S and Manby F R 2002 *Chem. Phys. Lett.* **354** 527
[226] Stechel E B, Williams A R and Feibelman P J 1994 *Phys. Rev.* B **49** 10088
[227] Niklasson A M N 2002 *Phys. Rev.* B **66** 155115
[228] Xiang H J, Liang W Z, Yang J, Hou J G and Zhu Q 2005 *J. Chem. Phys.* **123** 124105
[229] Niklasson A M N, Weber V and Challacombe M 2005 *J. Chem. Phys.* **123** 044107
[230] Mazziotti D A 2001 *J. Chem. Phys.* **115** 8305
[231] Mazziotti D A 2003 *Phys. Rev.* E **68** 066701
[232] Kőhalmi D, Szabados A and Surján P R 2005 *Phys. Rev. Lett.* **95** 013002
[233] Szakács P and Surján P 2008 *J. Math. Chem.* **43** 314
[234] Niklasson A M N, Tymczak C J and Challacombe M 2003 *J. Chem. Phys.* **118** 8611
[235] Rubensson E H and Sałek P 2005 *J. Comput. Chem.* **26** 1628
[236] Rubensson E H, Rudberg E and Sałek P 2008 *J. Chem. Phys.* **128** 074106
[237] Tsuchida E 2007 *J. Phys. Soc. Japan* **76** 034708
[238] Mauri F and Galli G 1994 *Phys. Rev.* B **50** 4316
[239] Yang W 1997 *Phys. Rev.* B **56** 9294
[240] Sternberg M, Galli G and Frauenheim T 1999 *Comput. Phys. Commun.* **118** 200
[241] Sasaki T, Ono T and Hirose K 2006 *Phys. Rev.* E **74** 056704
[242] Itoh S, Ordejón P and Martin R M 1995 *Comput. Phys. Commun.* **88** 173
[243] Shellman S D, Lewis J P, Glaesemann K R, Sikorski K and Voth G A 2003 *J. Comput. Phys.* **188** 1
[244] Stephan U, Drabold D A and Martin R M 1998 *Phys. Rev.* B **58** 13472
[245] Yang W 1991 *Phys. Rev. Lett.* **66** 1438
[246] Yang W 1991 *Phys. Rev.* A **44** 7823
[247] Yang W 1992 *J. Mol. Struct.: Theochem.* **255** 461

[248] Cortona P 1991 *Phys. Rev.* B **44** 8454
[249] Zhu T, Pan W and Yang W 1996 *Phys. Rev.* B **53** 12713
[250] Zhao Q and Yang W 1995 *J. Chem. Phys.* **102** 9598
[251] Kobayashi M, Kunisada T, Akama T, Sakura D and Nakai H 2011 *J. Chem. Phys.* **134** 034105
[252] Yang W and Lee T S 1995 *J. Chem. Phys.* **103** 5674
[253] Dixon S L and Merz K M Jr 1996 *J. Chem. Phys.* **104** 6643
[254] Ozaki T 2006 *Phys. Rev.* B **74** 245101
[255] Cankurtaran B O, Gale J D and Ford M J 2008 *J. Phys.: Condens. Matter* **20** 294208
[256] Shimojo F, Kalia R K, Nakano A and Vashishta P 2005 *Comput. Phys. Commun.* **167** 151
[257] Nakano A, Kalia R K, Nomura K, Sharma A, Vashishta P, Shimojo F, van Duin A C T, Goddard W A, Biswas R and Srivastava D 2007 *Comput. Mater. Sci.* **38** 642
[258] Shimojo F, Kalia R K, Nakano A and Vashishta P 2008 *Phys. Rev.* B **77** 085103
[259] Yao Y, Wang C, Zhang G, Ji M and Ho K 2009 *J. Phys.: Condens. Matter* **21** 235501
[260] Kobayashi M and Nakai H 2008 *J. Chem. Phys.* **129** 044103
[261] Ziółkowski M, Jansík B, Kjærgaard T and Jørgensen P 2010 *J. Chem. Phys.* **133** 014107
[262] Akama T, Kobayashi M and Nakai H 2007 *J. Comput. Chem.* **28** 2003
[263] Elliott P, Cohen M H, Wasserman A and Burke K 2009 *J. Chem. Theor. Comput.* **5** 827
[264] Elliott P, Burke K, Cohen M H and Wasserman A 2010 *Phys. Rev.* A **82** 024501
[265] Wang L W, Zhao Z and Meza J 2008 *Phys. Rev.* B **77** 165113
[266] Zhao Z, Meza J and Wang L W 2008 *J. Phys.: Condens. Matter* **20** 294203
[267] Wang L W 2002 *Phys. Rev. Lett.* **88** 256402
[268] Vukmirovic N and Wang L W 2008 *J. Chem. Phys.* **128** 121102
[269] Seijo L and Barandiarán Z 2004 *J. Chem. Phys.* **121** 6698
[270] Seijo L, Barandiarán Z and Soler J 2007 *Theor. Chim. Acta* **118** 541
[271] Kitaura K, Ikeo E, Asada T, Nakano T and Uebayasi M 1999 *Chem. Phys. Lett.* **313** 701
[272] Fedorov D G and Kitaura K 2007 *J. Phys. Chem.* A **111** 6904
[273] Ganesh V, Dongare R, Balanarayan P and Gadre S 2006 *J. Chem. Phys.* **125** 104109
[274] Szekeres Z, Mezey P G and Surján P R 2006 *Chem. Phys. Lett.* **424** 420
[275] Varga K 2010 *Phys. Rev.* B **81** 045109
[276] Ozaki T 2010 *Phys. Rev.* B **82** 075131
[277] Wang Y, Stocks G M, Shelton W A, Nicholson D M C, Szotek Z and Temmerman W M 1995 *Phys. Rev. Lett.* **75** 2867
[278] Abrikosov I A, Niklasson A M N, Simak S I, Johansson B, Ruban A V and Skriver H L 1996 *Phys. Rev. Lett.* **76** 4203
[279] Abrikosov I A, Simak S I, Johansson B, Ruban A V and Skriver H L 1997 *Phys. Rev.* B **56** 9319
[280] Zeller R, Dederichs P H, Újfalussy B, Szunyogh L and Weinberger P 1995 *Phys. Rev.* B **52** 8807
[281] Zeller R 1997 *Phys. Rev.* B **55** 9400
[282] Smirnov A V and Johnson D D 2001 *Phys. Rev.* B **64** 235129
[283] Smirnov A V and Johnson D D 2002 *Comput. Phys. Commun.* **148** 74
[284] Zeller R 2008 *J. Phys.: Condens. Matter* **20** 294215
[285] Ogura M and Akai H 2009 *J. Comput. Theor. Nanosci.* **6** 2483
[286] Heine V 1980 *Solid State Phys.* **35** 1
[287] Haydock R 1980 *Solid State Phys.* **35** 215
[288] Lanczos C 1950 *J. Res. Natl Bur. Stand.* **45** 255
[289] Ballentine L E and Kolar M 1986 *J. Phys. C: Solid State Phys.* **19** 981
[290] Baroni S and Giannozzi P 1992 *Europhys. Lett.* **17** 547
[291] Aoki M 1993 *Phys. Rev. Lett.* **71** 3842
[292] Horsfield A P, Bratkovsky A M, Pettifor D G and Aoki M 1996 *Phys. Rev.* B **53** 1656
[293] Horsfield A P, Bratkovsky A M, Fearn M, Pettifor D G and Aoki M 1996 *Phys. Rev.* B **53** 12694
[294] Ozaki T 1999 *Phys. Rev.* B **59** 16061
[295] Ozaki T, Aoki M and Pettifor D G 2000 *Phys. Rev.* B **61** 7972
[296] Ozaki T and Terakura K 2001 *Phys. Rev.* B **64** 195126
[297] Pettifor D G, Oleinik I I, Nguyen-Manh D and Vitek V 2002 *Comput. Mater. Sci.* **23** 33
[298] Goedecker S and Teter M 1995 *Phys. Rev.* B **51** 9455
[299] Goedecker S 1995 *J. Comput. Phys.* **118** 261
[300] Jay L O, Kim H, Saad Y and Chelikowsky J R 1999 *Comput. Phys. Commun.* **118** 21
[301] Silver R N, Roeder H, Voter A F and Kress J D 1996 *J. Comput. Phys.* **124** 115
[302] Voter A F, Kress J D and Silver R N 1996 *Phys. Rev.* B **53** 12733
[303] Röder H, Silver R N, Drabold D A and Dong J J 1997 *Phys. Rev.* B **55** 15382
[304] Liang W, Saravanan C, Shao Y, Baer R, Bell A T and Head-Gordon M 2003 *J. Chem. Phys.* **119** 4117
[305] Garcia-Cervera C, Lu J, Xuan Y and E W 2009 *Phys. Rev.* B **79** 115110
[306] Drabold D A and Sankey O P 1993 *Phys. Rev. Lett.* **70** 3631
[307] Sankey O F, Drabold D A and Gibson A 1994 *Phys. Rev.* B **50** 1376
[308] Wang L W 1994 *Phys. Rev.* B **49** 10154
[309] Krajewski F R and Parrinello M 2005 *Phys. Rev.* B **71** 233105
[310] Krajewski F R and Parrinello M 2006 *Phys. Rev.* B **73** 041105
[311] Krajewski F R and Parrinello M 2006 *Phys. Rev.* B **74** 125107
[312] Krajewski F R and Parrinello M 2007 *Phys. Rev.* B **75** 235108
[313] Ceriotti M, Kühne T D and Parrinello M 2008 *J. Chem. Phys.* **129** 024707
[314] Kenoufi A and Polonyi J 2004 *Phys. Rev.* B **70** 205105
[315] Rayson M 2007 *Phys. Rev.* B **75** 153203
[316] Gygi F 2009 *Phys. Rev. Lett.* **102** 166406
[317] Takayama R, Hoshi T and Fujiwara T 2004 *J. Phys. Soc. Japan* **73** 1519
[318] Takayama R, Hoshi T, Sogabe T, Zhang S and Fujiwara T 2006 *Phys. Rev.* B **73** 165108
[319] Hoshi T and Fujiwara T 2006 *J. Phys.: Condens. Matter* **18** 10787
[320] Fujiwara T, Hoshi T and Yamamoto S 2008 *J. Phys.: Condens. Matter* **20** 294202
[321] Teng H, Fujiwara T, Hoshi T, Sogabe T, Zhang S L and Yamamoto S 2011 *Phys. Rev.* B **83** 165103
[322] Haynes P D and Payne M C 1998 *Solid State Commun.* **108** 737
[323] Haynes P D and Payne M C 1999 *Phys. Rev.* B **59** 12173
[324] Adhikari S and Baer R 2001 *J. Chem. Phys.* **115** 11
[325] Wang L W and Teter M P 1992 *Phys. Rev.* B **46** 12798
[326] Wang L W and Teter M P 1992 *Phys. Rev.* B **45** 13196
[327] Pearson M, Smargiassi E and Madden P A 1993 *J. Phys.: Condens. Matter* **5** 3221
[328] García-Aldea D and Alvarellos J E 2007 *J. Chem. Phys.* **127** 144109
[329] Watson S, Jesson B J, Carter E A and Madden P A 1998 *Europhys. Lett.* **41** 37
[330] Smargiassi E and Madden P A 1994 *Phys. Rev.* B **49** 5220
[331] Foley M and Madden P A 1996 *Phys. Rev.* B **53** 10589
[332] Wang Y A, Govind N and Carter E A 1998 *Phys. Rev.* B **58** 13465
[333] Wang Y A, Govind N and Carter E A 1999 *Phys. Rev.* B **60** 16350
[334] Chai J D and Weeks J 2007 *Phys. Rev.* B **75** 205122
[335] Ho G, Ligneres V and Carter E A 2008 *Phys. Rev.* B **78** 045105
[336] Choly N and Kaxiras E 2002 *Solid State Commun.* **121** 281

[337] Gavini V, Bhattacharya K and Ortiz M 2007 *J. Mech. Phys. Solids* **55** 697
[338] Zhou B, Ligneres V L and Carter E A 2005 *J. Chem. Phys.* **122** 044103
[339] Huang C and Carter E A 2010 *Phys. Rev.* B **81** 045206
[340] Ho G, Ligneres V and Carter E A 2008 *Comput. Phys. Commun.* **179** 839
[341] Hung L and Carter E A 2009 *Chem. Phys. Lett.* **475** 163
[342] Cangi A, Lee D, Elliott P, Burke K and Gross E K U 2011 *Phys. Rev. Lett.* **106** 236404
[343] Helgaker T, Larsen H, Olsen J and Jørgensen P 2000 *Chem. Phys. Lett.* **327** 397
[344] Larsen H, Olsen J, Jørgensen P and Helgaker T 2001 *J. Chem. Phys.* **115** 9685
[345] Van Voorhis T and Head-Gordon M 2002 *Mol. Phys.* **100** 1713
[346] Shao Y, Saravanan C, Head-Gordon M and White C A 2003 *J. Chem. Phys.* **118** 6144
[347] Herbert J M and Head-Gordon M 2004 *J. Chem. Phys.* **121** 11542
[348] Hutter J, Parrinello M and Vogel S 1994 *J. Chem. Phys.* **101** 3862
[349] Raczkowski D, Fong C Y, Schultz P A, Lippert R A and Stechel E B 2001 *Phys. Rev.* B **64** 155203
[350] Weber V and Hutter J 2008 *J. Chem. Phys.* **128** 064107
[351] Vandevondele J and Hutter J 2003 *J. Chem. Phys.* **118** 4365
[352] Weber V, VandeVondele J, Hutter J and Niklasson A M N 2008 *J. Chem. Phys.* **128** 084113
[353] Niklasson A M N 2004 *Phys. Rev.* B **70** 193102
[354] Sherrill C D 2010 *J. Chem. Phys.* **132** 110902
[355] Ochsenfeld C, Kussmann J and Lambrecht D S 2007 *Reviews in Computational Chemistry* vol 23 ed K B Lipkowitz and T R Cundari (New York: Wiley) pp 1–82 chapter 1
[356] Ayala P Y and Scuseria G E 1999 *J. Chem. Phys.* **110** 3660
[357] Doser B, Lambrecht D, Kussmann J and Ochsenfeld C 2009 *J. Chem. Phys.* **130** 064107
[358] Schütz M, Hetzer G and Werner H J 1999 *J. Chem. Phys.* **111** 5691
[359] Werner H J and Manby F 2006 *J. Chem. Phys.* **124** 054114
[360] Lee M S, Maslen P E and Head-Gordon M 2000 *J. Chem. Phys.* **112** 3592
[361] Subotnik J, Sodt A and Head-Gordon M 2006 *J. Chem. Phys.* **125** 074116
[362] Flocke N and Bartlett R J 2004 *J. Chem. Phys.* **121** 10935
[363] Chwee T, Szilva A, Lindh R and Carter E A 2008 *J. Chem. Phys.* **128** 224106
[364] Kobayashi M, Imamura Y and Nakai H 2007 *J. Chem. Phys.* **127** 074103
[365] Yam C, Yokojima S and Chen G 2003 *Phys. Rev.* B **68** 153105
[366] Yokojima S and Chen G 1998 *Chem. Phys. Lett.* **292** 379
[367] Yokojima S and Chen G 1999 *Chem. Phys. Lett.* **300** 540
[368] Yokojima S and Chen G 1999 *Phys. Rev.* B **59** 7259
[369] Wang F, Yam C Y, Chen G, Wang X, Fan K, Niehaus T A and Frauenheim T 2007 *Phys. Rev.* B **76** 045114
[370] Wang F, Yam C Y, Chen G and Fan K 2007 *J. Chem. Phys.* **126** 134104
[371] Iitaka T, Nomura S, Hirayama H, Zhao X, Aoyagi Y and Sugano T 1997 *Phys. Rev.* E **56** 1222
[372] Iitaka T and Ebisuzaki T 2000 *Phys. Rev.* E **61** R3314
[373] Nomura S, Iitaka T, Zhao X, Sugano T and Aoyagi Y 1999 *Phys. Rev.* B **59** 10309
[374] Tretiak S, Isborn C M, Niklasson A M N and Challacombe M 2009 *J. Chem. Phys.* **130** 054111
[375] Lucero M J, Niklasson A M N, Tretiak S and Challacombe M 2008 *J. Chem. Phys.* **129** 064114
[376] Challacombe M 2010 arXiv:1001.2586
[377] Baroni S, Gebauer R, Malcıoğlu O B, Saad Y, Umari P and Xian J 2010 *J. Phys.: Condens. Matter* **22** 074204
[378] Goncharov V A and Varga K 2011 *Phys. Rev.* B **83** 035118
[379] Niklasson A M N and Challacombe M 2004 *Phys. Rev. Lett.* **92** 193001
[380] Weber V, Niklasson A M N and Challacombe M 2004 *Phys. Rev. Lett.* **92** 193002
[381] Xiang H J, Yang J, Hou J G and Zhu Q 2006 *Phys. Rev. Lett.* **97** 266402
[382] Ochsenfeld C and Head-Gordon M 1997 *Chem. Phys. Lett.* **270** 399
[383] Kussmann J and Ochsenfeld C 2007 *J. Chem. Phys.* **127** 054103
[384] Beer M and Ochsenfeld C 2008 *J. Chem. Phys.* **128** 221102
[385] Liang W, Zhao Y and Head-Gordon M 2005 *J. Chem. Phys.* **123** 194106
[386] Izmaylov A, Brothers E and Scuseria G 2006 *J. Chem. Phys.* **125** 224105
[387] Coriani S, Høst S, Jansík B, Thøgersen L, Olsen J, Jørgensen P, Reine S, Pawłowski F, Helgaker T and Sałek P 2007 *J. Chem. Phys.* **126** 154108
[388] Xiang H J, Yang J, Hou J G and Zhu Q 2007 *J. Chem. Phys.* **126** 244707
[389] Rubensson E H and Zahedi S 2008 *J. Chem. Phys.* **128** 176101
[390] Mortensen J J and Parrinello M 2001 *J. Phys.: Condens. Matter* **13** 5731
[391] Artacho E and Miláns del Bosch L 1991 *Phys. Rev.* A **43** 5770
[392] White C A, Maslen P, Lee M S and Head-Gordon M 1997 *Chem. Phys. Lett.* **276** 133
[393] Schweizer S, Kussmann J, Doser B and Ochsenfeld C 2007 *J. Comput. Chem.* **29** 1004
[394] Grote M J and Huckle T 1997 *SIAM J. Sci. Comput.* **18** 838
[395] Benzi M 2002 *J. Comput. Phys.* **182** 418
[396] Maslen P E, Ochsenfeld C, White C A, Lee M S and Head-Gordon M 1998 *J. Phys. Chem.* A **102** 2215
[397] Gibson A, Haydock R and LaFemina J P 1993 *Phys. Rev.* B **47** 9229
[398] Burger S K and Yang W 2008 *J. Phys.: Condens. Matter* **20** 294209
[399] Benzi M, Meyer C and Tuma M 1996 *SIAM J. Sci. Comput.* **17** 1135
[400] Ozaki T 2001 *Phys. Rev.* B **64** 195110
[401] Baer R and Head-Gordon M 1998 *J. Chem. Phys.* **109** 10159
[402] Rubensson E H, Bock N, Holmström E and Niklasson A M N 2008 *J. Chem. Phys.* **128** 104105
[403] Jansik B, Host S, Jørgensen P, Olsen J and Helgaker T 2007 *J. Chem. Phys.* **126** 124104
[404] Qiu S Y, Wang C Z, Ho K M and Chan C T 1994 *J. Phys.: Condens. Matter* **6** 9153
[405] Benzi M and Razouk N 2007 *Electron Trans. Numer. Anal.* **28** 16
[406] Canning A, Galli G, Mauri F, Vita A D and Car R 1996 *Comput. Phys. Commun.* **94** 89
[407] Challacombe M and Bock N 2010 arXiv:1011.3534
[408] Hackbusch W 1999 *Computing* **62** 89
[409] Hine N D M, Haynes P D, Mostofi A A and Payne M C 2010 *J. Chem. Phys.* **133** 114111
[410] Bowler D R, Miyazaki T and Gillan M J 2001 *Comput. Phys. Commun.* **137** 255
[411] Bowler D R and Miyazaki T 2010 *J. Phys.: Condens. Matter* **22** 074207
[412] Hine N, Haynes P, Mostofi A, Skylaris C K and Payne M 2009 *Comput. Phys. Commun.* **180** 1041
[413] Saravanan C, Shao Y, Baer R, Ross P N and Head-Gordon M 2003 *J. Comput. Chem.* **24** 618
[414] Rubensson E H, Rudberg E and Sałek P 2007 *J. Comput. Chem.* **28** 2531
[415] Hoshi T 2007 *J. Phys.: Condens. Matter* **19** 365243
[416] Challacombe M 2000 *Comput. Phys. Commun.* **128** 93
[417] Gan C K and Challacombe M 2003 *J. Chem. Phys.* **118** 9128

[418] Brázdová V and Bowler D R 2008 *J. Phys.: Condens. Matter* **20** 275223
[419] Skylaris C K, Haynes P D, Mostofi A A and Payne M C 2006 *Phys. Status Solidi* b **243** 973
[420] Sanz-Navarro C, Grima R, García A, Bea E, Soba A, Cela J and Ordejón P 2011 *Theor. Chem. Acc.* **128** 825
[421] Vincent J J, Dixon S L and Merz K M Jr 1998 *Theor. Chem. Acc.* **99** 220
[422] Watson S C and Carter E A 2000 *Comput. Phys. Commun.* **128** 67
[423] Goedecker S, Lancon F and Deutsch T 2001 *Phys. Rev.* B **64** 161102
[424] Fernández-Serra M V, Artacho E and Soler J M 2003 *Phys. Rev.* B **67** 100101
[425] Nemeth K, Coulaud O, Monard G and Angyan J 2000 *J. Chem. Phys.* **113** 5598
[426] Németh K and Challacombe M 2004 *J. Chem. Phys.* **121** 2877
[427] Niklasson A M N, Challacombe M, Tymczak C J and Németh K 2010 *J. Chem. Phys.* **132** 124104
[428] Bowler D R and Gillan M J 2002 *Chem. Phys. Lett.* **355** 306
[429] CONQUEST: http://www.order-n.org/
[430] ERGOSCF: http://ergoscf.org/
[431] Bock N, Challacombe M, Gan C K, Henkelman G, Nemeth K, Niklasson A M N, Odell A, Schwegler E, Tymczak C J and Weber V 2011 FREEON http://freeon.org/ Los Alamos National Laboratory (LA-CC 01-2; LA-CC-04-086) copyright University of California
[432] ONETEP: http://www.onetep.org/
[433] OPENMX: http://www.openmx-square.org/
[434] Hung L, Huang C, Shin I, Ho G S, Lignères V L and Carter E A 2010 *Comput. Phys. Commun.* **181** 2208
[435] SIESTA: http://www.icmab.es/siesta/
[436] Ordejón P 2000 *Phys. Status Solidi* b **217** 335
[437] Román-Pérez G and Soler J M 2009 *Phys. Rev. Lett.* **103** 096102
[438] Gulans A, Puska M and Nieminen R 2009 *Phys. Rev.* B **79** 201105
[439] Skylaris C K, Haynes P D, Mostofi A A and Payne M C 2005 *J. Chem. Phys.* **122** 084119
[440] Skylaris C K, Haynes P D, Mostofi A A and Payne M C 2005 *J. Phys.: Condens. Matter* **17** 5757
[441] Haynes P D, Skylaris C K, Mostofi A A and Payne M C 2006 *Phys. Status Solidi* b **243** 2489
[442] Mostofi A A, Haynes P D, Skylaris C K and Payne M C 2007 *Mol. Simul.* **33** 551
[443] Skylaris C K, Haynes P D, Mostofi A A and Payne M C 2008 *J. Phys.: Condens. Matter* **20** 064209
[444] Skylaris C K, Mostofi A A, Haynes P D, Diéguez O and Payne M C 2002 *Phys. Rev.* B **66** 035119
[445] Skylaris C K, Mostofi A A, Haynes P D, Pickard C J and Payne M C 2001 *Comput. Phys. Commun.* **140** 315
[446] Hine N D M, Robinson M, Haynes P D, Skylaris C K, Payne M C and Mostofi A A 2011 *Phys. Rev.* B **83** 195102
[447] Hill Q and Skylaris C K 2009 *Proc. R. Soc.* A **465** 669
[448] Weber V, Tymczak C J and Challacombe M 2006 *J. Chem. Phys.* **124** 224107
[449] Weber V, Daul C and Challacombe M 2006 *J. Chem. Phys.* **124** 214105
[450] Bowler D R and Gillan M J 2000 *Mol. Simul.* **25** 239
[451] Bowler D R, Bush I J and Gillan M J 2000 *Int. J. Quantum Chem.* **77** 831
[452] Bowler D R, Choudhury R, Gillan M J and Miyazaki T 2006 *Phys. Status Solidi* b **243** 989
[453] Gillan M J, Bowler D R, Torralba A S and Miyazaki T 2007 *Comput. Phys. Commun.* **177** 14
[454] Miyazaki T, Bowler D R, Choudhury R and Gillan M J 2004 *J. Chem. Phys.* **121** 6186
[455] Sena A M P, Miyazaki T and Bowler D R 2011 *J. Chem. Theory Comput.* **7** 884
[456] Ikeshoji T, Tsuchida E, Morishita T, Ikeda K, Matsuo M, Kawazoe Y and Orimo S i 2011 *Phys. Rev.* B **83** 144301
[457] de Pablo P J, Moreno-Herrero F, Colchero J, Gómez Herrero J, Herrero P, Baró A M, Ordejón P, Soler J M and Artacho E 2000 *Phys. Rev. Lett.* **85** 4992
[458] Bondesson L, Rudberg E, Luo Y and Sałek P 2007 *J. Phys. Chem.* B **111** 10320
[459] Sawada T, Fedorov D G and Kitaura K 2010 *J. Phys. Chem.* B **114** 15700
[460] Ishikawa T, Mochizuki Y, Imamura K, Nakano T, Mori H, Tokiwa H, Tanaka K, Miyoshi E and Tanaka S 2006 *Chem. Phys. Lett.* **430** 361
[461] Yeole S D and Gadre S R 2010 *J. Chem. Phys.* **132** 094102
[462] Lee K, Murray E D, Kong L, Lundqvist B I and Langreth D C 2010 *Phys. Rev.* B **82** 081101
[463] Grimme S 2006 *J. Comput. Chem.* **27** 1787
[464] Grimme S, Antony J, Ehrlich S and Krieg H 2010 *J. Chem. Phys.* **132** 154104
[465] Tkatchenko A and Scheffler M 2009 *Phys. Rev. Lett.* **102** 073005
[466] Ercolessi F and Adams J B 1994 *Europhys. Lett.* **26** 583
[467] Otsuka T, Miyazaki T, Bowler D R and Gillan M J 2011, in preparation
[468] Miyazaki T, Bowler D R, Gillan M J, Otsuka T and Ohno T 2009 *AIP Conf. Proc.* **1148** 685
[469] Mo Y W, Savage D E, Swartzentruber B G and Lagally M G 1990 *Phys. Rev. Lett.* **65** 1020
[470] Miyazaki T, Bowler D R, Choudhury R and Gillan M J 2007 *Phys. Rev.* B **76** 115327
[471] Miyazaki T, Bowler D R, Gillan M J and Ohno T 2008 *J. Phys. Soc. Japan* **77** 123706
[472] Tao N J 2006 *Nature Nanotechnol.* **1** 173
[473] Joachim C and Ratner M A 2005 *Proc. Natl Acad. Sci. USA* **102** 8801
[474] Biel B, García-Vidal F J, Rubio Á and Flores F 2008 *J. Phys.: Condens. Matter* **20** 294214
[475] Humbel S, Sieber S and Morokuma K 1996 *J. Chem. Phys.* **105** 1959
[476] Canfield P, Dahlbom M G, Hush N S and Reimers J R 2006 *J. Chem. Phys.* **124** 024301
[477] Vashishta P, Kalia R K and Nakano A 2006 *J. Phys. Chem.* B **110** 3727
[478] Shimojo F, Nakano A, Kalia R K and Vashishta P 2009 *Appl. Phys. Lett.* **95** 043114
[479] Ho G, Ong M T, Caspersen K J and Carter E A 2007 *Phys. Chem. Chem. Phys.* **9** 4951
[480] Hung L and Carter E A 2011 *J. Phys. Chem.* C **115** 6269
[481] Hung L and Carter E A 2011 *Modelling Simul. Mater. Sci. Eng.* **19** 045002
[482] Radhakrishnan B and Gavini V 2010 *Phys. Rev.* B **82** 094117
[483] Hung L and Carter E A 2011 personal communication
[484] Fernández-Serra M V and Artacho E 2004 *J. Chem. Phys.* **121** 11136
[485] Zonias N, Lagoudakis P and Skylaris C K 2010 *J. Phys.: Condens. Matter* **22** 025303
[486] Hoshi T and Fujiwara T 2003 *J. Phys. Soc. Japan* **72** 2429
[487] Hoshi T, Iguchi Y and Fujiwara T 2005 *Phys. Rev.* B **72** 075323
[488] Bowler D R, Fattebert J L, Gillan M J, Haynes P D and Skylaris C K 2008 *J. Phys.: Condens. Matter* **20** 290301