

Stochastic linear scaling for metals and nonmetals

Florian R. Krajewski and Michele Parrinello

Computational Science, Department of Chemistry and Applied Biosciences, ETH Zurich, USI Campus, Via Giuseppe Buffi 13,
CH-6900 Lugano, Switzerland

(Received 10 March 2005; published 17 June 2005)

Total energy electronic structure calculations, based on density functional theory or on the more empirical tight binding approach, are generally believed to scale as the cube of the number of electrons. By using the localization property of the high temperature density matrix we present exact deterministic algorithms that scale linearly in one dimension and quadratically in all others. We also introduce a stochastic algorithm that scales linearly with system size. These results hold for metallic and nonmetallic systems and are substantiated by numerical calculations on model systems.

DOI: 10.1103/PhysRevB.71.233105

PACS number(s): 71.15.-m, 31.15.-p

Total energy electronic structure calculations and molecular dynamics simulations based on density functional theory (DFT) or on the tight binding approach have been very successful in describing a large variety of phenomena and are finding increasing application in many fields of science. However even with present-day computer technology the size of the systems that can be studied is very restricted. This is due to the cubic dependence of the commonly used algorithms on the number of particles. This makes a daunting task of calculating the electronic properties of systems as large as those that are for instance of interest to nanotechnology and biochemistry.

Some 15 years ago it was realized that this need not be so and that linearly scaling algorithms could be devised.^{1,2} Up to now large number of linearly scaling algorithms have been proposed,³ but they are not devoid of problems and for a variety of reasons they are not yet routinely used. Most algorithms rely on the fact that the wave functions can be localized and have an exponential decay leading to a sparse Hamiltonian. This property does not hold when the gap between occupied and unoccupied levels vanishes, as in the case of metals for which it has proven difficult to obtain linearly scaling algorithms.

Here we propose a new approach to this problem that does not rely on an ability to localize the wave functions and is therefore equally applicable to metallic and nonmetallic systems. We introduce a series of algorithms which defy the commonly accepted wisdom that DFT calculations are of $O(N^3)$. In fact they are of $O(N^2)$ and even of $O(N)$. Furthermore we propose a stochastic algorithm that is linear scaling in all dimensions. A feature which sets our method apart from others is that it scales with the volume of the system and not with the number of electrons. A bonus, if one treats atomic species that are rich in electrons.

We work at finite temperature $1/\beta$ in the grand canonical ensemble, where the number of particles is controlled by the chemical potential μ and the relevant thermodynamic potential for spinless fermions is⁴

$$\Omega_f = -\frac{1}{\beta} \ln \det(1 + e^{\beta(\mu-H)}), \quad (1)$$

where H is a single particle Hamiltonian $H = -\frac{1}{2}\nabla^2 + V(r)$. In the $\beta \rightarrow \infty$ limit

$$\Omega_f = \sum_i \varepsilon_i - \mu N, \quad (2)$$

where the sum runs over the N lowest occupied states and N is the total number of particles. Limiting oneself to the consideration of free electrons in an external field is not as restrictive as it might at first seem. In fact it is well known that the full density functional can be obtained from the sum over the occupied orbital energy in Eq. (2), if the external potential of the Hamiltonian H is the self-consistent DFT potential and double counting terms are subtracted.⁵ Since the latter can be calculated with linear effort, finding a method for evaluating Ω_f in $O(N)$ operations solves the algorithmically hard part of the problem.

We make use of the identity:

$$(1 + e^{\beta(\mu-H)}) = \prod_{l=1}^P (1 + e^{i(\pi/P)(2l-1)} e^{(\beta/P)(\mu-H)}) \quad (3)$$

which is valid for any even P and the product goes over all the complex P^{th} roots of -1 . Using this decomposition one finds

$$\Omega_f = -\frac{1}{\beta} \sum_{l=1}^P \ln \det(\mathbf{M}(l)) \quad (4)$$

with $\mathbf{M}(l) = 1 + e^{i(\pi/P)(2l-1)} e^{(\beta/P)(\mu-H)}$. This is rather more complicated than Eq. (1) but has the advantage that it involves the propagator $e^{(\beta/P)(\mu-H)}$ rather than the more difficult $e^{\beta(\mu-H)}$. In fact if P is large enough one can use for $e^{(\beta/P)(\mu-H)}$ one of the many high-temperature representations of the exponential operator, such as the one based on the expansion in terms of Chebychev polynomials^{6,7} or the Trotter decomposition,⁸ as is commonly done in the numerical evaluation of path integrals. In this first implementation we shall employ the latter, which uses the identity

$$\langle r | e^{(\beta/P)(\mu-H)} | r' \rangle = e^{\beta/2P(\mu-V(r))} e^{-(mP/2\hbar^2\beta)(r-r')^2} e^{(\beta/2P)(\mu-V(r'))} \quad (5)$$

valid to $O((\beta/P)^3)$. For large P $e^{-(mP/2\hbar^2\beta)(r-r')^2}$ decays very

rapidly and most elements of $e^{(\beta/P)(\mu-H)}$ can be neglected, leading to a sparse matrix which is a key factor for obtaining a linear scaling algorithm. Since this property does not depend on the possibility of localizing the wave functions this

makes our method suitable also for metals. Alavi and Frenkel⁹ have shown that the high temperature density matrix in a cubic lattice of spacing δ can equally well be written as:

$$\langle i | e^{(\beta/P)(\mu-H)} | j \rangle = C \begin{cases} e^{(\beta/P)(\mu-V(i))} & i = j \\ e^{(\beta/2P)(\mu-V(i))} e^{-(mP\delta^2/2\hbar^2\beta)} e^{(\beta/2P)(\mu-V(j))} & \text{if } i \text{ and } j \text{ are first neighbors} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where i and j are lattice site indices, P is related to the lattice spacing by $P=3.432\hbar^2\beta/(m\delta^2)$ and C is a normalization constant. This expression is compatible with the approximations made so far, leads to an elegant lattice model with nearest neighbor interactions and at the same time does not make the problem any less complex, nor does it alter its scaling behavior.

Using this representation of the density matrix it is easily seen that in one dimension the matrices $\mathbf{M}(l)$ that appear in Eq. (4) are tridiagonal. Since the determinant of a tridiagonal matrix can be computed in $O(M)$ operations, we arrive at the interesting result that linear scaling is exact in 1D. Moving to higher dimensions the $\mathbf{M}(l)$ matrices become block tridiagonal where the dimension of each block is $m=M^{(d-1)/d}$. In spite of the fact that the blocks are very sparse we were unable to calculate the determinant of $\mathbf{M}(l)$ in less than $Mm^2=M^{3-2/d}$ operations. This is only marginally better than the standard M^3 scaling. Furthermore in $3d$ the resulting algorithm has a very unfavorable prefactor, which makes this approach unpractical unless substantially improved for instance by better exploiting the sparsity of $\mathbf{M}(l)$.

A more favorable scaling can be obtained if one focuses on the response of Ω_f with respect to appropriate parameters, which is a standard way of calculating physical quantities. For instance the number of particles is given by $\langle N \rangle = -(\partial\Omega_f/\partial\mu)$, the energy by $\langle E \rangle = (\partial(\beta\Omega_f)/\partial\beta) + \mu\langle N \rangle$ and so on. In general one can write for the value of a property A , conjugated to the field λ_A ,

$$\langle A \rangle = -\frac{1}{\beta} \sum_{l=1}^P \text{Tr} \left(\mathbf{M}(l)^{-1} \frac{\partial \mathbf{M}(l)}{\partial \lambda_A} \right), \quad (7)$$

which requires the inversion of the sparse matrices $\mathbf{M}(l)$ and not the calculation of its determinant. The inverse of $\mathbf{M}(l)$ can be found if one solves the M sets of linear equations $\mathbf{M}(l)\phi_j^l = \psi_j$, where $\{\psi_j\}$ is a complete orthonormal basis set and is given by $\mathbf{M}(l)^{-1} = \sum_{j=1}^M \phi_j^l \psi_j^\dagger$. Using a preconditioned biconjugate gradient method¹⁰ and the sparsity of $\mathbf{M}(l)$ we find that solving each linear equation takes $O(M)$ operations leading to an overall quadratic scaling also in 3D. An efficient preconditioner has proved to be the inverse of $\mathbf{M}_f(l)$ for free particles. Although in this case the full inverse can be evaluated exactly using Fourier transforms methods it is

computationally expedient to truncate $\mathbf{M}_f(l)^{-1}$ so as to give to $\mathbf{M}_f(l)^{-1}$ the same sparse structure as $\mathbf{M}(l)$. As we shall see below the theoretical $O(M^2)$ scaling can be demonstrated in practice, unfortunately at the cost of a large prefactor.

However linear scaling can be obtained if we use a stochastic approach. Taking our cue from what is done in quantum chromodynamics we introduce a random vector ψ . If the ψ_i are stochastically distributed such that their average satisfies

$$\langle \psi_i \psi_j^* \rangle = \delta_{ij} \quad \text{and} \quad \langle \psi_i \rangle = 0, \quad (8)$$

the inverse of \mathbf{M} can be written as an expectation value

$$\mathbf{M}(l)^{-1} = \langle \phi^l \psi^\dagger \rangle, \quad (9)$$

where the average is taken over the stochastic process and ϕ^l is the solution of the linear equation $\mathbf{M}(l)\phi^l = \psi$. In principle any distribution that fulfills Eq. (8) allows finding the inverse of $\mathbf{M}(l)$ as in Eq. (9). However in practice it has been noted that the statistical error does depend on the choice of the distribution. The one that gave the smaller noise was the S_1 distribution, in which the ψ_i are distributed as $\delta(\psi_i^* \psi_j - 1)$.¹¹ This means that the variables ψ_i can take random values $e^{i\alpha}$ on the unit circle with equal probability. Physical quantities can then be calculated according to Eq. (7) as

$$\langle Q_\lambda \rangle = -\frac{1}{\beta} \sum_{l=1}^P \text{Tr} \left(\langle \phi^l \psi^\dagger \rangle \frac{\partial \mathbf{M}(l)}{\partial \lambda} \right) \quad (10)$$

and therefore one finds overall linear scaling behavior.

We now substantiate the claims made on the scaling of the different algorithms introduced here with numerical calculations. Different model potentials have been investigated with satisfactory results. Here we report only one calculation done on a periodic potential constructed with a superposition of Gaussians $\sum_{l=1}^N -w e^{-(r-R_l)^2/\delta^2} \Theta(r_c^2 - (r-R_l)^2)$ where $w = 4$ a.u. and the cutoff radius is $r_c = 6$ a.u. The Gaussian centers R_l are arranged to form a cubic lattice and mimic a crystal of N atoms. The spacing between the atomic sites is taken to be 4δ with $\delta = 0.75$ a.u. Periodic boundary conditions are imposed throughout. The Trotter number is chosen to be $P = 256$ leading to an electronic temperature

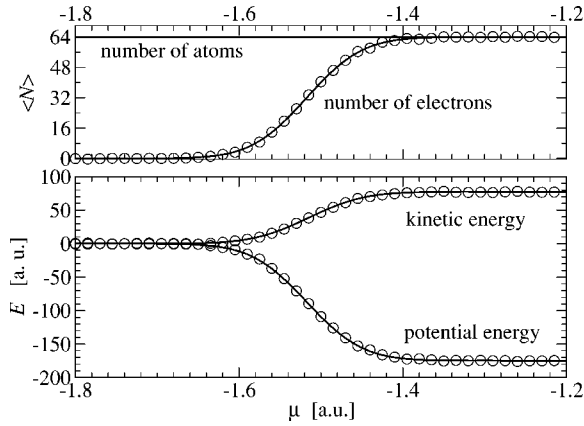


FIG. 1. Particle number, kinetic and potential energy of the Gaussian model as a function of the chemical potential. The circles show the results calculated with the stochastic linear scaling algorithm. The lines are obtained from the smallest eigenvalues calculated with an iterative diagonalization algorithm.

$T=7529$ K. This is rather small on the electronic energy scale and we have explicitly verified that it is close to the $T \rightarrow 0$ limit for the model.

The number of electrons, the kinetic and the potential energy as a function of the chemical potential are shown in Fig. 1. Upon increasing μ the states are filled with electrons. At $\mu \approx -1.5$ a.u. half of the states belonging to the first band are occupied and the system is metallic. For $\mu \approx -1.2$ a.u. the first band is completely filled and the model behaves as a large band gap insulator. The number of electrons, the kinetic and the potential energy are calculated with the stochastic algorithm and an iterative routine which calculates the largest relevant eigenvalues for comparison.^{12–15} The agreement is excellent. In order to reach the required precision, averages needed to be taken over about 100 independent configurations.

In a stochastic evaluation it is important to keep track of the error which for each property A is given by $\sigma_A / \sqrt{N_{\text{MC}}}$, where σ_A is its mean square fluctuation and N_{MC} the number of Monte Carlo steps. A number of physical observables together with their estimated σ_A is given in Table I. It is seen that the variances for the different observables do not differ qualitatively for metals and insulators. Thus the number of Monte Carlo steps does not have to be larger for metallic systems than for insulators. We have to mention that if l is close to $P/2$ the condition number of $\mathbf{M}(l)$ can be big in the metallic case. This could in principle increase the number of

TABLE I. Average values and variances for the particle number N , the kinetic energy T , the potential energy V , and the density at an atom site ρ .

	Insulator		Metal	
	$\langle A \rangle$	σ_A	$\langle A \rangle$	σ_A
N	64.0	2.7	32.0	4.5
T	77.2	5.4	36.7	4.1
V	-174.7	5.0	-86.1	4.2




FIG. 2. Log-log plot of the CPU time $\tau(N)$ versus the number of atoms N . The measured slopes confirm that the scaling is $O(N)$, $O(N^2)$, $O(N^{7/3})$, and $O(N^3)$ for the stochastic algorithm, the iterative inversion, the banded LU decomposition, and the partial iterative diagonalization, respectively. The unit of time is $\tau_\infty = \tau(N)/N$ measured at large N . That is the asymptotic incremental cost of an extra atom in the $N \rightarrow \infty$ limit. The computation was performed on a 1.7 GHz Pentium4 xenon processor.

iterations needed to solve $\mathbf{M}(l)\phi_j = \psi_j$. In practice however this did not lead to significant problems and the performance of the algorithm in the half-filled and filled case are very similar. For fixed Trotter number P the algorithm scales linearly with the number of grid points as the lattice constant δ is reduced.

We now compare the scaling of the different algorithms introduced here with a standard diagonalization procedure in which the largest relevant eigenvalues of are computed with sparse matrix diagonalization techniques.^{12–15} In Fig. 2 we see that the predictions made on the scaling of the different algorithms as a function of system size are confirmed by actual calculations.

Comparing the performance of the stochastic method with numerically exact methods is not easy since the performance of the method depends on the accuracy required. Here, in comparing the performance for different system sizes we have instead kept the number of Monte Carlo steps constant. Had we kept the relative accuracy constant this would have led to sublinear scaling due to the self-averaging properties of the larger systems. With this caveat from Fig. 2 one can see that our new algorithm based on stochastic matrix inversion techniques scales linearly, while the algorithm based on matrix diagonalization shows a cubic scaling. For the model chosen the crossing point at which our method becomes more efficient is $M \approx 38400$, which corresponds to 600 atoms. It must be stressed that we have taken the worst case scenario since we are here dealing with metals and P has to be taken larger to reach the $T=0$ limit. At full filling convergence is reached almost for $P=128$. For this P value the crossing point occurs at $N=450$ atoms. For different systems this value can vary because it might be necessary to choose a larger value for the Trotter number or the acceptable statistical error should be smaller. Still our method can be expected to become more efficient than cubic scaling algorithms at system sizes of a few hundred atoms. It should be also men-

tioned that the algorithm is trivially parallelizable. Furthermore increasing the number of electrons while keeping the volume constant does not increase the computational cost of our stochastic approach. In contrast the deterministic methods at constant volume scale quadratically with the number of electrons. Another advantage of the present method is memory saving, which grows linearly with volume and not as the product of the number of electrons times the volume.

In order to apply this method to fully self consistent DFT calculation one must take into account the fact that the evaluation of the electron density is affected by a statistical error. This will be considered at a later stage. As it is now the method can be profitably applied to tight binding calculations and to the Harris functional approximation¹⁶ to DFT. It can also be extended to include the ionic positions in the sampling so as to obtain an efficient Monte Carlo method.

-
- ¹W. Yang, Phys. Rev. Lett. **66**, 1438 (1991).
²G. Galli and M. Parrinello, Phys. Rev. Lett. **69**, 3547 (1992).
³S. Goedecker, Rev. Mod. Phys. **71**, 1085 (1999).
⁴N. D. Mermin, Phys. Rev. **137**, A1441 (1965).
⁵R. G. Parr and W. Yang, *Density-Functional Theory of Atoms and Molecules* (Oxford Science, Oxford, 1989), p. 147.
⁶S. Goedecker and L. Colombo, Phys. Rev. Lett. **73**, 122 (1994).
⁷S. Goedecker and M. Teter, Phys. Rev. B **51**, 9455 (1995).
⁸D. Chandler and P. G. Wolynes, J. Chem. Phys. **74**, 4078 (1981).
⁹A. Alavi and D. Frenkel, J. Chem. Phys. **97**, 9249 (1992).
¹⁰W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran 77* (Cambridge University Press, Cambridge, 1992).
¹¹Shao-Jing Dong and Keh-Fei Liu, Phys. Lett. B **328**, 130 (1994).
¹²NAG routine F02FJF based on the routine SIMTZ, P. J. Nikolai, ACM Trans. Math. Softw. **5**, 118 (1970).
¹³B. N. Parlett, *The Symmetric Eigenvalue Problem* (Prentice-Hall, Englewood Cliffs, NJ, 1980).
¹⁴H. Rutishauser, Numer. Math. **13**, 4 (1969).
¹⁵H. Rutishauser, Numer. Math. **16**, 205 (1970).
¹⁶J. Harris, Phys. Rev. B **31**, 1770 (1985).