

Implicit purification for temperature-dependent density matrices

Anders M. N. Niklasson

Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA

(Received 25 August 2003; published 30 December 2003)

An implicit purification scheme is proposed for calculation of the temperature-dependent, grand canonical single-particle density matrix, given as a Fermi-Dirac operator expansion in terms of the Hamiltonian. The computational complexity is shown to scale with the logarithm of the polynomial order of the expansion, or equivalently, with the logarithm of the inverse temperature. The system of linear equations that arise in each implicit purification iteration is solved efficiently by a conjugate gradient solver. The scheme is particularly useful in connection with linear scaling electronic structure theory based on sparse matrix algebra. The efficiency of the implicit temperature expansion technique is analyzed and compared to some explicit purification methods for the zero temperature density matrix.

DOI: 10.1103/PhysRevB.68.233104

PACS number(s): 71.15.-m, 31.15.-p, 65.90.+i, 31.90.+s

Linear scaling electronic structure theory in combination with tight-binding, self-consistent Hartree-Fock or density-functional theory has become a very powerful tool for studying complex large material systems.^{1,2} There are several ways to achieve a computational cost that scales linearly with system size. Here we focus on methods based on the single-particle density matrix for band-gap materials, the matrix elements of which decay exponentially with overlap distance. For large systems the number of matrix elements above some numerical threshold therefore scales as $O(N)$. In these schemes the two major steps are the construction of the tight-binding, Fockian or Kohn-Sham Hamiltonian $H(\mathbf{r}, \mathbf{r}')$ and the calculation of the density matrix $\rho(\mathbf{r}, \mathbf{r}')$. The present article concerns aspects of the second problem, the construction of the density matrix.

The relation between the density matrix at $T=0$ and the Hamiltonian is given by the Heaviside step function³

$$\rho = \theta(\mu I - H). \quad (1)$$

In density-matrix schemes this relation is approximated by constructing ρ from H using sparse matrix algebra, where each major operation computationally scales linearly with the system size, thanks to $O(N)$ matrix sparsity. This can be achieved through constrained minimization schemes,^{4,5} spectral projections or purification methods,⁶⁻¹² or by an expansion of the temperature dependent Fermi-Dirac function or similar step function approximations.¹³⁻¹⁸ Contour integral representations of the Fermi distribution with Padé polynomials as resolvents, that can be calculated with linear scaling effort³¹ as well as combinations of various approaches have also been explored.¹⁹⁻²⁴

The quadratically convergent purification techniques have turned out to be some of the most efficient approaches for the construction of the density matrix, both in memory and speed,^{8,11,12,24} with a computational complexity, in terms of number of matrix-matrix multiplications necessary to reach convergence that scales linearly with the logarithm of the inverse band gap and the degree of expansion, and with a numerical error that scales linearly with the threshold.^{11,12} The Fermi-Dirac operator expansion methods based on Chebyshev expansion techniques are generally much slower,

with the computational cost scaling at best with the square root of the degree of expansion.^{25,18} However, these methods have an important advantage; they can account for a finite temperature distribution of the density matrix. Here we propose an expansion scheme that combines the low logarithmic complexity and quadratic convergence of the purification schemes with the finite temperature Fermi-Dirac distribution. We show how this can be accomplished by an implicit purification scheme, based on a Padé approximation of the rescaled Fermi-Dirac function, with a computational complexity that scales logarithmically with the expansion order, or equivalently, the inverse temperature.

The Fermi-Dirac distribution,³

$$\Phi_{\text{FD}}(\varepsilon, \mu, \beta) = \frac{1}{e^{\beta(\varepsilon - \mu)} + 1}, \quad (2)$$

occurs in statistical mechanics as the occupational distribution of fermions at finite temperatures. It converges to a step function with the step formed at the chemical potential μ when $T \rightarrow 0$. The temperature-dependent grand canonical density matrix is formally given by the operator relation

$$\rho(\beta) = \Phi_{\text{FD}}(H, \mu, \beta). \quad (3)$$

The single-particle energy of a fermion system at a finite temperature is given by

$$E_s = \text{Tr}[H\rho(\beta)] = \sum_{i,j} \langle \phi_i | H | \phi_j \rangle \langle \phi_j | \rho(\beta) | \phi_i \rangle, \quad (4)$$

in some set of basis functions ϕ_i . In this formulation the expression for $\rho(\beta)$ does not have to be calculated explicitly; instead the Fermi-Dirac function can be expanded in Chebyshev polynomials. The Chebyshev expansion technique is one of the most efficient ways of approximating a function and the Chebyshev functions $T_n(x)$ obey a simple two-step recurrence formula ($T_0 = 1$, $T_1 = x$),

$$T_{n+1}(x) = 2xT_n(x) - T_{n-1}(x). \quad (5)$$

The products $\Phi_{\text{FD}}(H, \mu, \beta) | \phi_i \rangle$ can be calculated efficiently using the two-step recurrence formula using only matrix-

vector multiplications.^{13–17} The Chebychev expansion technique has many advantages: for example, the costly matrix-matrix multiplications are avoided, and error accumulation is small. However, compared to linear scaling purification techniques it is fairly inefficient.²⁴ The problem is the slow linear increase in polynomial order as a function of iterations in the two-step recurrence formula. To improve the efficiency, Liang *et al.*¹⁸ recently suggested an alternative approach, where the expansion polynomials are not calculated by the two-step formula, but by a direct expansion. This can be achieved with a computational complexity that scales with the square root of the polynomial order of the expansion $O(\sqrt{n})$. This limit is optimal for a general polynomial.²⁵ However, by choosing the expansion with a particular set of Padé polynomials, we will show how the computational complexity can be reduced even further, scaling only linearly with the logarithm of the polynomial order $O(\ln n)$ or the inverse temperature $O(\ln \beta)$.

There are at least 19 different ways to calculate matrix exponentials.²⁶ Here we use one particular technique based on a Padé approximation. Consider the exponential function

$$e^x = (e^{x/n})^n = \left(\frac{e^{x/(2n)}}{e^{-x/(2n)}} \right)^n. \quad (6)$$

A Taylor expansion to first order gives

$$e^x = \lim_{n \rightarrow \infty} \left(\frac{2n+x}{2n-x} \right)^n. \quad (7)$$

This Padé approximation can be used in the Fermi-Dirac function and for the rescaled chemical potential and inverse temperature,³ $\mu' = 1/2$ and $\beta' = 4n$,

$$\Phi_{\text{FD}}(x, \frac{1}{2}, 4n) \approx \frac{(1-x)^n}{x^n + (1-x)^n}. \quad (8)$$

At higher values of n the approximation becomes increasingly better. The choice $\mu' = 1/2$ is made to center the step of the Fermi-Dirac function at $x = 1/2$. In the interval $[0,1]$ the approximation is a continuously decreasing function with a maximum of 1 at $x=0$ and a minimum of 0 at $x=1$. This is the interval in which the temperature-dependent density matrix has its eigenvalues and it is the interval where the Fermi-Dirac distribution is well approximated already at fairly high temperatures. It is therefore the interval around which we chose to perform the expansion. This choice requires an initial rescaling of the Hamiltonian spectra around the interval $[0,1]$. Let

$$G_n(x) = \frac{x^n}{x^n + (1-x)^n}. \quad (9)$$

The Fermi-Dirac function for $\mu' = 1/2$ and $\beta' = 4n$ in the interval $[0,1]$ is approximated by

$$\Phi_{\text{FD}}(x, \frac{1}{2}, 4n) = [e^{4n(x-1/2)} + 1]^{-1} \approx 1 - G_n(x). \quad (10)$$

The polynomial order n of the Padé approximation is proportional to the inverse temperature since $n = \beta'/4$. This means

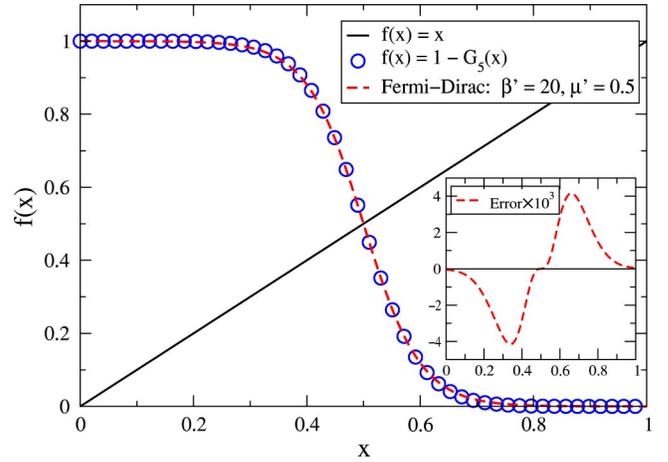


FIG. 1. The Fermi-Dirac distribution (dashed line) compared to the approximation $1 - G_5(x)$ (circles) in Eq. (10). The inset shows the error, $\text{Error} \times 10^3 = (\Phi_{\text{FD}}(x, 1/2, 20) - [1 - G_5(x)]) \times 10^3$.

that the lower the temperature the better the approximation. In practice, however, the approximation at the normalized energy interval $[0,1]$ is already very good at orders as low as $n \approx 5$. An example given in Fig. 1, which shows the Padé approximation $1 - G_5(x)$, is virtually identical to the corresponding Fermi-Dirac function with $\beta' = 20$ and $\mu' = 1/2$. The inset shows the error. With the interval $[0,1]$ equal to 1 Ry this example corresponds to a temperature of 7894 K. At lower temperatures the approximation becomes increasingly better. The Padé approximation in Eq. (10) is only one alternative, but, as will be shown below, it turns out to be particularly simple and efficient.

A major advantage with the Padé approximation in Eq. (10) is how efficiently we can calculate high orders of G_n . The computational complexity is very low, thanks to the iterative relation

$$G_{k \times l}(x) = G_k[G_l(x)]. \quad (11)$$

In an operator expansion this corresponds to purifications, projecting the eigenvalues towards 0 and 1. In contrast to a more general polynomial expansion such as the Chebychev expansion, which computationally scales at best with the square root of the polynomial order, $O(\sqrt{n})$,^{18,25} or as $O(n)$ if the two-step recurrence formula is used, the iterative relation above makes it possible to reach the same order of expansion in only $O(\ln n)$ steps. The same low logarithmic complexity is found generally in purification expansion schemes that are based on iterative spectral projections. The Padé approximation of the Fermi-Dirac distribution can thus be used in a highly efficient purification scheme which calculates the finite temperature density matrix $\rho(\beta)$ at a set of normalized inverse temperatures $\beta' = 4n$. The purification algorithm can be described by

$$X_1 = F(H, \mu),$$

$$X_{i+1} = G_m(X_i), \quad i = 1, 2, \dots, \log_m(n), \quad (12)$$

$$\rho(\beta) = I - X_{i+1}.$$

The expansion order n and thus the normalized inverse temperature β' must be chosen so that the number of iterations $\log_m(n)$, is an integer. The function

$$F(H, \mu) = \alpha(H - \mu I) + 0.5I \quad (13)$$

is a normalization function that rescales all the eigenvalues of H to the interval $[0,1]$, with the chemical potential μ shifted to $\mu' = 1/2$. The chemical potential and spectral bound must thus be known in advance. The normalization factor

$$\alpha \approx \frac{1}{2} \min[(\mu - H_{\min})^{-1}, (H_{\max} - \mu)^{-1}], \quad (14)$$

rescales the spectra and sets the temperature scale. The implicit purification scheme converges to the zero-temperature density matrix for any value of α , but the convergence is faster and the approximation is more accurate at higher temperatures if α is chosen to normalize the spectra around $[0,1]$. The spectral bounds H_{\max} and H_{\min} can be estimated by for example, Lanczos' algorithm or Gersgorin circles. Generally we have that the temperature³ $T = 1/(\alpha k_B 4n)$, where n is the accumulated expansion order in Eq. (12).

Because of the rational form of $G_m(X_i)$ the purification scheme is implicit. Assuming a finite orthogonal basis representation, a set of linear equations in Eq. (12) has to be solved in each step for $i = 1, 2, \dots, \log_m(n)$,

$$[X_i^m + (I - X_i)^m]X_{i+1} = X_i^m, \quad (15)$$

which is given from the second step in Eq. (12) and from the definition of G in Eq. (9) along with its nested iterative expansion property given in Eq. (11). Here we find another major advantage with our particular choice of Padé approximation. The left side system matrix $A_i = [X_i^m + (I - X_i)^m]$ is symmetric and positive definite for symmetric X_i 's with their spectra belonging to $[0,1]$. In fact, with increasing i , the system matrix A_i converges to the identity matrix I . The implicit equations are therefore very well suited for solutions with the linear conjugate gradient method,²⁷ that in turn, can efficiently exploit the close approximation of X_i to the unknown columns of X_{i+1} , which becomes increasingly more accurate and efficient towards the last iterations. Another possibly efficient alternative is the application of the sparse approximate inverse (AINV) (Refs. 28,19) that can be expected to work well for this particular problem. However, this approach has not been explored in the present study.

Alternative implicit purification schemes can also be derived from various sign matrix expansions.^{29,30} Sign matrix expansions are equivalent to purification. The only difference is that spectral projections are performed in the interval $[-1,1]$ instead of $[0,1]$, as in the case of purification.

To analyze the efficiency of the algorithm compared with explicit purification schemes, we have chosen an $N \times N$ model Hamiltonian with N random diagonal elements. The overlap elements decay exponentially as a function of site separation on a randomly distorted simple cubic lattice. This model represents a Hamiltonian of an insulator that might occur, for example, with a Gaussian basis set in density-functional theory or in various tight-binding schemes. The convergence is mainly determined by the occupation and the

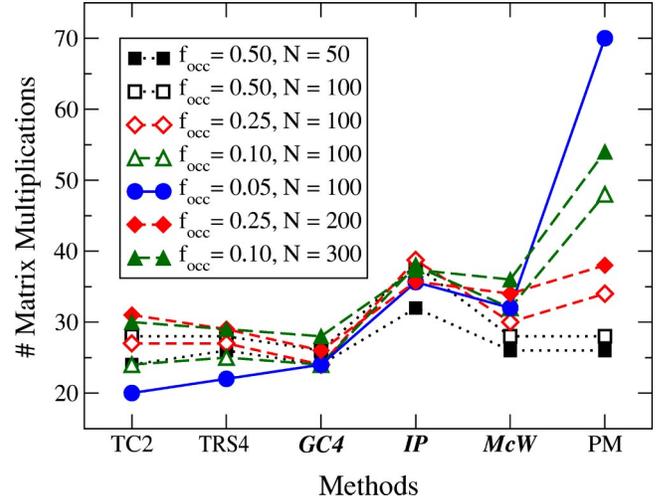


FIG. 2. Computational cost for various schemes at $T \approx 0$. Grand canonical schemes requiring prior knowledge of μ are written with bold italics. The N eigenvalues are uniformly distributed in $[0,1]$ and the band gaps are therefore $\Delta g = 1/N$ independent of the fractional occupation $f_{\text{occ}} = N_e/N$.

band gap. The test Hamiltonian was therefore modified such that all N eigenvalues were uniformly distributed on $[0,1]$. In this case the band gap $\Delta g = 1/N$, independent of the fractional occupation $f_{\text{occ}} = N_e/N$. This simplifies the analysis and comparison of the different methods, which otherwise are hard to perform for a less idealized set of material systems. After each iteration a numerical threshold $\tau = 1.0 \times 10^{-7}$ was applied and convergence was determined when the error in energy $|E_{\text{approx}} - E_{\text{exact}}| < 1.0 \times 10^{-5}$.³ The convergence criterion corresponds in practice to $T = 0$. A comparison at $T \approx 0$ is necessary since the explicit purification schemes used in the comparison only give the zero-temperature density matrix. At room temperature the computational effort with the implicit purification scheme is only slightly reduced because of the rapid convergence. The computational complexity was measured in number of matrix-matrix multiplications, where N conjugate gradient steps, i.e., N matrix-vector multiplications, are counted as one matrix-matrix multiplication.

Figure 2 shows the computational cost for various occupation factors. The implicit purification scheme of order two IP, i.e., with $m = 2$ in Eq. (15), using X_i as initial approximations to X_{i+1} , is compared to the trace correcting scheme with second-order polynomials (TC2) by Niklasson,¹¹ the trace resetting asymmetric fourth-order method (TRS4) by Niklasson *et al.*,¹² the grand canonical scheme with fourth-order projections (GC4) by Niklasson,¹¹ the grand canonical McWeeny purification scheme (McW) (Refs. 6,8), and finally the canonical scheme (PM) by Palser and Manolopolous.⁸ The grand canonical schemes that require prior knowledge of the chemical potential are indicated in the figure by bold italics.

For small band gaps, i.e., high values of N , and with prior knowledge of μ , the asymmetric GC4 method is the most efficient technique. The best performing schemes that require no prior knowledge of μ are the TC2 and TRS4 schemes.

The TC2 scheme is more memory efficient since it only needs to calculate a second-order polynomial in each iteration and intermediate storage needed in higher-order expansions is avoided. However, it cannot deal with degeneracy and fractional occupancy, which are addressed with the TRS4 scheme.¹² At low occupation the PM scheme becomes very inefficient. This sensitivity is not seen for any of the other schemes. The proposed implicit purification scheme is slower than the alternative explicit purification schemes except for the PM scheme at low occupancies. However, it is the only method that, for only a slightly increased computational cost, correctly gives the temperature-dependent Fermi-Dirac distribution of the single-particle eigenstates. The implicit purification scheme scales with the logarithm of the expansion order. This is an important improvement over previous Fermi-Dirac operator expansion methods. Whereas in general, a polynomial can be calculated with a computational

cost scaling at best with the square root of the order of the polynomial,²⁵ we restrict the polynomial approximation to a nested form $f(f(\dots f(x)\dots))$. In this case a high order can be reached much more efficiently than for the general form. This is the key idea behind purification expansions.

In summary, we have proposed an implicit purification scheme for the calculation of the temperature-dependent single-particle density matrix given as a Fermi-Dirac operator expansion in terms of the Hamiltonian. The method is useful in connection with linear scaling electronic structure theory and it has a computational complexity that scales with the logarithm of the inverse temperature $O(\ln \beta)$ or as the logarithm of the polynomial expansion order $O(\ln n)$.

Discussions with Matt Challacombe, Eric Chisolm, Siobhan Corish, Stefan Goedecker, C. J. Tymczak, and John Wills are gratefully acknowledged.

-
- ¹S. Goedecker, *Rev. Mod. Phys.* **71**, 1085 (1999).
²S.Y. Wu and C.S. Jayanthi, *Phys. Rep.* **358**, 1 (2002).
³The notation θ is the Heaviside step function, μ is the chemical potential, and μ' is the shifted chemical potential ($\mu'=1/2$), $\beta=1/(k_B T)$ is the inverse temperature, and the normalized discrete $\beta'=4n=\beta/\alpha$, where α is the scaling factor in Eq. (14). $I=\delta(r-r')$, $H=H(\mathbf{r},\mathbf{r}')$, $\rho=\rho(\mathbf{r},\mathbf{r}')$, $\text{Tr}(A)=\int d\mathbf{r}A(\mathbf{r},\mathbf{r})$, $AB=\int d\mathbf{r}''A(\mathbf{r},\mathbf{r}'')B(\mathbf{r}'',\mathbf{r}')$, the density $n(\mathbf{r})=\rho(\mathbf{r},\mathbf{r})$, the single particle energy $E_s=\text{Tr}(H\rho)$. N is the total number of states and N_e is the number of occupied states. The fractional occupation factor $f_{\text{occ}}=N_e/N$. For nonorthogonal representations see, for example, Ref. 19.
⁴X.-P. Li, W. Nunes, and D. Vanderbilt, *Phys. Rev. B* **47**, 10 891 (1993).
⁵W. Kohn, *Phys. Rev. Lett.* **76**, 3168 (1996).
⁶R. McWeeny, *Rev. Mod. Phys.* **32**, 335 (1960).
⁷W.L. Clinton, A.J. Galli, and L.J. Masa, *Phys. Rev.* **177**, 7 (1969).
⁸A.H.R. Palser and D.E. Manolopoulos, *Phys. Rev. B* **58**, 12 704 (1998).
⁹A. Holas, *Chem. Phys. Lett.* **340**, 552 (2001).
¹⁰A.M.N. Niklasson, C.J. Tymczak, and H. Röder, *Phys. Rev. B* **66**, 155120 (2002).
¹¹A.M.N. Niklasson, *Phys. Rev. B* **66**, 155115 (2002).
¹²A.M.N. Niklasson, C.J. Tymczak, and M. Challacombe, *J. Chem. Phys.* **118**, 8611 (2003).
¹³S. Goedecker and L. Colombo, *Phys. Rev. Lett.* **73**, 122 (1994).
¹⁴R.N. Silver and H. Roder, *Int. J. Mod. Phys. C* **5**, 735 (1994).
¹⁵L.W. Wang, *Phys. Rev. B* **49**, 10 154 (1994).
¹⁶R.N. Silver, H. Roder, A.F. Voter, and J.D. Kress, *J. Chem. Phys.* **124**, 115 (1996).
¹⁷R. Baer, and M. Head-Gordon, *J. Chem. Phys.* **107**, 10003 (1997).
¹⁸A. Liang, C. Saravanan, Y. Shao, R. Baer, A. T. Bell, and M. Head-Gordon, *J. Chem. Phys.* **119**, 4117 (2003).
¹⁹M. Challacombe, *J. Chem. Phys.* **110**, 2332 (1999).
²⁰D.R. Bowler and M.J. Gillan, *Comput. Phys. Commun.* **120**, 95 (1999).
²¹T. Helgaker, H. Larsen, J. Olsen, P. Jorgensen, *Chem. Phys. Lett.* **327**, 397 (2000).
²²M. Head-Gordon, Y.H. Shao, C. Saravanan, and C.A. White, *Mol. Phys.* **101**, 37 (2003).
²³Y.H. Shao, C. Saravanan, M. Head-Gordon, and C.A. White, *J. Chem. Phys.* **118**, 6144 (2003).
²⁴A.D. Daniels and G.E. Scuseria, *J. Chem. Phys.* **110**, 1321 (1999).
²⁵M.S. Paterson and L.J. Stockmyer, *SIAM J. Comput.* **2**, 60 (1973).
²⁶C. Moler and C. Van Loan, *SIAM Rev.* **45**, 3 (2003).
²⁷M.R. Hestenes and E. Stiefel, *J. Res. Natl. Bur. Stand.* **49**, 409 (1952).
²⁸M. Benzi, C.D. Meyer, and M. Tuma, *SIAM J. Sci. Comput. (USA)* **17**, 1135 (1996).
²⁹C.S. Kenney and A.J. Laub, *SIAM J. Matrix Anal. Appl.* **2**, 273 (1991).
³⁰G. Beylkin, N. Coult, and M. Mohlenkamp, *J. Comput. Phys.* **152**, 32 (1999).
³¹S. Goedecker, *Phys. Rev. B* **48**, 17 573 (1993).