

Assessment of density matrix methods for electronic structure calculations*

Elias Rudberg[†] and Emanuel H. Rubensson[‡]

*Division of Scientific Computing, Department of Information Technology,
Uppsala University, Box 337, SE-751 05 Uppsala, Sweden*

(Dated: October 12, 2010)

Purification and minimization methods for computation of the one-particle density matrix are compared. This is done by considering the work needed by each method to achieve a given accuracy in terms of the difference to the exact solution. Simulations employing orthogonal as well as non-orthogonal versions of the methods are performed using both element magnitude and cutoff radius based truncation approaches. The results indicate that purification is considerably more efficient than the studied minimization methods even when a good starting guess for minimization is available. The computational cost of the studied minimization methods is observed to be significantly more sensitive to small band gaps than purification. An $\mathcal{O}(\sqrt{1/\xi})$ dependence on the band gap ξ is observed for minimization which can be compared to the $\mathcal{O}(\ln(1/\xi))$ dependence for purification. Minimization is found to perform at its best at 50% occupancy. Error control and stopping criteria are also discussed.

* To be published in the Technical Report series from the Department of Information Technology at Uppsala University, see <http://www.it.uu.se/research/publications/reports/>

[†]Electronic address: elias.rudberg@it.uu.se

[‡]Electronic address: emanuel.rubensson@it.uu.se

I. INTRODUCTION

Calculation of the one-electron density matrix for a given effective Hamiltonian matrix is one of the most computationally demanding operations in large scale electronic structure calculations using methods such as Hartree-Fock and Kohn-Sham density functional theory. For small systems, well-known linear algebra algorithms can be used to calculate the density matrix via a diagonalization of the effective Hamiltonian matrix. However, for large systems the cubical scaling of the diagonalization operation makes this approach unfeasible. Therefore, density matrix methods based on sparse matrix-matrix multiplications that can scale linearly with system size have been developed.

Basically two types of density matrix methods exist. In the first type, a minimization problem is solved using steepest descent or conjugate gradient iterations. Several such methods are based on the functional form proposed by Li et al. and Daw [1–4]. Other minimization variants exist, including the curvy step [5–7] and penalty functional [8–10] approaches. The second type of density matrix methods are based on the observation that the density matrix at zero temperature is given by a step function of the effective Hamiltonian matrix. This matrix function is approximated by a serial or recursive polynomial expansion. The serial expansion is typically given by a Chebyshev series [11–14]. In methods employing a recursive expansion, commonly referred to as density matrix purification [15–21] or sign matrix methods [22, 23], the polynomial is constructed by recursive application of low-order polynomials. Here, we shall focus on the recursive expansions or purification methods as they give a much more rapid convergence compared to the serial expansion approach [24].

Both minimization and purification methods allow for computation of the density matrix with linear scaling complexity. However, it is not obvious which method should be chosen in order to obtain a result with satisfactory accuracy with least possible use of computer resources. On one hand, the convergence of purification methods is generally much faster than that of minimization methods. On the other hand, in purification methods, contact with the effective Hamiltonian matrix which alone defines the exact solution, is lost already in the first iteration. This means that any accuracy that is lost during the purification iterations cannot be recovered. This is a problem the minimization methods do not have. So, which method is most efficient, minimization or purification, or perhaps a hybrid between the two?

We believe that when comparing density matrix methods, it is important to use a clear problem formulation that allows the methods to be compared on an equal footing. A consistent way to measure the error should be used for all methods, for example comparing the density matrix approximation to the exact solution. Previous comparisons exist in the literature [25–27]. However, one has typically not made sure that the computed results from different methods are at an equal level of accuracy, which manifests itself by different numbers of cycles in the self-consistent field optimizations. This makes it difficult to draw definite conclusions even for a specific system under study. Although the number of self-consistent field cycles may give some indication of the accuracy of density matrix methods, in our view this indirect way of assessing the accuracy is unsatisfactory. Instead, we focus on the isolated problem of constructing the density matrix for a given effective Hamiltonian matrix with as low computational cost as possible while achieving a desired accuracy.

This paper is organized as follows. In the following section, we formulate the problem and define the error measure used throughout the paper. After that, in Section III, we briefly describe the methods under consideration. Then, in Section IV we present test calculations comparing the considered methods. The results are further discussed in Section V where also hybrid schemes are

considered. Finally, our conclusions are presented in Section VI.

II. PROBLEM FORMULATION

We consider in this article the problem of computing the density matrix D for a given effective Hamiltonian F at zero electronic temperature defined by

$$\begin{aligned} F_{\perp} &= Z^T F Z \\ D_{\perp} &= \theta(\mu I - F_{\perp}) \\ D &= Z D_{\perp} Z^T \end{aligned} \quad (1)$$

where $\theta(x)$ is the Heaviside step function and μ is the so-called chemical potential, a value between eigenvalues corresponding to occupied and unoccupied eigenstates, respectively. The matrix Z is an inverse factor of the overlap matrix S such that $Z^T S Z = I$. Note that for an orthogonal basis set, $S = I$ and $D = \theta(\mu I - F)$. The conditioning of the problem depends on the gap ξ between eigenvalues corresponding to occupied and unoccupied eigenstates, commonly referred to as the band gap or HOMO-LUMO gap. For a fixed spectral width of F , the condition number is proportional to $1/\xi$ [28].

The minimization and purification methods discussed in this paper provide means to compute D with an effort that increases only linearly with system size. These methods can scale linearly provided that the involved matrices are sparse enough which usually requires that many matrix elements are neglected during the computation. The outcome is an approximate matrix \tilde{D} whose quality depends on the stopping criterion of the iterative process and on which matrix elements have been neglected.

We study the performance of each method by considering the computational cost needed to reach a certain accuracy in the computed density matrix \tilde{D} measured by

$$\|\tilde{D} - D\|_2 \leq \gamma. \quad (2)$$

Here, $\|\cdot\|_2$ denotes the spectral norm and γ is the desired accuracy. Measuring the accuracy in terms of the norm of the error matrix is consistent with the framework for error control in the self-consistent field procedure proposed in Ref. [29].

III. METHODS

In this section, we describe the minimization and purification methods considered in this article. Both classes of methods can be formulated to operate in orthogonal or non-orthogonal basis. In case of methods that operate in orthogonal basis, preparatory and final steps, consisting of the congruence transformations $F_{\perp} = Z^T F Z$ and $D = Z D_{\perp} Z^T$, are needed. For these methods, an inverse factor Z can be computed using for example inverse Cholesky decomposition [3, 4, 30, 31] or recursive inverse factorization [32]. In this paper, the inverse Cholesky factor is used. For methods operating in non-orthogonal basis, no inverse factor Z is used. Instead, the inverse overlap matrix S^{-1} enters the calculation.

In the following we use the notation D and D_{\perp} for the exact solution to (1) in non-orthogonal and orthogonal basis, respectively, while X and X_{\perp} denote the current density matrix approximation during the course of an iterative method.

A. LNV minimization methods

In 1993, Li, Nunes, and Vanderbilt (LNV) proposed minimization of the functional

$$\Omega(X_{\perp}) = \text{Tr}[(3X_{\perp}^2 - 2X_{\perp}^3)(F_{\perp} - \mu I)] \quad (3)$$

as a method to compute the density matrix [1]. A similar approach was independently proposed by Daw [2]. Since then, several variants have been presented that generalize or improve the original method in different ways. This includes reformulations to operate directly in non-orthogonal basis [33] and methods that automatically adjust the electron count without need to know the chemical potential μ in advance [3, 34]. In this paper, however, we consider the chemical potential μ as known and apply the LNV minimization methods to the shifted matrix $F - \mu I$. In our benchmark calculations we consider the original LNV method and a simplified version by Challacombe and their non-orthogonal counterparts denoted as follows:

- *Standard*: The original LNV method [1].
- *Standard (non-orth)*: The Nunes and Vanderbilt method of Ref. [33] with modifications as suggested by White et al. [35].
- *Simplified*: The simplified variant of LNV minimization proposed by Challacombe [4].
- *Simplified (non-orth)*: The non-orthogonal variant by Challacombe [4].

In all four variants, we employ the Fletcher-Reeves conjugate gradient method with analytic line searches. It is noteworthy that the analytic line searches constitute a significant part of the computational effort. Throughout this article the $\frac{1}{2}I$ starting guess is used for the LNV methods, unless stated otherwise.

An appealing feature of LNV methods is that errors do not accumulate during the iterations. However, errors are still introduced through the neglect of small matrix elements needed to obtain a linear scaling method.

Other minimization methods exist which use a parametrization of the density matrix that satisfies the idempotency and occupation count constraints. One approach is to use an exponential parametrization [5] but other choices exist as well [36]. So-called curvy steps are then taken on the idempotent manifold given by the parametrization. In case of an exponential parametrization, the new density matrix is evaluated by use of a truncated Baker-Campbell-Hausdorff expansion. If this expansion is evaluated to first order, the curvy step is in the same direction as an LNV step [7]. Evaluations to higher orders in principle lead to improved conservation of idempotency at the expense of a higher computational cost per step. Contrary to the LNV method, this method is not able to recover any idempotency that is lost during the iterations. The curvy step and related approaches are not included in our benchmark calculations.

B. Purification methods

Already in 1956, McWeeny suggested that roughly idempotent matrices could be refined using the polynomial $x^2(3 - 2x)$ [37]. In 1998, Palser and Manolopoulos [15] proposed purification as a method to compute the one-electron density matrix given by (1). In this method, McWeeny purification is applied after an appropriate initial transformation of the Hamiltonian matrix. Palser and Manolopoulos also presented an alternative, referred to as canonical purification, that does not

	prep	iter	final
<i>LNV minimization methods</i>			
Standard	4	5(2.5)	1.5
Standard (non-orth)	5	8.5(3.5)	0
Simplified	2.5	2(1)	1.5
Simplified (non-orth)	1.5	6(4.5)	0
<i>Purification methods</i>			
Trace-correcting	1.5	0.5	1.5
Trace-correcting (non-orth)	1.5	1.5	0
Canonical	1.5	1	1.5
Canonical (non-orth)	1.5	2	0

TABLE I: Number of matrix-matrix multiplications used by the considered methods per iteration (iter) and in preparatory (prep) and final (final) steps. A matrix-matrix multiplication with a product that is known to be symmetric is counted as 0.5 multiplications. The numbers within parentheses for the LNV minimization methods represent multiplications used for analytical line searches.

require prior knowledge of the chemical potential. Several variants of density matrix purification have been suggested [16–19, 21, 28]. In our benchmark calculations we will consider the canonical purification method and the trace correcting variant by Niklasson [16] and their non-orthogonal counterparts denoted as follows:

- *Trace correcting*: Second order trace-correcting purification by Niklasson employing the polynomials x^2 and $2x - x^2$ [16].
- *Trace correcting (non-orth)*: The non-orthogonal variant [38] of Niklasson’s second order trace correcting scheme using the initial transformation given by equation (31) of Ref. [38].
- *Canonical*: Canonical purification by Palser and Manolopoulos [15].
- *Canonical (non-orth)*: The non-orthogonal variant of canonical purification [15].

For both variants of trace correcting purification, the purifying polynomial is in each iteration chosen such as to minimize the occupation error in the updated density matrix approximation [39].

All purification methods start from a scaled and shifted effective Hamiltonian matrix. Low order polynomials are then recursively applied to construct an expansion of the step function in (1). In exact arithmetics, the occupied subspace of the effective Hamiltonian matrix is preserved throughout this procedure, since the polynomials only affect the eigenvalues. However, in order to obtain a linear scaling method, matrix elements have to be neglected in each iteration of this procedure. This filtering of matrix elements affects the accuracy of the occupied subspace and, in contrast to the LNV methods, any accuracy that is lost cannot be recovered.

C. Truncation

All linear scaling density matrix methods rely on matrix sparsity, which means that some matrix elements must be neglected. In the following we refer to the procedure of removing or neglecting matrix elements as *truncation*.

Two truncation approaches have dominated the literature. In the first approach, matrix elements with magnitude smaller than a predefined threshold are neglected [3, 17, 36, 40]. In the second approach one neglects matrix elements corresponding to distances between basis function centers larger than a predefined cutoff radius [1, 7, 15, 41]. There are also methods to select threshold values for truncation automatically to guarantee a certain accuracy in terms of the norm of the error matrix [42, 43]. This is particularly useful for purification methods where it allows for controlling the forward error of the whole purification process [28].

In this paper, however, we wish to focus on the performance of LNV minimization and purification methods without the additional complications that arise due to difficulties in selecting threshold values or cutoff radii. Therefore, we scan a set of parameters to make sure that we have near optimal threshold values or cutoff radii for each specific case. Although this is not realistic for practical large calculations, this procedure gives us information about the performance that would be possible if the parameters were chosen in a near optimal manner. In Section IV we present several example calculations using both the element magnitude and cutoff radius based truncation schemes. The primary purpose of including examples using both these approaches is not to compare them. Such comparison should take difficulties with selection of threshold values into account. Instead, the purpose is to show that the conclusions of this paper are largely independent of the used truncation approach.

D. Stopping criteria

In LNV minimization methods, both the magnitude of the gradient [3] and the energy difference between two consecutive iterations [27] have been used to decide when to stop the conjugate gradient iterations. We are not aware of any way to relate such stopping criteria to the error in the density matrix so that (2) can be guaranteed. In purification methods, it is possible to directly relate the stopping criterion to the requested accuracy [28]. This is done by stopping the iterations as soon as the eigenvalues are sufficiently close to 0 and 1 so that the error is dominated by the error in the occupied subspace. Combined with rigorous norm-based truncation schemes [43, 44], this allows for controlling errors in such way that (2) is fulfilled. In this paper, however, we avoid the issue of selecting stopping criteria; instead we directly compare to a precomputed highly accurate reference solution and stop the LNV minimization or purification iterations as soon as a requested accuracy as measured by (2) is reached. In this way, we obtain information about the performance that would be possible if the stopping criteria were chosen in a near optimal manner.

E. Utilizing matrix symmetry

In both LNV minimization and purification methods, the performance can be significantly improved by utilizing that many of the involved matrices are symmetric. For example, a common operation is the computation of A^2 for a symmetric matrix A . The result matrix is then known to be symmetric which makes it possible to reduce the computational cost by nearly a factor of 2, see for example Ref. [31]. Also, when expressions like $BA + AB$ are to be computed for symmetric matrices A and B , one can utilize that $AB = (BA)^T$, saving again a factor of 2. Table I shows the number of matrix-matrix multiplications needed per iteration for each of the LNV minimization and purification methods. The number of multiplications needed in preparatory and final phases is also shown. In all reports of matrix-matrix multiplication counts presented in this article, a multiplication with product that is known to be symmetric is counted as 0.5 multiplications.

IV. EXAMPLES

In this section, we compare the performance of LNV minimization and purification methods. The effective Hamiltonian and overlap matrices used in our tests come from a spin-restricted Hartree-Fock calculation for a linear alkane molecule $C_{50}H_{102}$ with 402 electrons occupying 201 orbitals. The standard Gaussian minimal basis set labeled STO-3G is used, giving 352 basis functions. The length of the studied alkane chain is 119 a.u. and its electronic structure is well localized so that accurate sparse matrix approximations can be employed using either element magnitude or cutoff radius based truncation schemes. The eigenvalue spectrum spans from -11.03 to 0.84 a.u. with a band gap of 0.87 a.u. located symmetrically around 0.13 a.u.

Both the orthogonal and non-orthogonal variants of each method are investigated. For all methods, the element magnitude approach for truncation of small matrix elements is tested. For the non-orthogonal variants the cutoff radius approach is tested as well. We recall from Section III that methods operating in orthogonal basis need an inverse factor Z while methods operating in non-orthogonal basis use the inverse overlap matrix S^{-1} . In our examples, the cost of multiplications with Z or S^{-1} is taken into account. However, the cost of computing Z or S^{-1} is not included. Often, the density matrix problem is to be solved several times without changing the basis set so that a precomputed Z or S^{-1} can be reused.

It is common practice to measure performance in density matrix methods by the number of matrix-matrix multiplications needed to obtain the result. Here, we also want to take into account the fact that the cost of each sparse matrix-matrix multiplication depends on the degree of sparsity at each step of the calculation. Therefore, our tests are performed using an implementation of sparse matrix-matrix multiplication in which every single floating point multiplication is counted. This makes it possible to measure the computational cost of each method in a computer platform independent way.

For each method, performance is measured by the total number of scalar floating point multiplications needed to achieve an accuracy of $\|X - D\|_2 \leq 10^{-2}$. Parameters for truncation were selected by carrying out a search for the threshold values or cutoff radii that give the smallest number of floating point operations while achieving the requested accuracy. When the element magnitude approach is used, we use a single neglect threshold τ for all involved matrices. When the cutoff radius approach is used, we employ one radius R_S for the overlap matrix S , one radius R_F for the effective Hamiltonian matrix F and one radius R_X for the current density matrix approximation X . The radius R_X is also used for the inverse of the overlap matrix S^{-1} . In all methods we only perform truncation between iterations and not between intermediate steps of one iteration. For example, when canonical purification is employed, the matrix X_{\perp}^2 is not truncated before computation of X_{\perp}^3 from X_{\perp} and X_{\perp}^2 . Symmetry of matrices is utilized whenever possible as described in Section III E. When the cutoff radius approach is used, the nonzero structures of certain product matrices are known. Such information is used to reduce the number of operations in matrix-matrix multiplications.

The performance using the element magnitude truncation approach is illustrated in Table II, while Table III shows the performance when using the cutoff radius approach. In both cases, the trace correcting purification method is superior both with respect to the number of matrix-matrix multiplications and the number of floating point multiplications. Trace correcting purification in orthogonal basis is the most efficient method, outperforming all LNV methods by more than an order of magnitude.

The different convergence behaviors of the LNV minimization and purification methods are illustrated in Figure 1. Note that as the requested accuracy is increased, the difference between

	τ	iters	mmul	ops
<i>LNV minimization methods</i>				
Standard	1.6e-4	11	60.5	109.4
Standard (non-orth)	3.3e-4	12	107.0	392.3
Simplified	4.1e-5	10	24.0	77.2
Simplified (non-orth)	4.1e-5	10	61.5	369.4
<i>Purification methods</i>				
Trace-correcting	1.6e-4	12	9.0	7.4
Trace-correcting (non-orth)	1.6e-4	12	19.5	28.6
Canonical	8.2e-5	11	14.0	16.8
Canonical (non-orth)	1.6e-4	10	21.5	45.1

TABLE II: Computational effort needed for different methods to reach an accuracy of $\|X - D\|_2 \leq 1e-2$. Truncation is performed using the element magnitude based approach. The table shows the threshold value (τ) used for truncation, the number of iterations (iters), matrix-matrix multiplications (mmul), and millions of floating point multiplications (ops). A matrix-matrix multiplication with symmetric product is counted as 0.5 multiplications.

	R_S	R_F	R_X	iters	mmul	ops
<i>LNV minimization methods</i>						
Standard (non-orth)	9.0	8.0	17.0	9	81.5	316.8
Simplified (non-orth)	6.0	10.0	15.0	9	55.5	179.9
<i>Purification methods</i>						
Trace-correcting (non-orth)	6.0	8.0	15.0	11	18.0	26.0
Canonical (non-orth)	9.0	10.0	15.0	11	23.5	58.9

TABLE III: Computational effort needed for different methods to reach an accuracy of $\|X - D\|_2 \leq 1e-2$. Truncation is performed using the cutoff radius approach. The table shows the cutoff radii (R_S , R_F , and R_X in a.u.) used for truncation, the number of iterations (iters), matrix-matrix multiplications (mmul), and millions of floating point multiplications (ops). A matrix-matrix multiplication with symmetric product is counted as 0.5 multiplications.

purification and LNV methods becomes more and more pronounced, due to the more rapid convergence of the purification methods. The convergence behavior for non-orthogonal trace-correcting purification seen in Figure 1 is in line with previous results for orthogonal trace-correcting and canonical purification [39]. The temporary error increase in the purification simulation labeled A after around 25 multiplications is due to repeated application of the polynomial $2x - x^2$ causing eigenvalues to drift away from 0. Such behavior can in principle result in problems with stability for purification methods that use matrix traces to choose polynomials. This can be remedied by using HOMO/LUMO-based choices of polynomials as discussed in Ref. [28].

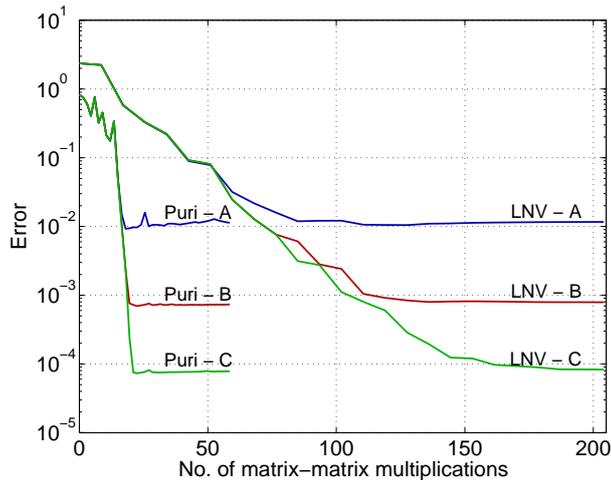


FIG. 1: Convergence of non-orthogonal trace correcting purification (labeled Puri) and non-orthogonal standard LNV minimization (labeled LNV) for different sets of cutoff radii. The radius sets $\{R_S, R_F, R_X\}$ in the simulations labeled A, B, and C were chosen as $\{8, 8, 15\}$, $\{10, 13, 27\}$, and $\{11, 21, 33\}$ a.u. respectively. The error is measured by $\|X - D\|_2$ where X is the current density matrix approximation in the iterative process.

A. Comparison when starting guess is available

Sometimes a density matrix close to the desired result is available, for example from a previous self-consistent field iteration. An advantage of the LNV minimization methods is that such a density matrix can then be used as a starting guess. Then, the number of LNV iterations needed to reach the desired result can be significantly reduced.

The simplified LNV method, which is the most efficient LNV variant according to Tables II and III, requires that the starting guess commutes with F [4]. Unfortunately, available guesses typically do not commute with F .

To study the performance of the LNV method when starting guesses of varying quality are available, we consider the density matrix construction at different stages of the Hartree-Fock self-consistent field calculation for the alkane test system. In each density matrix construction procedure, we use the density matrix from the previous self-consistent field iteration as starting guess. For comparison, the orthogonal variant of the trace-correcting purification method is also employed for the same set of Fock matrices.

Table IV shows the work needed to reach an accuracy of $\|X - D\|_2 \leq 1e-4$. At the last self-consistent field iteration, the starting guess density matrix is close to the exact result with $\|D_{\text{guess}} - D\|_2 = 2.2e-4$. Even under such favorable circumstances, the standard LNV method is more than an order of magnitude more expensive than the trace-correcting purification method in orthogonal basis. Note that the purification method does not use any starting guess. Therefore, the work needed for purification is essentially the same in each line in Table IV, with differences only due to variations in the Fock matrix.

$\ \Delta D\ _2$	LNV (non-orth)			Puri		
	iters	mmul	ops	iters	mmul	ops
3.1e-1	(failed to converge)			15	10.5	45.3
5.1e-2	20	175.0	1953.8	14	10.0	41.8
6.9e-3	12	107.0	1191.5	14	10.0	41.4
9.1e-4	8	73.0	810.3	14	10.0	41.5
2.2e-4	5	47.5	524.4	14	10.0	41.5

TABLE IV: Computational effort needed by the standard LNV method in non-orthogonal basis to reach an accuracy of $\|X - D\|_2 \leq 1e-4$ when starting guesses of varying quality are employed. Truncation is performed using the cutoff radius set $\{R_S, R_F, R_X\} = \{11, 21, 33\}$. Corresponding results for the orthogonal variant of trace-correcting purification are shown for reference. In the purification case, truncation is performed using the element magnitude based approach with a matrix element neglect threshold of $\tau = 6.4e-7$. The columns show the deviation $\|D_{\text{guess}} - D\|_2$ of the starting guess density from the exact result ($\|\Delta D\|_2$), the number of iterations (iters), matrix-matrix multiplications (mmul), and millions of floating point multiplications (ops). A matrix-matrix multiplication with symmetric product is counted as 0.5 multiplications.

B. Dependence on eigenvalue distribution

In this section we will investigate how the number of matrix-matrix multiplications depends on the eigenvalue distribution of the effective Hamiltonian matrix. Considering that the condition number of the density matrix construction problem is proportional to $1/\xi$, all methods are expected to require more work as ξ is decreased. Previous work has shown that purification methods require $\mathcal{O}(\ln(1/\xi))$ iterations to converge [16]. We have not come across any corresponding assessment of the band gap dependence of the LNV methods in the literature, although examples where a small band gap leads to large numbers of LNV iterations have been reported [27].

To check the band gap dependence, we apply the LNV and purification methods to large diagonal Hamiltonian test matrices with band gaps of varying size located at the center of the eigen-spectrum. The results are shown in Figure 2. For purification methods, our results confirm the $\mathcal{O}(\ln(1/\xi))$ dependence reported previously. For the LNV method an $\mathcal{O}(\sqrt{1/\xi})$ dependence on the gap is observed, indicating extremely slow convergence for small gaps compared to purification. It should be noted that the test Hamiltonian matrices had all eigenvalues in the $[0, 1]$ interval. Generally, the spectral width $\Delta\epsilon$ of the Hamiltonian is not equal to 1. Then, the number of matrix multiplications scale as $\mathcal{O}(\ln(\Delta\epsilon/\xi))$ and $\mathcal{O}(\sqrt{\Delta\epsilon/\xi})$ for the purification and LNV methods, respectively.

Apart from the dependence on the band gap, another important aspect is how sensitive the method is to variations in occupancy and the location of the chemical potential. To investigate this, we apply the different methods to large diagonal Hamiltonian test matrices with fixed gap but varying occupancy and location of the chemical potential. Results are shown in Figure 3. For purification methods, Figure 3 confirms previous results on canonical and trace-correcting purification [15, 16]. For the LNV method we observe that the most rapid convergence is achieved for $\mu = \frac{n_{\text{occ}}}{n} = 0.5$.

Since the simulations shown in Figures 2 and 3 are performed for diagonal test matrices and assuming an orthogonal basis set, all four LNV methods in Section III A behave identically, the only difference being the number of matrix-matrix multiplications per iteration. We choose to

demonstrate the multiplication counts for the cheapest method of the four, namely the simplified LNV method in orthogonal basis. For purification, we demonstrate the orthogonal variants of trace-correcting and canonical purification.

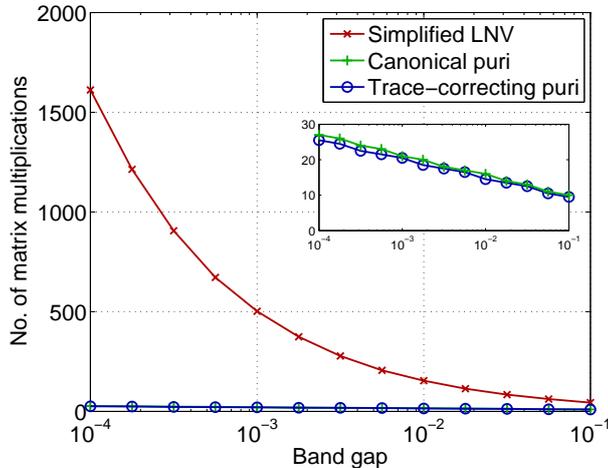


FIG. 2: Number of matrix-matrix multiplications for varying band gap. The test Hamiltonian is a diagonal 100000×100000 matrix with eigenvalues uniformly distributed between 0 and 1 with exception for the band gap which is placed symmetrically around 0.5. The requested accuracy is set to $\|X - D\|_2 \leq 1e-4$.

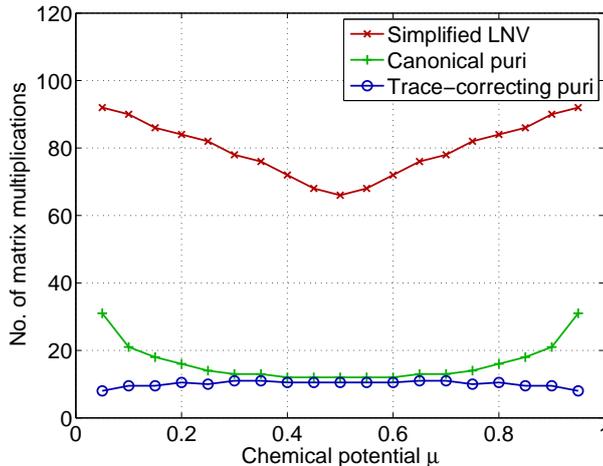


FIG. 3: Number of matrix-matrix multiplications for varying occupancy n_{occ}/n and location of the chemical potential μ . Here, n_{occ} is the number of occupied orbitals and $n = 100000$ is the matrix dimension. The test Hamiltonian is a diagonal matrix with eigenvalues uniformly distributed between 0 and 1 with exception for the band gap which is of size 0.05 and located symmetrically around μ . The chemical potential and occupancy are varied together so that $\mu = n_{\text{occ}}/n$. The requested accuracy is set to $\|X - D\|_2 \leq 1e-4$.

V. COMMENTS AND DISCUSSION

The results of the comparison between LNV and purification methods in the previous section are clearly in favor of the purification approach, which in the test calculations reached the

desired accuracy with less than 10% of the computational cost needed for LNV methods. The rapid convergence of the purification methods make them even more superior when high accuracy is desired. Furthermore, purification methods are generally more economic than LNV methods when it comes to memory usage. This is because purification methods use fewer intermediate matrices, and to a larger extent work with symmetric matrices. In particular, the second order trace-correcting purification scheme in orthogonal basis requires only two symmetric matrices in memory.

Several improvements and extensions of the canonical and trace-correcting purification methods have been suggested. Examples include HOMO-LUMO based purification allowing for rigorous control of the forward error and a natural stopping criterion [28], purification for spin-unrestricted calculations [21], efficient calculation of band edge states [45], and application to perturbation theory [46].

In principle, LNV methods have an advantage in that the number of iterations can be reduced when a good starting guess is available. One might think that, in such cases, LNV methods would be more efficient than purification in spite of their slower convergence. However, given the high computational cost of each LNV iteration, the LNV method would have to converge in very few iterations to be competitive with purification. In our test calculations presented in Table IV, the LNV method required 5 iterations to improve the accuracy by a factor of 2.2. These 5 iterations required more computational work than a whole corresponding purification process. In practice, an additional difficulty when using minimization methods is to relate the stopping criterion to the desired accuracy. An inadequate stopping criterion could lead to superfluous iterations or errors larger than the desired accuracy.

A. Hybrid schemes

Several hybrid schemes have been proposed, in which the LNV and purification approaches are combined in various ways [4, 10, 27]. Here, we will consider the hybrid method by Bowler and Gillan [27] which is used in the Conquest code [41, 47]. In that hybrid method, canonical purification is used to generate a starting guess for the LNV method, which reduces the number of LNV iterations needed to reach a heuristically set stopping criterion in terms of the fractional change in energy between LNV iterations.

Before further considering this hybrid method, we want to point out that fundamentally different perspectives can be taken with respect to the “error” of an LNV calculation. One way to measure the error is to consider the difference to the exact solution D as we have done in this paper. Another possibility is to search for the matrix L that minimizes the functional in (3) for a given set of cutoff radii. This is a well defined matrix, but it is different from D . Thus, the forward error of a LNV method can be defined either as $\|X - D\|$ or as $\|X - L\|$. Lowering the error with respect to L does not necessarily reduce the error with respect to D .

We have implemented the hybrid method by Bowler and Gillan and applied it to our alkane test system. In each iteration, we have computed both the error with respect to D and the difference with respect to L for the used set of cutoff radii. The results, shown in Figure 4, show that the error is reduced rapidly during the first 12 purification iterations, after which it stagnates. During the LNV iterations, the result slowly approaches L , but the error with respect to D is not improved. Thus, when the error is measured as in (2), there is apparently no point in adding LNV iterations after the purification phase. However, if for some reason L is considered to be the desired result it could make sense to perform LNV iterations after purification as in the hybrid method.

Other hybrid schemes employ purification between the conjugate gradient steps to move to-

wards idempotency and avoid runaway solutions [10, 25]. This approach is not tested in the present work.

B. Three-dimensional molecular systems

The test calculations in this work were performed for a linear alkane system small enough to allow comparisons to nearly exact results as well as using a sparse matrix multiplication implementation that counts operations. When density matrix methods are applied to three-dimensional systems, the number of nonzero matrix elements grows much faster as a result of matrix operations. This effect is clearly illustrated in Ref. [48] where filling fractions for the matrices X , XS , and XSX are compared for a three-dimensional system. In that case, the matrix XSX contains more than 13 times as many elements as X . In our own calculations applying purification in orthogonal basis for three-dimensional water clusters [49, 50], we have observed that the matrix X_{\perp}^2 typically contains more than 5 times as many elements as X_{\perp} . Of course, if no truncation is applied between matrix operations, the number of elements will grow more for each matrix multiplication. This effect can be seen in the test calculations in this work by comparing the number of operations per matrix-matrix multiplication for the various methods; each multiplication is cheapest for the trace-correcting purification in orthogonal basis since that method only needs to compute X_{\perp}^2 from X_{\perp} , no additional multiplications involving X_{\perp}^2 are needed. For three-dimensional systems this effect will be much more pronounced. This is a disadvantage for the LNV methods that need to evaluate long products such as $XSX S$. Thus, for three-dimensional systems the difference in performance between purification and LNV methods is expected to be even larger than indicated in our test calculations.

C. Orthogonal or non-orthogonal formulation

Regarding the choice between orthogonal and non-orthogonal formulations of each method, our test calculations indicate that the orthogonal formulation is in general more efficient, provided that an inverse factor Z of the overlap matrix can be readily obtained. However, it should be noted that the choice between orthogonal and non-orthogonal formulation can be considered a separate issue from the choice between purification and minimization approaches. Both purification and LNV minimization methods are available in both orthogonal and non-orthogonal formulations. In cases where an inverse factor Z is not available, trace-correcting purification in non-orthogonal basis should be considered as an alternative to LNV minimization.

VI. CONCLUSION

The comparison of purification and LNV minimization methods for constructing the density matrix presented in this paper indicates that the purification approach is significantly more efficient. Considering that the difference in performance is expected to be even larger for three-dimensional systems, it is likely that the application of purification can improve the performance with more than an order of magnitude compared to the LNV method. Therefore, we recommend users of the LNV method to consider modern purification algorithms as a way to significantly

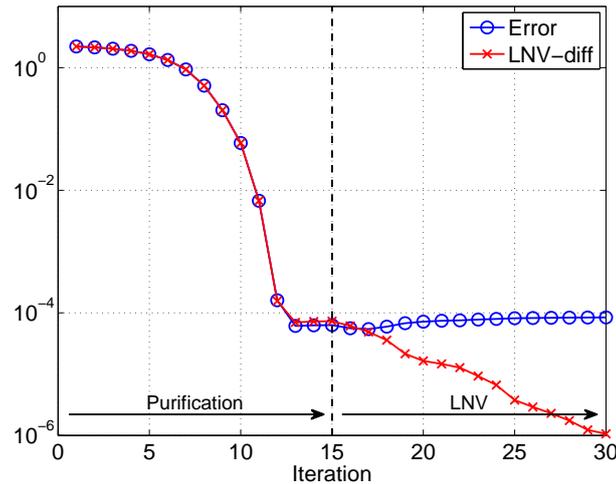


FIG. 4: Convergence of a non-orthogonal hybrid purification – LNV minimization method. The method starts with canonical purification and switches to LNV minimization on the first increase of $\text{Tr}[XF]$ from one purification iteration to the next. The radii $\{R_S, R_F, R_X\}$ were chosen as $\{11, 21, 33\}$ a.u.. The figure shows the error as measured by $\|X - D\|_2$ (Error) and the difference $\|X - L\|_2$ to the exact LNV solution L for this particular set of radii (LNV-diff).

improve performance.

-
- [1] X.-P. Li, R. W. Nunes, and D. Vanderbilt, *Phys. Rev. B* **47**, 10891 (1993).
 - [2] M. S. Daw, *Phys. Rev. B* **47**, 10895 (1993).
 - [3] J. M. Millam and G. E. Scuseria, *J. Chem. Phys.* **106**, 5569 (1997).
 - [4] M. Challacombe, *J. Chem. Phys.* **110**, 2332 (1999).
 - [5] T. Helgaker, H. Larsen, J. Olsen, and P. Jørgensen, *Chem. Phys. Lett.* **327**, 397 (2000).
 - [6] H. Larsen, J. Olsen, P. Jørgensen, and T. Helgaker, *J. Chem. Phys.* **115**, 9685 (2001).
 - [7] Y. Shao, C. Saravanan, M. Head-Gordon, and C. A. White, *J. Chem. Phys.* **118**, 6144 (2003).
 - [8] W. Kohn, *Phys. Rev. Lett.* **76**, 3168 (1996).
 - [9] P. D. Haynes and M. C. Payne, *Phys. Rev. B* **59**, 12173 (1999).
 - [10] P. Haynes, C.-K. Skylaris, A. Mostofi, and M. Payne, *J. Phys.: Condens. Matter* **20**, 294207 (2008).
 - [11] S. Goedecker and L. Colombo, *Phys. Rev. Lett.* **73**, 122 (1994).
 - [12] R. Baer and M. Head-Gordon, *J. Chem. Phys.* **107**, 10003 (1997).
 - [13] W. Liang, C. Saravanan, Y. Shao, R. Baer, A. T. Bell, and M. Head-Gordon, *J. Chem. Phys.* **119**, 4117 (2003).
 - [14] S. Goedecker and M. Teter, *Phys. Rev. B* **51**, 9455 (1995).
 - [15] A. H. R. Palser and D. E. Manolopoulos, *Phys. Rev. B* **58**, 12704 (1998).
 - [16] A. M. N. Niklasson, *Phys. Rev. B* **66**, 155115 (2002).
 - [17] A. M. N. Niklasson, C. J. Tymczak, and M. Challacombe, *J. Chem. Phys.* **118**, 8611 (2003).
 - [18] A. Holas, *Chem. Phys. Lett.* **340**, 552 (2001).
 - [19] D. A. Mazziotti, *Phys. Rev. E* **68**, 066701 (2003).
 - [20] A. M. N. Niklasson, C. J. Tymczak, and H. Røder, *Phys. Rev. B* **66**, 155120 (2002).

- [21] H. J. Xiang, W. Z. Liang, J. Yang, J. G. Hou, and Q. Zhu, *J. Chem. Phys.* **123**, 124105 (2005).
- [22] K. Németh and G. E. Scuseria, *J. Chem. Phys.* **113**, 6035 (2000).
- [23] G. Beylkin, N. Coult, and M. J. Mohlenkamp, *J. Comput. Phys.* **152**, 32 (1999).
- [24] E. H. Rubensson, Ph.D. thesis, Department of Theoretical Chemistry, Royal Institute of Technology, Stockholm (2008).
- [25] D. K. Jordan and D. A. Mazziotti, *J. Chem. Phys.* **122**, 084114 (2005).
- [26] A. D. Daniels and G. E. Scuseria, *J. Chem. Phys.* **110**, 1321 (1999).
- [27] D. Bowler and M. Gillan, *Comp. Phys. Commun.* **120**, 95 (1999).
- [28] E. H. Rubensson, E. Rudberg, and P. Sałek, *J. Chem. Phys.* **128**, 074106 (2008).
- [29] E. H. Rubensson, E. Rudberg, and P. Sałek, *J. Math. Phys.* **49**, 032103 (2008).
- [30] M. Benzi, C. D. Meyer, and M. Tuma, *SIAM J. Sci. Comput.* **17**, 1135 (1996).
- [31] E. H. Rubensson, E. Rudberg, and P. Sałek, *J. Comput. Chem.* **28**, 2531 (2007).
- [32] E. H. Rubensson, N. Bock, E. Holmström, and A. M. N. Niklasson, *J. Chem. Phys.* **128**, 104105 (2008).
- [33] R. W. Nunes and D. Vanderbilt, *Phys. Rev. B* **50**, 17611 (1994).
- [34] A. D. Daniels, J. M. Millam, and G. E. Scuseria, *J. Chem. Phys.* **107**, 425 (1997).
- [35] C. A. White, P. Maslen, M. S. Lee, and M. Head-Gordon, *Chem. Phys. Lett.* **276**, 133 (1997), ISSN 0009-2614.
- [36] A. Dyan, P. Dubot, and P. Cenedese, *Phys. Rev. B* **72**, 125104 (pages 6) (2005).
- [37] R. McWeeny, *Proc. R. Soc. London Ser. A* **235**, 496 (1956).
- [38] A. M. N. Niklasson, V. Weber, and M. Challacombe, *J. Chem. Phys.* **123**, 044107 (2005).
- [39] A. M. N. Niklasson and V. Weber, *J. Chem. Phys.* **127**, 064105 (2007).
- [40] P. E. Maslen, C. Ochsenfeld, C. A. White, M. S. Lee, and M. Head-Gordon, *J. Phys. Chem. A* **102**, 2215 (1998).
- [41] D. R. Bowler and T. Miyazaki, *J. Phys.: Condens. Matter* **22**, 074207 (2010).
- [42] E. H. Rubensson and P. Sałek, *J. Comput. Chem.* **26**, 1628 (2005).
- [43] E. H. Rubensson, E. Rudberg, and P. Sałek, *J. Comput. Chem.* **30**, 974 (2009).
- [44] E. H. Rubensson and E. Rudberg, to be published (2010).
- [45] E. H. Rubensson and S. Zahedi, *J. Chem. Phys.* **128**, 176101 (2008).
- [46] A. M. N. Niklasson and M. Challacombe, *Phys. Rev. Lett.* **92**, 193001 (2004).
- [47] D. R. Bowler, R. Choudhury, M. J. Gillan, and T. Miyazaki, *phys. stat. sol. (b)* **243**, 989 (2006).
- [48] N. Hine, P. Haynes, A. Mostofi, C.-K. Skylaris, and M. Payne, *Comp. Phys. Commun.* **180**, 1041 (2009).
- [49] E. Rudberg, E. H. Rubensson, and P. Sałek, *J. Chem. Phys.* **128**, 184106 (2008).
- [50] E. Rudberg, E. H. Rubensson, and P. Sałek, to be published (2010).