PHYSICAL REVIEW B **66**, 155115 (2002)

Expansion algorithm for the density matrix

Anders M. N. Niklasson
Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

(Received 28 June 2002; published 24 October 2002)

A purification algorithm for expanding the single-particle density matrix in terms of the Hamiltonian operator is proposed. The scheme works with a predefined occupation and requires less than half the number of matrix-matrix multiplications compared to existing methods at low ($<10\%$) and high ($>90\%$) occupancies. The expansion can be used with a fixed chemical potential, in which case it is an asymmetric generalization of and a substantial improvement over grand canonical McWeeny purification. It is shown that the computational complexity, measured as the number of matrix multiplications, essentially is independent of system size even for metallic materials with a vanishing band gap.

DOI: 10.1103/PhysRevB.66.155115 PACS number(s): 71.15.−m, 02.70.−c, 31.15.−p

I. INTRODUCTION

Theoretical predictions of material properties of complex systems consisting of millions of atoms are often limited not by theory but by calculational techniques. Recently there has been a large effort to develop numerical methods that computationally scale linearly with system size.[1] The techniques may play an important role in a broad spectrum of science such as molecular biology, materials science, chemistry, and nanotechnology. Several of the linear scaling schemes are based on the single-particle density matrix that can be used in order to calculate the energies and densities that occur in self-consistent field theories. The construction of the density matrix is used as an alternative to solving an eigenvalue problem. For large complex systems within a sparse matrix representation this approach can be performed more efficiently and instead of a cubic scaling the computational cost scales linearly with the system size.[1] The matrix sparsity is essential for the success of density matrix schemes. For materials with a band gap the real-space representation of the density matrix is sparse[2–4] due to a finite correlation length, which is usually referred to as nearsightedness.[5] However, within other representations, such as a multiresolution wavelet basis or a group-renormalization representation, the density matrix is sparse also for metallic systems.[6–9]

Most techniques for constructing the density matrix can be seen as a polynomial expansion of the density matrix ρ_0 in terms of the Hamiltonian operator H. In an iterative approach this expansion can be formulated as

$$X_0 = P_0(H),$$

$$X_n = P_n(H, X_{n-1}), \quad n = 1,2,\ldots,$$

$$\rho_0 = \lim_{n\to\infty} X_n . \tag{1}$$

The projection polynomials $P_n(H,X_n)$ are chosen to achieve a rapid convergence under the conditions of commutation, $[H,\rho_0]=0$, idempotency, $\rho_0^2=\rho_0$, and particle conservation, $\mathrm{Tr}[\rho_0]=N_e$. They may either be chosen from a constrained conjugate gradient minimization of the energy functional $\mathrm{Tr}[\rho H]$ (Refs. 5 and 10–20) or as a fast expansion of the step function $\theta(\mu I - H)$ centered at the chemical potential μ or, for finite temperatures, the Fermi-Dirac distribution.[9,19–25] Each computational step consists of matrix-matrix additions, subtractions, and multiplications. The problem is to find a rapidly convergent expansion that minimizes the number of matrix-matrix multiplications, since these operations are the most time consuming.[26,27]

The efficiency of the different density-matrix schemes varies depending on the particular characteristics of the problem such as the existence of a band gap, a predefined chemical potential, filling factor, self-consistency cycles, thresholding, basis set, and system size. In this paper we propose a purification algorithm for the construction of the density matrix that is simple, general, and rapidly convergent also for very large metallic systems. The method works with a predefined occupation and does not need the input or adjustment of the chemical potential. Only one previous purification strategy, recently developed by Palser and Manolopoulos (PM),[24] exists for this important problem. By using a starting guess X_0 with the trace equal to the occupation number and thereafter performing trace-conserving spectral projections, X_n converges to the correct density matrix without prior knowledge of the chemical potential. The PM scheme has an excellent performance compared to other methods.[24,20] However, due to the constraint of trace conservation, the method is inefficient at low and high partial occupancies. This is of great concern, for example, when using a multiresolution wavelet representation for metallic problems, since the fractional filling in this case is low. The same problem occurs with a minimal basis set at both high and low occupancies. A simple general algorithm that avoids this particular problem and still converges as or more rapidly, especially for very large problems, would therefore be of great interest.

II. TRACE-CORRECTING PURIFICATION

The method we propose is based on the continuously increasing purification polynomials with stationary end points in [0,1]:

$$P_m^{(a)}(x) = 1-(1-x)^m[1+mx],$$

$$P_m^{(b)}(x) = x^m[1+m(1-x)]. \tag{2}$$

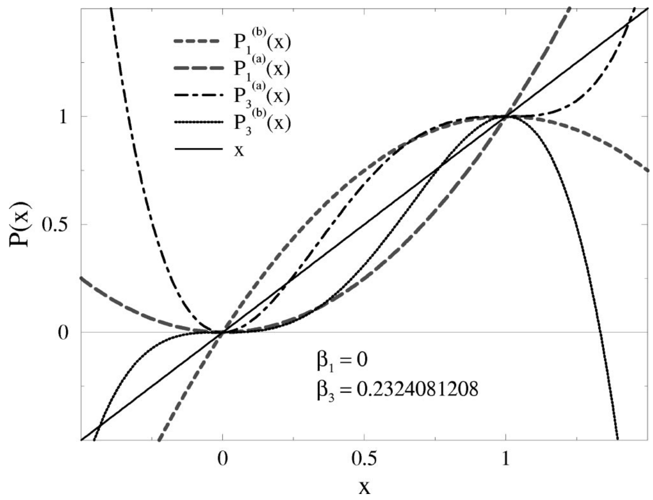

FIG. 1. Different projection polynomials for the adaptive expansion of the step function in Eq. (2).

Two examples of $P_m^{(a)}(x)$ and $P_m^{(b)}(x)$, for $m=1$ and 3, are displayed in Fig. 1. It can be shown that any combination of these polynomials in an iterative expansion converges to a step function for $x \in [0,1]$, i.e.,

$$\theta(x-\xi) = \cdots (P_m^{(a/b)}(P_m^{(a/b)}(x))\cdots, \tag{3}$$

with the step $\xi \in [\beta_m, 1-\beta_m]$. Here β_m is the inflection point of $P_m^{(a)}(x)$, i.e., where $P_m^{(a)}(\beta_m) = \beta_m$, $0 < \beta_m < 1$, and $(1-\beta_m)$ is the inflection point of $P_m^{(b)}(x)$. The convergence towards a step function can be understood from the fact that for each new iteration the new function will still be continuously increasing, but with an increasing number of vanishing derivatives at the end points. The asymmetry in the number of vanishing derivatives determines the position of the step. The choice $m=2$ corresponds to the McWeeny polynomial.[21] In this symmetric case $\beta_2 = 1/2$, and a step can only be formed at $\xi = 0.5$. The occupation of an operator X can be modified such that

$$\mathrm{Tr}[P_m^{(a)}(X)] \geqslant \mathrm{Tr}[X], \quad \varepsilon(X) \in [\beta_m, 1],$$

$$\mathrm{Tr}[P_m^{(b)}(X)] \leqslant \mathrm{Tr}[X], \quad \varepsilon(X) \in [0, 1-\beta_m], \tag{4}$$

where $\varepsilon(X)$ are the eigenvalues of X. For $m=1$ the reverse situation holds, with switched inequalities. With $m \neq 2$ we can apply the polynomials of Eq. (2) in the expansion of the density matrix, Eq. (1), such that each step adjusts for the occupation of X_n. In this way an expansion is created that converges to the density matrix with the correct occupation, i.e., $\rho_0 = \theta(\mu I - H)$ with $\mathrm{Tr}[\rho_0] = N_e$, but without *a priori* knowledge of μ. The algorithm (for $m>2$) is given by this pseudocode:

```
function $\rho_0(H, N_e, \text{error limit})$
estimate $\varepsilon_0(H)$, $\varepsilon_N(H)$
$X_0 = (1-2\beta_m)(\varepsilon_N I - H)/(\varepsilon_N - \varepsilon_0) + \beta_m I$
while error > error limit
    if $\mathrm{Tr}[X_n] - N_e < 0$
        $X_{n+1} = P_m^{(a)}(X_n)$
        else
        $X_{n+1} = P_m^{(b)}(X_n)$
        end
        estimate error
    end
```

$$\rho_0 = X_n. \tag{5}$$

For $m=1$ the trace condition has to be reversed to ">."

The scheme can be described as follows: First the Hamiltonian is normalized to an initial matrix X_0 with all its eigenvalues $\varepsilon(X_0) \in [\beta_m, 1-\beta_m]$. The constants ε_0 and ε_N are the lowest and highest eigenvalues of H, respectively. These can be approximated by, for example, Gersgorin estimates or the Lanczos method, with only a small extra computational cost.[20,24] A necessary criterion for convergence is that the unknown chemical potential of the normalized initial matrix $\mu(X_0) \in [\beta_m, 1-\beta_m]$. For intermediate occupancy, provided $\mu(X_0) \in [\beta_m, 1-\beta_m]$, β_m can be set to zero in the starting guess. This usually reduces the number of iterations by one or two steps. The improvement has not been used in the present study. After initializing X_0 the projections $P_m^{(a)}(X_n)$ or $P_m^{(b)}(X_n)$ are performed, adjusting the occupancy and expanding a step function at the same time. The iteration stops when some appropriate error estimate is less than a predefined error limit. Note that a high-order expansion may lead to too fast a convergence, making an adjustment of the occupation impossible. We may also use combinations with different values of m as well as other asymmetric purification polynomials. Any set of asymmetric continuously increasing polynomials in $[0,1]$ with stationary points at 0 and 1 can be used equivalently. The presented algorithm cannot handle problems with degenerate eigenstates at μ. The algorithm would still converge, but to the wrong density matrix, since the degenerate states would split due to numerical noise. This differs the presented trace correcting purification scheme from the PM scheme, which correctly can treat the case of degeneracy.

III. GRAND CANONICAL PURIFICATION

Since any combination of the expansion polynomials in Eq. (2) converges to a step function, we can use a predefined fixed expansion combination of $\theta(\mu I - H)$. In this case μ must be known, but the efficiency might be slightly improved compared to schemes working with a predefined occupancy. For example, we may use repetitions of the combination $P_{\mathrm{ex}}^{\mathrm{GC}}(x) = P_2^{(b)}(P_1^{(b)}(P_3^{(a)}(x)))$ or (for $m>1$) only $P_m^{(a)}(x)$, or only $P_m^{(b)}(x)$, in the combination

$$P_m^{\mathrm{GC}}(x,\overline{\mu})=\begin{cases}P_m^{(a)}(x), & \overline{\mu}\geqslant 1/2,\\ P_m^{(b)}(x), & \overline{\mu}\leqslant 1/2,\end{cases} \tag{6}$$

where $\overline{\mu}$ is the normalized chemical potential, $\overline{\mu}=(\mu-\varepsilon_0)(\varepsilon_N-\varepsilon_0)^{-1}$. We can now, as in Eq. (1), perform the expansion using the fixed repeated combination of $P_m^{(a)}(x)$ and $P_m^{(b)}(x)$ with the starting guess

$$X_0=\alpha(\mu I-H)+\beta I,$$

$$\alpha=\min\{\beta[\varepsilon_N-\mu]^{-1},(1-\beta)[\mu-\varepsilon_0]^{-1}\}. \tag{7}$$

The constant β is here determined by the inflection point of the repeated fixed polynomial combination, e.g., β_m or $(1-\beta_m)$ for $P_m^{\mathrm{GC}}(x,\overline{\mu})$. The approach can be seen as an asymmetric generalization of the grand canonical McWeeny purification scheme.[21,24,25] With $P_2^{\mathrm{GC}}(x,\overline{\mu})$ they are equivalent. The method is directly related to matrix sign function expansions.[23] The matrix sign function expansion is equivalent to the purification scheme via a trivial linear transform where the step function expansion is performed between -1 and $+1$ in the interval $[-1,1]$.

If the chemical potential is unknown, the density matrix may have to be recalculated with different values of μ until the occupation is correct. However, since the density matrix can be described as a superposition of outer products of the occupied eigenstates, we can adjust the occupation by adding or subtracting Hamiltonian eigenstates close to the chemical potential. A few of these states can be calculated efficiently, for example, using inverse power iterations.[28] In this way the occupation can be adjusted without a complete recalculation of the density matrix with a new shifted chemical potential. Moreover, in the case of a material with a band gap we do not need to have a very precise prior knowledge of μ as long as the estimate is somewhere in the gap. The grand canonical approach, with a predefined fixed μ, may thus be an efficient alternative in some special cases.

IV. EXAMPLES

To illustrate the efficiency of the expansion techniques we have constructed an $N\times N$ Hamiltonian test matrix $\overline{H}(i,j)$ with randomized off-diagonal elements decaying as $|i-j|^{-2}$ and with a uniform distribution of eigenvalues in [0,1]. Only the eigenvalue distribution of the Hamiltonian is of importance for the convergence. With $\overline{H}$ we have that the occupation factor $\lambda=N_e/N=\overline{\mu}=\mu$ and it is easy to compare grand canonical schemes with the trace-correcting or trace-conserving methods.

Figure 2 shows the number of matrix multiplications necessary to achieve an error $||X_n-\rho_0||_2\leqslant 10^{-9}$ as a function of the filling factor λ or chemical potential μ. The PM trace-conserving purification scheme is slow at low and high occupancy (since the slope at the inflection point tends to 1 as the inflection point approaches 0 or 1), whereas the new trace-correcting expansion algorithm, with $m=1$ (P_1), $m=3$ (P_3), and $m=5$ (P_5), has an overall fast convergence. For example, at 10% occupancy the new scheme with $m=3$ is about twice as fast compared to the PM scheme. There are essentially three reasons for the improved convergence: (i) a more optimized ratio between the number of matrix multiplications and the polynomial order, (ii) a faster increase of the number of vanishing derivatives at the stationary end points, and (iii) as will be shown below, a steeper slope at the inflection points, partly due to the asymmetry of the purification polynomials. The generalized grand canonical expansion with $m=3$ (P_3^{GC}) converges several steps faster compared to the grand canonical McWeeny (McW) method for the same reasons.

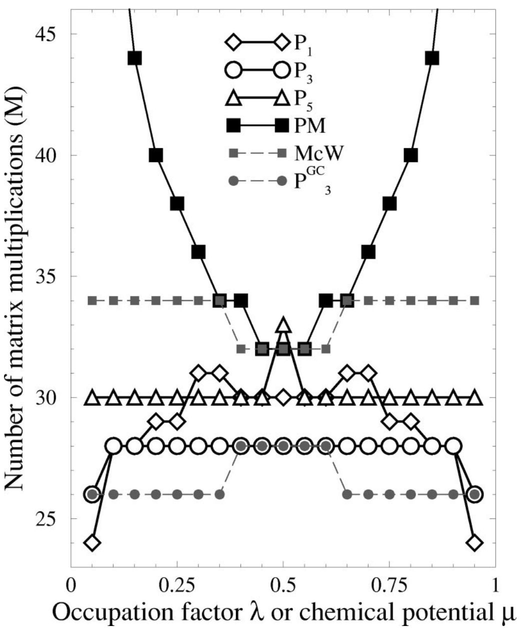

FIG. 2. The number of matrix-matrix multiplications M necessary to achieve a convergence $||X_n-\rho_0||_2\leqslant 10^{-9}$, as a function of the filling factor λ or, equivalently, the chemical potential μ, for $\overline{H}$ with $N=100$. PM corresponds to the result using the canonical trace-conserving purification scheme by Palser and Manolopoulos (Ref. 24). The open symbols P_m show the result of the expansion algorithm, Eq. (5). The small squares indicate the result of the grand canonical WcWeeny purification (McW) and $P_3^{\mathrm{GC}}(x,\overline{\mu})$ in Eq. (6).

V. SCALING

By varying N, i.e., the size of the Hamiltonian test matrix, we may see how the number of matrix multiplications necessary for convergence scales with system size or, equivalently, with the inverse gap at the chemical potential $\Delta\varepsilon_\mu=1/N$. The behavior is crucial for very large systems, especially if we wish to construct an expansion scheme that computationally scales linearly with the system size for metallic materials with a vanishing band gap. Figure 3 displays the number of necessary matrix multiplications as a function of ln(N). In the upper graph the results of McWeeny purification (McW) and the trace-conserving canonical purification (PM) are on top of each other. For this particular symmetrical case the two schemes are identical.[24] The graph indicates a stepwise linear relationship between the number of matrix multiplications M necessary for convergence and the logarithm

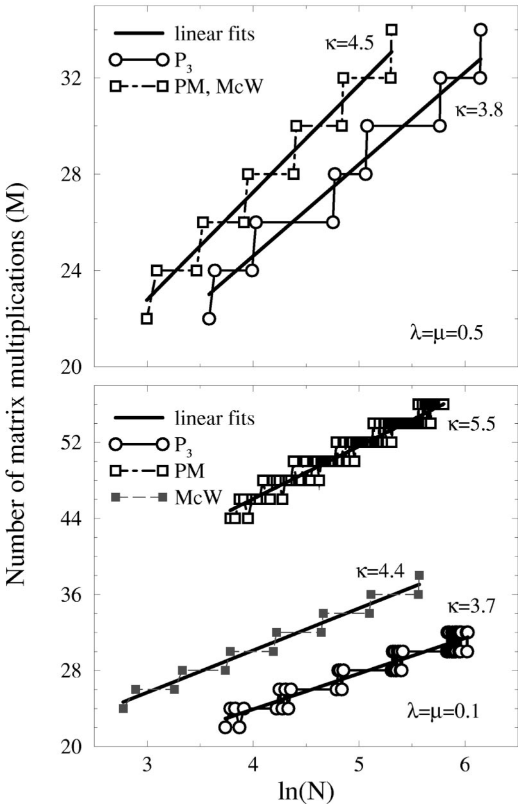

FIG. 3. The number of matrix-matrix multiplications M necessary to achieve a convergence $|E_n - E_0|/N \leqslant 10^{-9}$, as a function of ln($N$). Here $E_n = \mathrm{Tr}[HX_n]$.

of the system size. The relation can be approximated by the linear formula

$$M(\mu,N)=\alpha(\mu)+\kappa(\mu)\ln(N). \tag{8}$$

The least-squares fits of $M(\mu,N)$ are shown together with the values of κ. The expansions P_1, P_5, and P_3^{GC} perform equally well or slightly worse compared to P_3 and are not shown. The slopes determine the efficiency for very large systems and we find the best scaling for P_3.

The convergence is determined by the slowest converging eigenvalue $\gamma(X_0)$, closest above or below the chemical potential of the normalized initial matrix $\mu(X_0)$. This particular eigenvalue should either converge to 1 or 0. Since the purification polynomials are continuously increasing, preserving the order of the eigenvalues, all other eigenvalues converge faster. In the case of grand canonical purification, with a uniform distribution of N eigenvalues, $\gamma(N,X_0)\approx\beta \pm 1/(2N)$. By means of a linearization of the purification polynomial around β it can be shown that $\gamma(N,X_0) \approx \gamma(kN,X_1)$, where k is the derivative of the purification polynomial at the inflection point, e.g., $k = P_m^{\mathrm{GC}}{}'(\beta_m,\overline{\mu})$. Thus, increasing the number of states by $\Delta N = N(k-1)$ and adding one extra iteration leads to the same error in γ. If $\tilde{M}_\gamma$ is the number of multiplications necessary to achieve a fixed error of γ and p is the number of matrix multiplication in one iteration, we have that

$$[\tilde{M}_\gamma(N+\Delta N)-\tilde{M}_\gamma(N)]/(\Delta N)\approx p[N(k-1)]^{-1}. \tag{9}$$

Let M_γ be a continuous version of $\tilde{M}_\gamma$ such that

$$dM_\gamma/dN = p[N(k-1)]^{-1}. \tag{10}$$

Integration gives

$$M_\gamma(N)=\alpha+\frac{p}{k-1}\ln(N), \tag{11}$$

for some constant α. This approximate formula explains the linear relation in Eq. (8) and is a useful measure in optimizations of grand canonical expansions. Purification polynomials should be optimized on the criterion of matrix multiplications versus vanishing derivatives at the stationary fixed points and slope $\kappa_{\mathrm{est}} = p(k-1)^{-1}$ as estimated in Eq. (11). For example, the purification $P_3{}^{\mathrm{GC}}(X)$ requires two matrix multiplications in each iteration, it has three vanishing fixed point derivatives, and $\kappa_{\mathrm{est}} = 3.1$. This should be compared to McWeeny purification, which also requires two matrix multiplications in each iteration, but with only two vanishing fixed point derivatives and with a $\kappa_{\mathrm{est}} = 4.0$.

For materials with a band gap, the value $1/N$ should be replaced by the gap at the chemical potential. In this case the scaling with the logarithm of the system size vanishes and the number of necessary matrix multiplication for a predefined convergence accuracy is constant.

Notice that one may use different criteria for convergence such as the error per state, per atom, or total error. However, this has only a minor effect on the number of necessary matrix multiplications because of the very rapid rate of convergence close to idempotency, which, for example, is quadratic in the WcWeeny case.

VI. FIRST-PRINCIPLES PERFORMANCE

To further illustrate the performance of the expansion scheme we show the result of an implementation in the MONDOSCF suite of linear scaling self-consistent field programs.[17,29–31] Figure 4 displays the number of necessary matrix multiplications for clusters and strings of Li atoms and for different molecules. In the case of Li clusters and strings of Li atoms the systems are metallic in the sense that the gap vanishes in the limit of infinite number of atoms. This is thus a good test to check the linear relationship between the number of matrix multiplications and the logarithm of the system size. The P_1 scheme is efficient compared to the PM scheme, especially in the case of SiF_4 where the occupancies λ is high. This particular example illustrates the inefficiency of the PM scheme at high and low occupancies which is avoided with the trace-correcting algorithm. The weak logarithmic dependence between computational cost and system size for metallic systems, illustrated by the dashed lines in the figure, is also confirmed. Notice that an actual linear scaling is reached only if the number of nonzero elements of the density matrix grows linearly with system size. This can generally not be achieved within a real-space representation for metallic systems. Instead, as mentioned above, a multiresolution wavelet basis or a group-

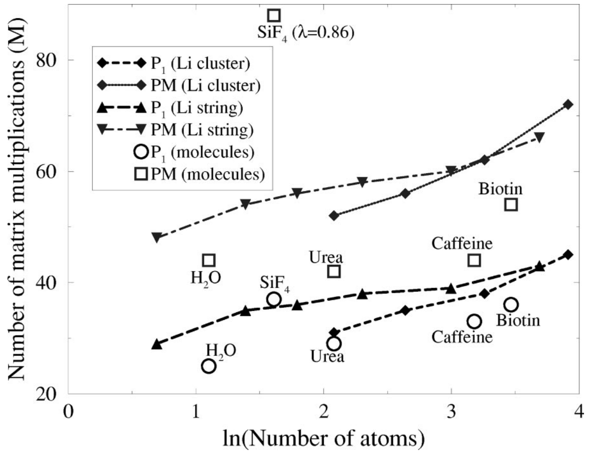

FIG. 4. The number of matrix multiplications M (after three self-consistency cycles with STO-3G or STO-6G basis sets) necessary to achieve a convergence $|E_n - E_{n-1}| \leqslant 10^{-7}$ a.u., as a function of the logarithm of the number of atoms. Occupation $\lambda \in [0.3,0.7]$ except for NiF$_4$. For the Li strings and clusters, as indicated by the dashed lines, the gap is vanishing, i.e., metallic, in the limit "number of atoms" $\rightarrow \infty$.

renormalization approach has to be applied. However, the number of matrix multiplications necessary for convergence should not be affected by a change of representation since the vanishing gap around the chemical potential will remain the same regardless of the representation, given, in the case of wavelets, via a biorthogonal transformation of the Hamiltonian.[8,9]

The P_1 scheme implemented in the MONDOSCF programs has two major advantages compared to other schemes: (i) It requires less memory compared to higher-order schemes since only second-order polynomials are used and intermediate matrix products do not have to be stored. (ii) It is less complex and only matrix squares have to be calculated. A specially designed algorithm for matrix squares can possibly be made more efficient than a general matrix product algorithm.

VII. DISCUSSION

The expansion scheme and the convergence analysis illustrated and argued for here provide a basis for understanding purification algorithms and their efficiency, and it shows that the computational complexity, as measured in the number of matrix multiplications, essentially is independent of system size even for metallic systems. However, if the additional problem of thresholding is included, which can be performed either via a finite cutoff radius truncation or via a numerical threshold, the computational complexity as a function of system size, within some required numerical accuracy, becomes far more difficult to analyze and practical experience may be the only way to understand the efficiency.

In the alternative construction of the density matrix using a constrained functional minimization, as devised by Li *et al.*,[11] McWeeny purification is used to impose idempotency. The asymmetric polynomial expansions proposed here may serve as a possible alternative.

VIII. SUMMARY

In summary we have proposed an algorithm for expanding the single-particle density matrix in terms of the Hamiltonian that is simple, general, and with a computational complexity essentially independent of system size even for very large metallic systems with a vanishing band gap. If the expansion is used together with a fixed chemical potential, it was shown to be an asymmetric generalization of grand canonical McWeeny purification. The algorithm is a substantial improvement of previous schemes and provides together with the presented convergence analysis a framework for the understanding and optimization of purification.

ACKNOWLEDGMENTS

Discussions and support from Matt Challacombe and C. J. Tymczak are gratefully acknowledged. I am also very thankful for editorial help from Eric Chisolm.

---

[1] S. Goedecker, Rev. Mod. Phys. **71**, 1085 (1999).
[2] W. Kohn, Phys. Rev. **115**, 809 (1959).
[3] R. Baer and M. Head-Gordon, Phys. Rev. Lett. **79**, 3962 (1997).
[4] U. Stephan, R.M. Martin, and D.A. Drabold, Phys. Rev. B **62**, 6885 (2000).
[5] W. Kohn, Phys. Rev. Lett. **76**, 3168 (1996).
[6] G. Beylkin, R. Coifman, and V. Rokhlin, Commun. Pure Appl. Math. **44**, 141 (1991).
[7] S.R. White, Phys. Rev. Lett. **69**, 2863 (1992).
[8] S. Goedecker and O.V. Ivanov, Phys. Rev. B **59**, 7270 (1999).
[9] A.M.N. Niklasson, C.J. Tymczak, and H. Röder (unpublished).
[10] A.H. Sameh and J.A. Wisniewski, SIAM (Soc. Ind. Appl. Math.) J. Numer. Anal. **19**, 1243 (1982).
[11] X.-P. Li, R.W. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10 891 (1993).
[12] A.E. Carlsson, Phys. Rev. B **51**, 13 935 (1995).
[13] E. Hernandez, M.J. Gillan, and C.M. Goringe, Phys. Rev. B **53**, 7147 (1996).
[14] A.D. Daniels, J.M. Millam, and G.E. Scuseria, J. Chem. Phys. **107**, 425 (1997).
[15] W. Yang, Phys. Rev. B **56**, 9294 (1997).
[16] U. Stephan and D.A. Drabold, Phys. Rev. B **57**, 6391 (1998).
[17] M. Challacombe, J. Chem. Phys. **110**, 2332 (1999).
[18] P.D. Haynes and M.C. Payne, Phys. Rev. B **59**, 12 173 (1999).

[19] D.R. Bowler and M.J. Gillan, Comput. Phys. Commun. **120**, 95 (1999).

[20] A.D. Daniels and G.E. Scuseria, J. Chem. Phys. **110**, 1321 (1999).

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

[22] S. Goedecker and L. Colombo, Phys. Rev. Lett. **73**, 122 (1994).

[23] C.S. Kenney and A.J. Laub, IEEE Trans. Autom. Control **40**, 1330 (1995).

[24] A.H.R. Palser and D.E. Manolopoulos, Phys. Rev. B **58**, 12 704 (1998).

[25] A. Holas, Chem. Phys. Lett. **340**, 552 (2001).

[26] M. Challacombe, Comput. Phys. Commun. **128**, 93 (2000).

[27] D.R. Bowler, T. Miyazaki, and M. Gillan, Comput. Phys. Commun. **137**, 255 (2001).

[28] I.C.F. Ipsen, SIAM Rev. **39**, 254 (1997).

[29] M. Challacombe and E. Schwengler, J. Chem. Phys. **106**, 5526 (1997).

[30] E. Schwengler, M. Challacombe, and M. HeadGordon, J. Chem. Phys. **106**, 9708 (1997).

[31] M. Challacombe, J. Chem. Phys. **113**, 10 037 (2000).