8   SCIENTIFIC HIGHLIGHT OF THE MONTH: "ONETEP: linear-scaling density-functional theory with plane waves"

ONETEP: linear-scaling density-functional theory with plane waves

Peter D. Haynes[1], Chris-Kriton Skylaris[2], Arash A. Mostofi[3] and Mike C. Payne[1]

[1] *Theory of Condensed Matter, Cavendish Laboratory, University of Cambridge, U.K.*

[2] *Physical and Theoretical Chemistry Laboratory, University of Oxford, U.K.*

[3] *Department of Materials Science and Engineering, Massachusetts Institute of Technology, U.S.A.*

Abstract

An overview of the ONETEP (Order-N Electronic Total Energy Package) code is presented, focusing on the twin aims of overall linear scaling and controlled accuracy. The method is outlined, including a description of the density-matrix formulation of density-functional theory, and the optimisation procedures for both the density-kernel and the local orbitals or non-orthogonal generalised Wannier functions. Results of applying the method to a variety of systems are presented to demonstrate the accomplishment of the original aims.

1   Introduction

Density-functional theory (DFT) [1, 2] has made a unique impact [3] on science, well beyond the traditional realms of quantum-mechanical simulations into disciplines such as microelectronics [4], biochemistry [5] and geology [6]. This broad multidisciplinary appeal has its origin in the ability of DFT to provide a sufficiently accurate description of electron correlation for many purposes at a computational cost which scales favourably (as the cube of the system size N) compared with correlated wave function methods (which typically exhibit N^5 to N^7 scaling).

The origin of the N^3 asymptotic scaling of DFT can be understood with reference to Fig. 1, which shows one of the solutions $\psi_n(\mathbf{r})$ to the single-particle Kohn-Sham equation

$$\hat{H}_{\mathrm{s}}\psi_n(\mathbf{r}) = \left[-\tfrac{1}{2}\nabla^2 + V_{\mathrm{eff}}[n](\mathbf{r})\right]\psi_n(\mathbf{r}) = \varepsilon_n\psi_n(\mathbf{r}). \tag{1}$$

In addition to the potential due to the ion cores, the effective potential $V_{\mathrm{eff}}[n](\mathbf{r})$ describes electron-electron interactions in a mean field manner (via the electron density $n(\mathbf{r})$). As a result, Eq. 1 must be solved self-consistently along with

$$n(\mathbf{r}) = 2 \sum_n f_n |\psi_n(\mathbf{r})|^2 , \qquad (2)$$

where the occupation numbers $\{f_n\}$ are unity (zero) for states below (above) the chemical potential at absolute zero and the factor of two accounts for spin degeneracy. For small basis sets, Eq. 1 may be solved by direct matrix diagonalisation which scales as N^3. For large-scale simulations or large basis sets (not least the plane-wave pseudopotential method [7,8], the leading workhorse for DFT calculations) iterative methods [9] are more efficient. In this case, the orthonormality of the Kohn-Sham orbitals $\{\psi_n(\mathbf{r})\}$ must be imposed directly:

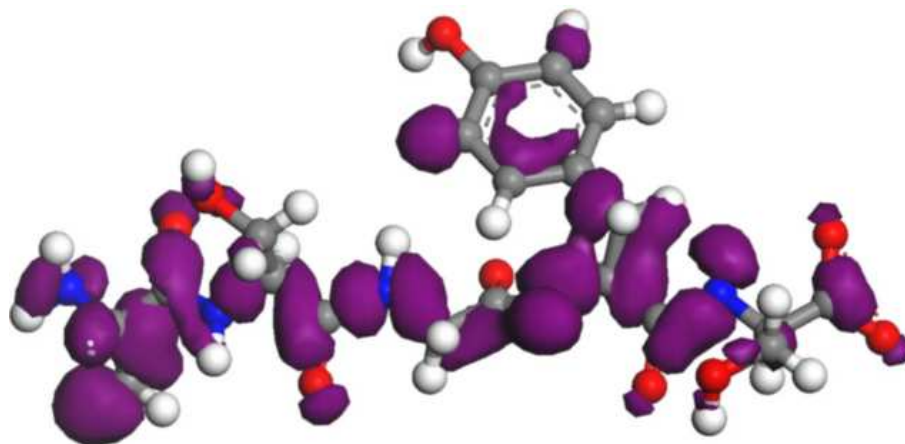$$\int \mathrm{d}^3 r \, \psi_m^*(\mathbf{r}) \psi_n(\mathbf{r}) = \delta_{mn}. \qquad (3)$$

Figure 1: A Kohn-Sham orbital for an oligopeptide molecule: in general, each orbital extends over the entire system.

The number of orbital pairs and thus the number of constraints (3) is proportional to N^2. As illustrated in Fig. 1, each Kohn-Sham orbital extends over the entire system so that the overlap integral in (3) requires a computational effort which scales linearly with N (in the same manner as the volume of the system increases). Thus the total cost of enforcing the orthonormality constraints scales as N^3.

While relatively benign compared to other methods, this N^3 scaling still presents a bottleneck which restricts the size of simulations to a few hundred atoms, even with the most powerful supercomputers. In recent years there has therefore been considerable interest in the development of linear-scaling or order-N methods (see the review articles [10,11] and references therein). These offer the potential to revolutionise the scope and scale of first-principles simulations based on DFT to include entire biological molecules and nanostructures containing many thousands of atoms.

While many different computational schemes have been proposed, this article focuses on one more recent method which has been implemented in the ONETEP code [12, 13]. The twofold aim of this work has been to develop a general purpose scheme which exhibits true linear scaling and controlled accuracy.

In all proposed order-N schemes the computational effort associated with the dominant part of the calculation (e.g. evaluation of the Hamiltonian matrix elements or application of the Hamiltonian to a trial vector) scales linearly with system-size. However this does not guarantee that the total calculation time will scale so well. For example, while the time for a single self-consistent iteration may scale linearly, various forms of ill-conditioning [14] can cause the number of iterations required to reach convergence to increase with system-size. The aim of true linear scaling thus requires that the total effort to calculate the desired physical properties of a system should scale linearly.

In order to obtain any linear-scaling method, it is necessary to make further approximations that exploit the "nearsightedness" [15, 16] of quantum many-particle systems, in addition to those common to all DFT methods (exchange-correlation functional, finite basis set, discrete Brillouin zone sampling and choice of pseudopotential if appropriate). Pursuing Kohn's analogy of myopia, this involves deliberately impairing the vision of the simulated quantum system so that the local density is blind to distant perturbations in the potential. As the range of sight is decreased, the calculation will become cheaper, but with a consequent loss of accuracy. Therefore this approximation must be controlled to reduce any error to an acceptable level if the results are to be used with confidence.

The accuracy of a calculation is heavily influenced by the choice of basis set. Linear-scaling methods can be divided into two categories: those which use relatively small basis sets of atomic-like orbitals such as numerical atomic orbitals [17–21], Slater type orbitals [22], localised spherical waves [23], or Gaussian type orbitals [24,25]; and those which use larger sets of simpler functions such as B-splines or blip functions [26], finite elements [27], or methods based on finite-difference [28] and multigrid [29] techniques. To satisfy the requirement for controlled accuracy, it is necessary to choose a basis set which may be improved systematically, ideally via a single parameter, and which spans the range from minimal atomic sets to plane-wave accuracy.

In Sec. 2 an overview of the ONETEP linear-scaling method is given, and in Sec. 3 an assessment is made of representative results for a variety of systems in the light of the above aims.

2 Overview of the ONETEP method

2.1 Density-matrix formulation

The set of Kohn-Sham orbitals $\{\psi_n(\mathbf{r})\}$ provides a complete description of the fictitious system of non-interacting particles in DFT. An equivalent description may also be given by the single-particle density-matrix $\rho(\mathbf{r}, \mathbf{r}')$,

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_n f_n \, \psi_n^*(\mathbf{r}) \, \psi_n(\mathbf{r}'), \tag{4}$$

which possesses the property of idempotency, namely

$$\rho^2(\mathbf{r}, \mathbf{r}') = \int \mathrm{d}^3 r'' \, \rho(\mathbf{r}, \mathbf{r}'') \, \rho(\mathbf{r}'', \mathbf{r}') = \rho(\mathbf{r}, \mathbf{r}'), \tag{5}$$

which requires the orthonormality of the orbitals $\{\psi_n(\mathbf{r})\}$ according to Eq. 3 and the *Auf-bau* principle of singly occupying all states up to the chemical potential, which itself follows straightforwardly from the Pauli exclusion principle. The density-matrix is thus the position representation of the projection operator onto the space of occupied states $\hat{\rho}$.

The density $n(\mathbf{r})$ may be obtained from the diagonal elements of the density-matrix,

$$n(\mathbf{r}) = 2\rho(\mathbf{r}, \mathbf{r}), \tag{6}$$

where the factor of two again accounts for spin degeneracy (assuming no spin polarisation). The total energy of the non-interacting system E_s may be obtained via

$$E_s = 2\,\mathrm{Tr}(\hat{\rho}\hat{H}_s) \tag{7}$$

and thus the energy of the real interacting system calculated by applying the usual double-counting corrections to the Hartree and exchange-correlation terms. The solution to Eq. 1 may therefore be found by minimising the energy with respect to the density-matrix, subject to the constraints of idempotency (5) and normalisation,

$$2\int \mathrm{d}^3r\, \rho(\mathbf{r}, \mathbf{r}) = N_e, \tag{8}$$

i.e. the density-matrix corresponds to a system of N_e electrons.

Since the number of occupied states is directly proportional to N and each state extends over the whole system (Fig. 1), the amount of information in the density-matrix defined by Eq. 4 therefore scales as N^2. Any calculation involving manipulation of this density-matrix will therefore scale quadratically with system-size at best. In order to obtain a linear-scaling method, it is necessary to exploit the nearsightedness of many-body quantum mechanics [15] mentioned above.

Both analytical [30, 31] and numerical [32] studies have demonstrated that, for an insulating system, both the Wannier functions and density-matrix decay exponentially, so that

$$\rho(\mathbf{r}, \mathbf{r}') \sim \exp(-\gamma\,|\mathbf{r} - \mathbf{r}'|) \to 0 \text{ as } |\mathbf{r} - \mathbf{r}'| \to \infty. \tag{9}$$

This means that for any given position $\mathbf{r}$ the density-matrix $\rho(\mathbf{r}, \mathbf{r}')$ differs significantly from zero only for points $\mathbf{r}'$ within a finite volume around $\mathbf{r}$. Since the decay rate γ depends only on the energy gap between the highest occupied and lowest unoccupied states, this volume is independent of system-size. Thus the total amount of significant information in the density-matrix only scales linearly with N.

In practice this is exploited by writing the density-matrix in separable form:

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_{\alpha\beta} \phi_\alpha(\mathbf{r}) K^{\alpha\beta} \phi_\beta^*(\mathbf{r}'), \tag{10}$$

where the $\{\phi_\alpha(\mathbf{r})\}$, illustrated in Fig. 2, are a set of spatially localised non-orthogonal functions which span a superspace of the Hilbert space of the set of occupied Kohn-Sham orbitals
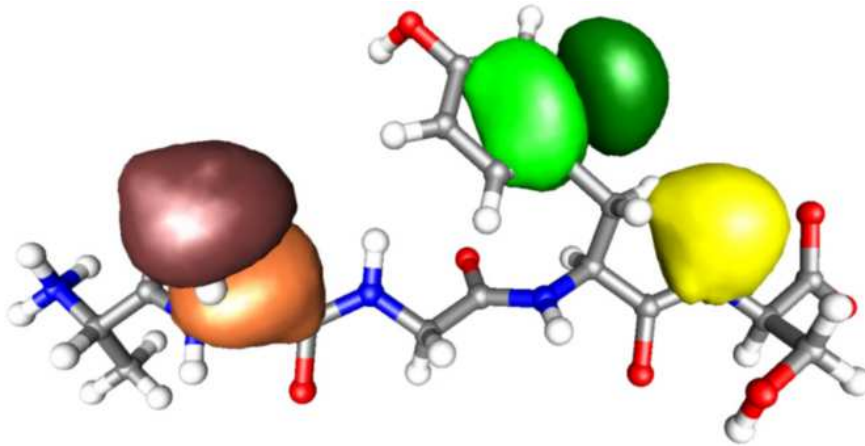
Figure 2: Three localised orbitals (non-orthogonal generalised Wannier functions) generated by ONETEP for the same oligopeptide molecule as Fig. 1.

$\{\psi_n(\mathbf{r}); \varepsilon_n < \mu\}$ and have been called support functions [33] or non-orthogonal generalised Wannier functions (NGWFs) [34] in this work. The matrix $K^{\alpha\beta}$, known as the density-kernel [35] is in fact the representation of the density-matrix in the set of duals of the NGWFs $\{\phi^\alpha(\mathbf{r})\}$ defined by

$$\langle \phi_\beta | \phi^\alpha \rangle = \int \mathrm{d}^3 r \, \phi_\beta^*(\mathbf{r}) \phi^\alpha(\mathbf{r}) = \delta_\beta{}^\alpha. \tag{11}$$

The total amount of information contained in the density-matrix defined by Eq. 10 must of course still scale as N^2. However the advantage of this form is that it allows the nearsightedness to be exploited with the use of spatial cut-offs.

First, the NGWFs which are exponentially localised are truncated, by allowing them to be non-vanishing only in spherical regions of fixed radii $\{r_\alpha\}$ and centred at positions $\{\mathbf{R}_\alpha\}$. In ONETEP, a number of NGWFs are associated with each atom in the system, so that the regions are centred on atoms and their radii essentially depend only on the atomic species.

Second, the density-kernel is required to be a sparse matrix, by discarding elements $K^{\alpha\beta}$ corresponding to NGWFs centred further apart than some cut-off r_K. Note that since the density-kernel is related directly to the duals of the NGWFs, rather than the NGWFs themselves, the density-kernel cut-off r_K is not simply the sum of the radii $r_\alpha + r_\beta$. This is a consequence of the non-orthogonality of the NGWFs.

The approach adopted in ONETEP of truncating the NGWFs and predetermining the sparsity pattern of the density-kernel contrasts with approaches based on thresholding [36] which are common in methods based on Gaussian basis sets. The former approach has the advantage that the overlap matrix for NGWFs remains strictly positive definite, whereas the latter approach is more flexible and hence potentially more efficient.

Imposing spatial cut-offs on the NGWFs and the density-kernel results in a density-matrix whose information content scales linearly with system-size, an approximation which is controlled by adjusting the $\{r_\alpha\}$ and r_K. In practice, these cut-offs are increased until the desired physical

properties of the system converge.

2.2 Density-kernel optimisation

Most linear-scaling methods fall into two categories, typically along the same lines as the basis set division above. The first is equivalent to optimising the density-kernel only [22, 37–39] for a fixed, but potentially large, set of local orbitals and the second involves both density-kernel and NGWF optimisation [26, 34, 40]. These two approaches are further compared in the next section, while this section concerns the common point of density-kernel optimisation.

Functionally at least, this stage of the calculation proceeds along the lines of an *ab initio* tight-binding [41] calculation and linear-scaling methods for tight-binding [42] may be applicable. While self-consistency may be treated independently of density-kernel and NGWF optimisation [43], the approach in ONETEP is closely related to the ensemble DFT method for metallic systems [44] and performs both density-kernel and NGWF optimisation self-consistently.

To impose the idempotency constraint a combination of two methods is used: a penalty-functional method [45] and the method of Li, Nunes and Vanderbilt [47], and independently Daw [48], based on McWeeny's purification transformation [35]. The latter has been implemented in both its original non-orthogonal formulation [49] and the Millam and Scuseria variant [50].
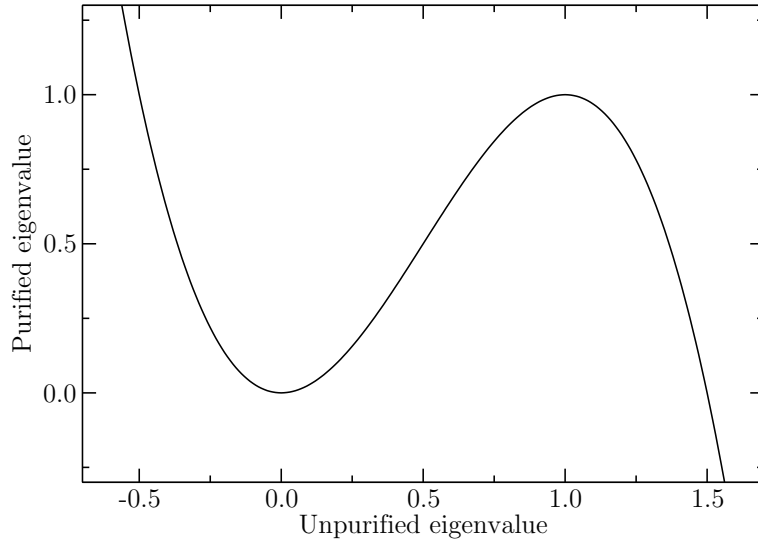

Figure 3: The purification transformation illustrated in terms of its effect on the density-matrix eigenvalues (occupation numbers).

The purification transformation may be defined in terms of an auxiliary matrix $\sigma(\mathbf{r}, \mathbf{r}')$ by:

$$\rho(\mathbf{r}, \mathbf{r}') = 3\sigma^2(\mathbf{r}, \mathbf{r}') - 2\sigma^3(\mathbf{r}, \mathbf{r}'). \tag{12}$$

Clearly $\rho(\mathbf{r}, \mathbf{r}')$ and $\sigma(\mathbf{r}, \mathbf{r}')$ commute and so the transformation is best illustrated as in Fig. 3 in terms of their eigenvalues, which for $\rho(\mathbf{r}, \mathbf{r}')$ are the occupation numbers $\{f_n\}$ of Eq. 4. Under repeated iteration of this transformation, eigenvalues around zero vanish quadratically

and those close to unity converge to the maximum there. Expressing the density-matrix in this way and optimising the auxiliary matrix thus applies the constraint of idempotency, so long as the occupation numbers remain in the interval $\left[\frac{1-\sqrt{5}}{2}, \frac{1+\sqrt{5}}{2}\right]$ where the transformation is stable. Should the calculation become unstable (generally at the start of a calculation) then the globally convergent penalty-functional method is used instead.

The purification transformation may also be derived as the result of a steepest descents minimisation of the positive definite functional $P[\rho]$ given by

$$P[\rho] = \int \int \mathrm{d}^3 r \, \mathrm{d}^3 r' \left[\rho^2(\mathbf{r}, \mathbf{r}') - \rho(\mathbf{r}, \mathbf{r}')\right] \left[\rho^2(\mathbf{r}', \mathbf{r}) - \rho(\mathbf{r}', \mathbf{r})\right], \tag{13}$$

which gives a quantitative measure of the departure of $\rho(\mathbf{r}, \mathbf{r}')$ from idempotency (for nearly-idempotent density-matrices it is proportional to the mean square deviation of the occupation numbers from zero or unity). $P[\rho]$ may thus be used as a penalty functional, some proportion of which is added to the energy functional to enforce idempotency [45].
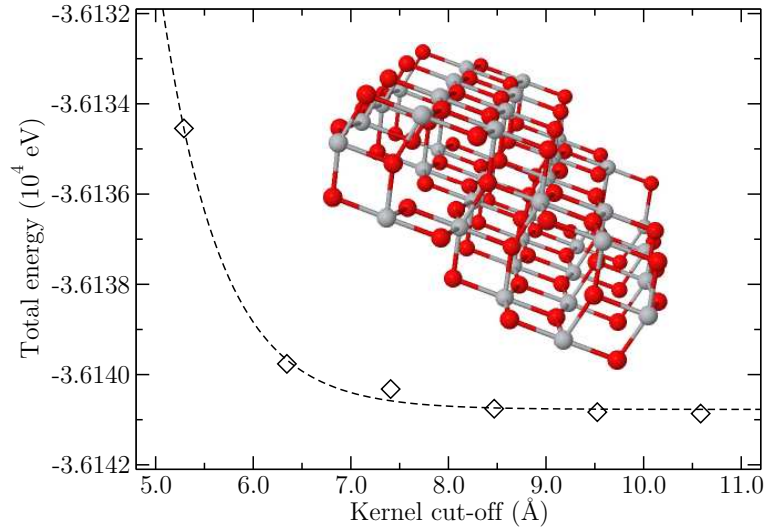

Figure 4: Total energy convergence with respect to density-kernel cut-off r_K for a $Ti_{38}O_{76}$ cluster [46].

Figure 4 shows a particular example of how the total energy varies as the density-kernel cut-off r_K is increased, and this behaviour is qualitatively similar to that observed in all insulating systems. Physical properties based on energy differences converge even more rapidly with r_K. For example, in crystalline silicon the errors from truncating the density-kernel to $r_K = 10$ Å are 0.3% and 0.1% in the lattice parameter and bulk modulus respectively [13].

2.3 NGWF optimisation

Linear-scaling methods which use a fixed set of local orbitals generally use atomic-type functions, typically solutions to the Schrödinger equation for isolated atoms, possibly immersed in a confining potential. The smallest minimal orbital sets allow calculations which are quick but relatively inaccurate. To improve their accuracy, the size of the orbital set is increased, gen-

erating split valence or multiple-zeta sets, with optional polarisation functions to describe the response of the atom to an applied electric field.

The second approach, which is the one adopted in ONETEP, instead optimises these local orbitals in the environment of the system being studied, to generate NGWFs. Thus it is no longer necessary to increase the size of the orbital set to improve transferability. While extra computational effort is required to perform this optimisation, some saving is obtained since the orbital set used is of minimal size, while high accuracy is maintained.
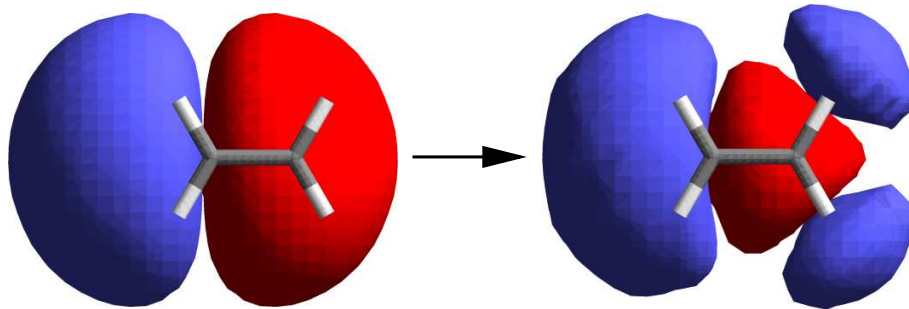

Figure 5: A single NGWF centred on one carbon atom of an ethene molecule before (left) and after (right) optimisation in ONETEP.

Figure 5 illustrates this optimisation for the case of a single orbital centred on the left-hand carbon atom in an ethene molecule. The orbital is initialised to a truncated p_x atomic orbital. After the orbital has been adapted *in situ* to generate the NGWF on the right-hand side of the figure, it is clear that the NGWF now reflects the molecular environment.

In order to perform this optimisation, it is necessary to expand the NGWFs in some underlying basis set of primitive functions. In ONETEP, these are chosen to be periodic cardinal sine (psinc) functions [34, 51], also known as Dirichlet or Fourier Lagrange functions [52, 53]. One such function is illustrated in Fig. 6, and there is one function centred on each grid point of a regular mesh commensurate with the simulation cell.

Since the psinc functions are related to plane-waves by Fourier transform, they enable the kinetic energy to be calculated accurately and efficiently [54, 55] using fast Fourier transforms (FFTs). They are also orthogonal by construction and highly localised, being non-zero at only one grid point (although they do oscillate between grid points). The basis set quality may also be controlled by a single energy cut-off parameter, related to that used for plane-wave basis sets [13]. Thus in ONETEP it is possible to vary the accuracy from that of a minimal atomic-type basis (by performing no NGWF optimisation) to full plane-wave accuracy.

The scheme as described so far involves expanding each NGWF in terms of the psinc basis functions whose centres lie within the truncation volume of that NGWF. This is a small subset of the total psinc basis set, as illustrated schematically in Fig. 7. However the FFTs required to calculate the kinetic energy remain global operations, so that the time to calculate the kinetic energy of a single NGWF scales as $O(N \log N)$, giving an overall scaling higher than N^2.

To overcome this problem the FFT box technique has been introduced [54], which is illustrated in Fig. 8. A smaller FFT box is associated with each NGWF, which is centred on that NGWF.
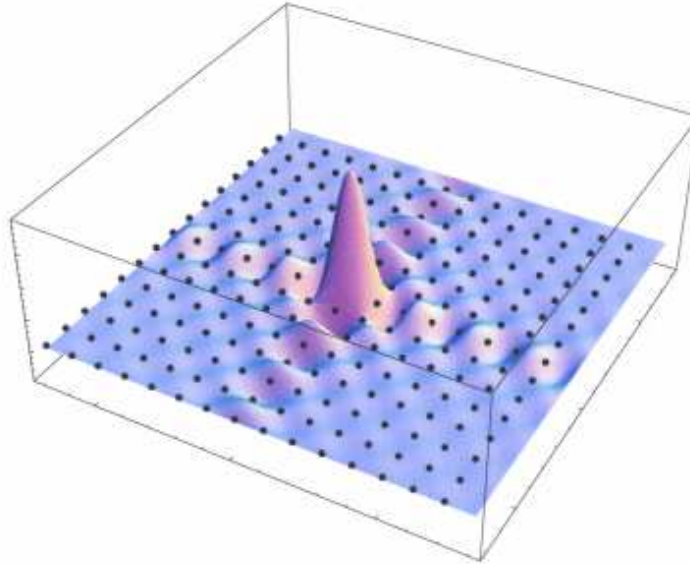
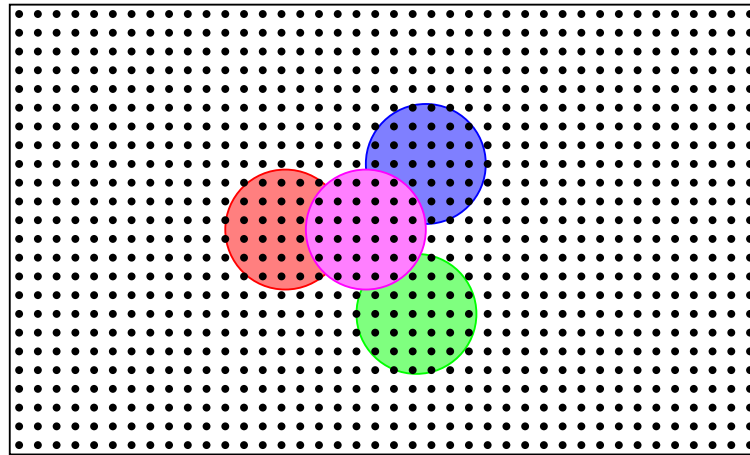Figure 6: A psinc basis function used to expand the NGWFs in ONETEP.

Figure 7: The regular psinc grid covering the whole simulation cell.

The FFT boxes for all NGWFs are of a universal shape and size, and are chosen to be sufficiently large to include the central NGWF and all of its overlapping neighbours. This guarantees the consistent action of operators and the hermiticity of the calculated Hamiltonian.

A psinc basis set may be defined for each FFT box, and each NGWF is now expanded in those functions centred in its truncation volume. In effect, the periodicity of the NGWFs is altered, but since the NGWFs are localised within volumes smaller than the FFT box, this is a very good approximation. Equivalently, since the NGWFs are localised in real-space, their Fourier transforms are broad in reciprocal-space and thus the coarser sampling in reciprocal-space provided by the FFT box compared to the cell remains adequate. The volume of the FFT box is independent of system-size, depending only upon the cut-off radii of the NGWFs, and so the computational effort of an FFT within the FFT box does not increase as the system becomes larger, and linear scaling is recovered.
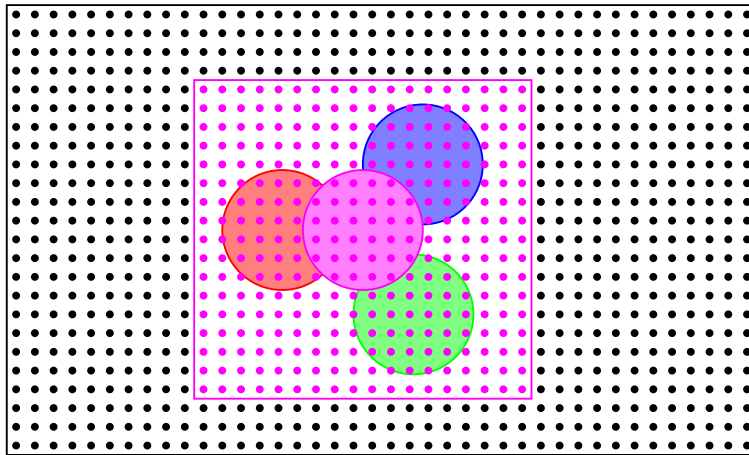
Figure 8: Illustration of the FFT box technique used in ONETEP.

In ONETEP the FFT box is used not only to calculate the kinetic energy, but also to interpolate the NGWFs before the calculation of the density (to avoid aliasing) and to evaluate every term in the Hamiltonian [56].

The FFT box also allows an estimate of the cross-over (the point at which a linear-scaling method is more efficient than traditional cubic-scaling methods). Comparing ONETEP to a traditional plane-wave pseudopotential code shows that both methods spend considerable time performing FFTs. For the traditional code these are performed over the entire simulation cell. In the case of ONETEP, fine grid FFTs in the FFT box dominate the computational effort. Since both methods perform similar numbers of FFTs, the cross-over is expected to occur when the FFT grids are about the same size i.e. the volume of the FFT box and the simulation cell are comparable. The number of atoms that this corresponds to depends upon the nature of the system. ONETEP has a particular advantage for low-dimensional structures such as molecules, polymers, nanotubes and surfaces, since regions of vacuum can be included essentially for free [12]. For solids, the cross-over will be rather higher in the number of atoms.

3   Results and assessment

The ONETEP method has been parallelised, and results are available elsewhere [12, 57]. By construction, the cost of a single self-consistent iteration in the determination of the ground state scales linearly with system-size. However, as pointed out in Sec. 1, this is not sufficient to guarantee that the calculation of physical properties scales linearly with system-size, since the number of iterations required has been observed to increase in some methods.

Figure 9 shows the convergence of the ONETEP method with respect to iterations of the NGWF optimisation procedure. A wide variety of systems with different properties have been chosen: a supercell of crystalline silicon, a hydrogen-bonded formaldehyde-water complex, a (20,0) carbon nanotube, a protein complex and a supercell of the zeolite ZSM5. Both the gain in total energy and the length of the residual vector are plotted. In all cases, the convergence criteria are satisfied within 10-20 iterations, independently of the number of atoms in the system. Thus
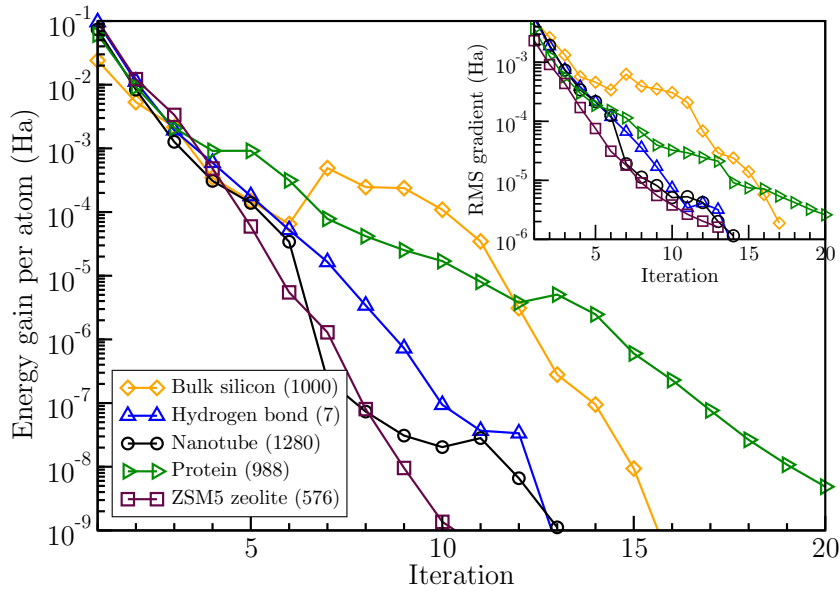
Figure 9: Energy (inset RMS gradient) convergence of the NGWF optimisation for a variety of systems – figures in parentheses in legend indicate number of atoms.

in ONETEP, the total calculation time, not just the time per iteration, scales linearly with system-size. This is believed to result from the choice of an orthogonal basis set and the careful and consistent treatment of all terms in the energy.
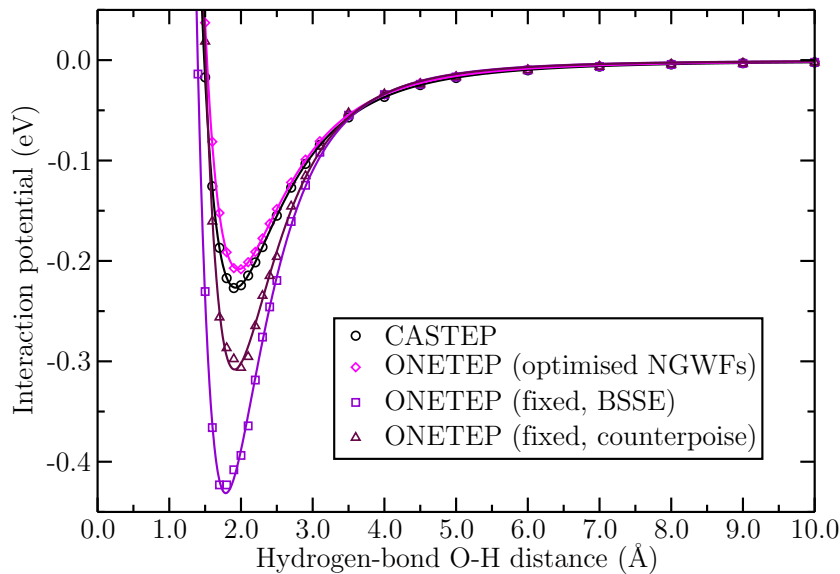

Figure 10: Interaction potentials (negated binding energies) for a hydrogen-bonded water dimer.

The accuracy of ONETEP calculations has been compared with traditional plane-wave pseu-dopotential and all-electron methods elsewhere [13]. Figure 10 shows the interaction potential for a hydrogen-bonded water dimer as calculated by the CASTEP [58] code and by ONETEP using the same pseudopotential, gradient-corrected exchange-correlation functional [59] and equivalent energy cut-offs for the basis sets [60]. In two of the ONETEP calculations, only the density-

kernel has been optimised, the NGWFs have been fixed as a minimal set of atomic-type orbitals. This calculation exhibits basis set superposition error (BSSE) [61] for which in this system the counterpoise correction [62] is only partially successful. While ONETEP is capable of performing calculations at this level, for weakly bound systems such as this one, high accuracy is required. When the NGWFs are optimised in addition to the density-kernel, no BSSE is observed and the agreement is vastly improved. With fully converged basis sets, the equilibrium bond lengths and binding energies obtained by ONETEP and CASTEP agree to 0.3% and 2% respectively.

Together with the results presented in Fig. 4, it can be seen that the accuracy of calculations in ONETEP can be controlled, and that plane-wave accuracy can be achieved routinely. The twin aims of overall linear scaling and controlled accuracy have thus been achieved.

Acknowledgments

References

[1] P. Hohenberg and W. Kohn, *Phys. Rev.* **136**, 864 (1964).

[2] W. Kohn W and L. J. Sham, *Phys. Rev.* **140**, 1133 (1965).

[3] S. Redner, *Phys. Today* **58**, 49 (June 2005).

[4] J. Junquera and P. Ghosez, *Nature* **422**, 506 (2003).

[5] M. T. Green, J. H. Dawson, and H. B. Gray, *Science* **304**, 1653 (2004).

[6] D. Alfè, M. J. Gillan, and G. D. Price, *Nature* **401**, 462 (1999).

[7] J. Ihm, A. Zunger, and M. L. Cohen, *J. Phys. C* **12**, 4409 (1979).

[8] P. J. H. Denteneer and W. van Haeringen, *J. Phys. C* **18**, 4127 (1985).

[9] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, *Rev. Mod. Phys.* **64**, 1045 (1992).

[10] G. Galli, *Curr. Opin. Solid St. M.* **1**, 864 (1996).

[11] S. Goedecker, *Rev. Mod. Phys.* **71**, 1085 (1999).

[12] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne, *J. Chem. Phys.* **122**, 084119 (2005).

[13] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne, *J. Phys.: Condens. Matter* **17**, 5757 (2005).

[14] D. R. Bowler and M. J. Gillan, *Comput. Phys. Commun.* **112**, 103 (1998).

[15] W. Kohn, *Phys. Rev. Lett.* **76**, 3168 (1996).

[16] E. Prodan and W. Kohn, *Proc. Nat. Acad. Sci. U.S.A* **102**, 11635 (2005).

[17] O. F. Sankey and D. J. Niklewski, *Phys. Rev. B* **40**, 3979 (1989).

[18] B. Delley, *J. Chem. Phys.* **113**, 7756 (2000).

[19] S. D. Kenny, A. P. Horsfield, and H. Fujitani, *Phys. Rev. B* **62**, 4899 (2000).

[20] J. Junquera, Ó. Paz, D. Sánchez-Portal, and E. Artacho, *Phys. Rev. B* **64**, 235111 (2001).

[21] E. Anglada, J. M. Soler, J. Junquera, and E. Artacho, *Phys. Rev. B* **66**, 205101 (2002).

[22] C. F. Guerra, J. G. Snijders, G. te Velde, and E. J. Baerends, *Theor. Chem. Acc.* **99**, 391 (1998).

[23] P. D. Haynes and M. C. Payne, *Comput. Phys. Commun.* **102**, 17 (1997).

[24] C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, *Chem. Phys. Lett.* **253**, 268 (1996).

[25] M. C. Strain, G. E. Scuseria, and M. J. Frisch, *Science* **271**, 51 (1996).

[26] E. Hernández, M. J. Gillan, and C. M. Goringe, *Phys. Rev. B* **55**, 13485 (1997).

[27] J. E. Pask and P. A. Sterne, *Modelling Simul. Mater. Sci. Eng.* **13**, R71 (2005).

[28] E. Hernández, M. J. Gillan, and C. M. Goringe, *Phys. Rev. B* **53**, 7157 (1996).

[29] M. B. Nardelli, J. L. Fattebert, and J. Bernholc, *Phys. Rev. B* **64**, 245423 (2001).

[30] W. Kohn, *Phys. Rev.* **115**, 809 (1959).

[31] J. des Cloizeaux, *Phys. Rev.* **135**, A685 and A698 (1964).

[32] S. Ismail-Beigi and T. Arias, *Phys. Rev. Lett.* **82**, 2127 (1999).

[33] E. Hernández and M. J. Gillan, *Phys. Rev. B* **51**, 10157 (1995).

[34] C.-K. Skylaris, A. A. Mostofi, P. D. Haynes, O. Diéguez, and M. C. Payne, *Phys. Rev. B* **66**, 035119 (2002).

[35] R. McWeeny, *Rev. Mod. Phys.* **32**, 335 (1960).

[36] E. Schwegler and M. Challacombe, *J. Chem. Phys.* **105**, 2726 (1996).

[37] G. E. Scuseria, *J. Phys. Chem. A* **103**, 4782 (1999).

[38] M. Challacombe, *J. Chem. Phys.* **110**, 2332 (1999).

[39] J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal, *J. Phys.: Condens. Matter* **14**, 2745 (2002).

[40] J. L. Fattebert and J. Bernholc, *Phys. Rev. B* **62**, 1713 (2000).

[41] A. P. Horsfield and A. M. Bratkovsky, *J. Phys.: Condens. Matter* **12**, R1 (2000).

[42] D. R. Bowler, M. Aoki, C. M. Goringe, A. P. Horsfield, and D. G. Pettifor, *Modelling Simul. Mater. Sci. Eng.* **5**, 199 (1997).

[43] D. R. Bowler, T. Miyazaki, and M. J. Gillan, *J. Phys.: Condens. Matter* **14**, 2781 (2002).

[44] N. Marzari, D. Vanderbilt, and M. C. Payne, *Phys. Rev. Lett.* **79**, 1337 (1997).

[45] P. D. Haynes and M. C. Payne, *Phys. Rev. B* **59**, 12173 (1999).

[46] P. D. Haynes, A. A. Mostofi, C.-K. Skylaris, and M. C. Payne, to appear in *J. Phys.: Conf. Ser.* (2005).

[47] X.-P. Li, R. W. Nunes, and D. Vanderbilt, *Phys. Rev. B* **47**, 10891 (1993).

[48] M. S. Daw, *Phys. Rev. B* **47**, 10895 (1993).

[49] R. W. Nunes and D. Vanderbilt, *Phys. Rev. B* **50**, 17611 (1994).

[50] J. M. Millam and G. E. Scuseria, *J. Chem. Phys.* **106**, 5569 (1997).

[51] A. A. Mostofi, P. D. Haynes, C.-K. Skylaris, and M. C. Payne, *J. Chem. Phys.* **119**, 8842 (2003).

[52] D. Baye and P.-H. Heenen, *J. Phys. A: Math. Gen.* **19**, 2041 (1986).

[53] K. Varga, Z. Zhang, and S. T. Pantelides, *Phys. Rev. Lett.* **93**, 176403 (2004).

[54] C.-K. Skylaris, A. A. Mostofi, P. D. Haynes, C. J. Pickard, and M. C. Payne, *Comput. Phys. Commun.* **140**, 315 (2001).

[55] C.-K. Skylaris, O. Diéguez, P. D. Haynes, and M. C. Payne, *Phys. Rev. B* **66**, 073103 (2002).

[56] A. A. Mostofi, C.-K. Skylaris, P. D. Haynes, and M. C. Payne, *Comput. Phys. Commun.* **147**, 788 (2002).

[57] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne, submitted to *Phys. Stat. Sol. B*.

[58] M. D. Segall, P. J. D. Lindan, M. J. Probert, C. J. Pickard, P. J. Hasnip, S. J. Clark, and M. C. Payne, *J. Phys.: Condens. Matter* **14**, 2717 (2002).

[59] J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).

[60] P. D. Haynes, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne, submitted.

[61] S. F. Boys and F. Bernardi, *Mol. Phys.* **19**, 553 (1970).

[62] H. B. Jansen and P. Ros, *Chem. Phys. Lett.* **3**, 140 (1969).