<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper published in *SIAM Journal on Scientific Computing*. This paper has been peer-reviewed but does not include the final publisher proof-corrections or journal pagination.

Citation for the original published paper (version of record):

Rubensson, E H., Niklasson, A M. (2014)

Interior eigenvalues from density matrix expansions in quantum mechanical molecular dynamics

SIAM Journal on Scientific Computing, 36: B147-B170

<https://doi.org/10.1137/130911585>

Access to the published version may require subscription.

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-220198>

INTERIOR EIGENVALUES FROM DENSITY MATRIX EXPANSIONS IN QUANTUM MECHANICAL MOLECULAR DYNAMICS*

EMANUEL H. RUBENSSON[†] AND ANDERS M. N. NIKLASSON[‡]

Abstract. An accelerated polynomial expansion scheme to construct the density matrix in quantum mechanical molecular dynamics simulations is proposed. The scheme is based on recursive density matrix expansions, e.g. [Phys. Rev. B. 66 (2002), p. 155115], which are accelerated by a scale-and-fold technique [J. Chem. Theory Comput. 7 (2011), p. 1233]. The acceleration scheme requires interior eigenvalue estimates, which may be expensive and cumbersome to come by. Here we show how such eigenvalue estimates can be extracted from the recursive expansion by a simple and robust procedure at a negligible computational cost. Our method is illustrated with density functional tight-binding Born–Oppenheimer molecular dynamics simulations, where the computational effort is dominated by the density matrix construction. In our analysis we identify two different phases of the recursive polynomial expansion, the conditioning and purification phases, and we show that the acceleration represents an improvement of the conditioning phase, which typically gives a significant reduction of the computational cost.

Key words. Born–Oppenheimer molecular dynamics, density functional theory, density matrix purification, Fermi operator expansion, homo-lumo gap, interior eigenvalues, linear scaling electronic structure theory, matrix functions, non-monotonic expansion, recursive expansion, scale-and-fold acceleration, trace-correcting purification

AMS subject classifications. 65F15, 65F60, 70E55, 81-08

1. Introduction. With the fast growth of computational processing power, atomistic simulations based on quantum mechanical calculations of the electronic structure have become a powerful approach to the study of a broad range of problems in materials science, chemistry, and biology [26, 33, 64]. Nevertheless, the computational cost associated with electronic structure calculations normally limits applications to fairly small systems. In particular, the cubic, $O(N^3)$, scaling of the computational cost as a function of the number of atoms, N , for the regular solution of the quantum mechanical eigenvalue problem is considered to be a most limiting factor. A number of different electronic structure technologies have therefore been developed that circumvent this bottleneck with a computational effort that scales only linearly with the system size [6, 18]. The reduction in the cost is typically achieved by utilizing sparse matrix algebra in an iterative construction of the density matrix, which avoids the full regular solution of the quantum mechanical eigenvalue problem. The matrix sparsity arise in localized atomic basis set representations due to the short-range character of the electronic wavefunctions for non-metallic materials [3, 27, 28]. With linear scaling techniques it is today possible to perform electronic structure calculations involving millions of atoms [5, 65].

An ideal linear scaling method should combine several important properties: *i)* it should be well tailored for high-performance on heterogeneous multi-core architec-

*Support from the Göran Gustafsson foundation, the Swedish research council (grant no. 621-2012-3861), the Lisa and Carl–Gustav Esseen foundation, the Swedish national strategic e-science research program (eSSENCE), United States Department of Energy Office of Basic Energy Sciences, and the International Ten Bar Java group of T. Peery, as well as discussions with M. Cawkwell are gratefully acknowledged.

[†]Division of Scientific Computing, Department of Information Technology, Uppsala University, Box 337, SE-751 05 Uppsala, Sweden (emanuel.rubensson@it.uu.se).

[‡]Theoretical division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA (amn@lanl.gov).

tures, *ii*) require little intermediate memory storage, *iii*) allow well controlled numerical accuracy, and *iv*) have a low computational prefactor. Among several linear scaling techniques [2, 4, 10, 11, 22, 24, 25, 31, 34, 45], the recursive second-order spectral projection method (SP2) [37], which sometimes also is referred to as the trace-correcting purification algorithm, represents a surprisingly simple and efficient choice [57]. The SP2 algorithm allows for control of errors arising from an approximate sparse matrix algebra [41, 54, 55] and it has a low computational prefactor. Even in the dense (non-sparse) limit the SP2 algorithm outperforms regular LAPACK diagonalization techniques on multi-core platforms, both with respect to speed and accuracy [9]. The computational cost of the SP2 algorithm is dominated by a sequence of matrix-matrix multiplications; typically between 20 to 40 multiplications are needed. Highly optimized linear algebra subroutines can perform dense matrix-matrix multiplications with close to peak performance on multi-core platforms. Thus, even without utilizing linear scaling sparse matrix algebra, the SP2 algorithm has excellent performance and is well suited for modern computational architectures [9].

Recently Rubensson proposed an acceleration technique that can further boost the performance of the SP2 algorithm, as well as other recursive polynomial expansion schemes [52]. The acceleration is based on the idea to allow for nonmonotonicity in the interval of interest, which gives additional flexibility in the choice of recursive expansion polynomials. By the use of a scale and fold technique, i.e. by shifting and re-scaling the SP2 projection polynomials, the number of matrix-matrix multiplications required by the SP2 algorithm is significantly reduced. This boosting technique requires estimates of interior eigenvalues that may be expensive and cumbersome to come by. In this article we show how such eigenvalue estimates can be extracted by a simple procedure from the recursive polynomial expansion at negligible computational cost.

Our scheme is particularly useful in molecular dynamics simulations. A molecular dynamics simulation often involves hundreds of thousands of time steps and therefore requires a low computational cost for the calculation of each new configuration. This is provided by the accelerated SP2 algorithm and we show how eigenvalue estimates that are necessary for the acceleration can be extracted from previous time steps at practically no extra cost. The scheme is illustrated with density functional tight-binding molecular dynamics simulations, where the computational effort is dominated by the construction of the electronic density matrix through the recursive polynomial expansion. All the simulations are performed using dense matrix algebra in the computational framework of extended Lagrangian Born–Oppenheimer molecular dynamics [38, 61], which provides accurate and stable molecular trajectories with long-term energy conservation.

The article is outlined as follows: first we give some background to the novel developments in the present work. We describe the problem and give a brief overview of methods for computation of the density matrix, focusing on recursive polynomial expansions and the acceleration technique. Thereafter, we show how the required interior eigenvalues can be estimated from the recursion. We discuss considerations underlying the new algorithm and its expected behavior for large systems. The article goes on with the application of the new method in quantum mechanical molecular dynamics. The extended Lagrangian formulation of Born–Oppenheimer molecular dynamics is presented and we illustrate the performance of our eigenvalue estimates and the accelerated recursion in this context. We end the article with a discussion and some concluding remarks.

2. Background.

2.1. The density matrix. The single-particle density matrix at zero electronic temperature can be defined as a matrix function of the effective Hamiltonian matrix,

$$(2.1) \quad D = \theta(\mu I - H),$$

where $\theta(x)$ is the Heaviside step function and μ is the chemical potential or Fermi level. We restrict our applications to non-metallic materials and assume that H does not have degenerate eigenvalues at μ . We further assume that all matrices are Hermitian. We will refer to the smallest eigenvalue above μ as the lowest unoccupied molecular orbital (lumo) eigenvalue and to the largest eigenvalue below μ as the highest occupied molecular orbital (homo) eigenvalue. The gap between the homo and lumo eigenvalues is referred to as the homo-lumo gap. See Figure 2.1 for an illustration.

The density matrix is the matrix for orthogonal projection onto the occupied subspace spanned by the eigenvectors of H that correspond to eigenvalues smaller than μ . The dimension of the occupied subspace is equal to the number of occupied electron orbitals, N_{occ} . Usually, the number of electrons is given as input to the program which has to automatically adjust the chemical potential so that a correct occupation number is achieved.

In general, the method of choice for computation of matrix functions of symmetric matrices is diagonalization [20, p. 84]. In short, an eigendecomposition $H = V\Lambda V^T$ is computed (with Λ diagonal) and the matrix function can then be computed as $f(H) = Vf(\Lambda)V^T$. This is the standard method to solve (2.1) in the context of electronic structure calculations.

Since D is a projection matrix, an eigendecomposition, which possesses information about all eigenpairs of H , is not really needed to form D . Any basis of the occupied subspace could be used to construct D . Important for the accuracy however is to resolve the step at the chemical potential. In other words, while rotations of the eigenvectors *within* the occupied (or unoccupied) subspace do not affect the result at all, rotations corresponding to a leakage between the occupied and unoccupied subspaces do affect the accuracy, and are more likely to occur between eigenvectors corresponding to eigenvalues around the chemical potential. One would therefore expect the problem to become more difficult when the homo-lumo gap decreases. This is reflected by the condition number for the problem.

An appropriate condition number for the problem of computing the density matrix can be defined as [53]

$$(2.2) \quad \kappa = \lim_{h \rightarrow 0} \sup_{A: \|A\| = \Delta\epsilon} \frac{\|\theta(\mu I - (H + hA)) - \theta(\mu I - H)\|}{h}.$$

Since $D = \theta(\mu I - H)$ is the matrix for orthogonal projection onto the occupied subspace, the condition number can be analyzed in terms of the canonical angles between the exact and perturbed subspaces. By essentially following the analysis in [55], it can be shown that $\kappa = \Delta\epsilon/\xi$, where $\Delta\epsilon = \lambda_{\max} - \lambda_{\min}$ is the spectral width and ξ is the homo-lumo gap of H . In the assessment of density matrix methods, it is therefore important to consider how the computational cost scales with the homo-lumo gap.

2.1.1. The self-consistent field procedure. In electronic structure methods such as Hartree–Fock [36, 51] or density functional theory [21, 29], the effective single-

FIG. 2.1. Eigenvalues of the effective Hamiltonian matrix H corresponding to occupied (red circles) and unoccupied (blue crosses) orbitals. $[\lambda_{\text{homo}}^{\text{out}}, \lambda_{\text{homo}}^{\text{in}}]$ and $[\lambda_{\text{lumo}}^{\text{in}}, \lambda_{\text{lumo}}^{\text{out}}]$ are intervals containing the homo and lumo eigenvalues, respectively. The chemical potential μ is a number between λ_{homo} and λ_{lumo} . $\lambda_{\min}$ and $\lambda_{\max}$ are either the extremal eigenvalues or lower and upper bounds thereof. We will refer to $\lambda_{\text{homo}}^{\text{in}}$ and $\lambda_{\text{lumo}}^{\text{in}}$ as inner bounds, and $\lambda_{\text{homo}}^{\text{out}}$ and $\lambda_{\text{lumo}}^{\text{out}}$ as outer bounds.

particle Hamiltonian matrix H depends on the electronic configuration, which is determined by the density matrix D . If the density matrix used to construct the Hamiltonian matrix is equal to the density matrix given by (2.1), we say that the solution is self-consistent. Finding the correct self-consistent electronic ground state therefore requires an iterative procedure, with repeated computation of Hamiltonian and density matrices,

$$(2.3) \quad D_0 \rightarrow H_1 \rightarrow D_1 \rightarrow H_2 \rightarrow D_2 \rightarrow \dots$$

In its simplest form, this self-consistent field procedure is a fixed-point iteration, where in each iteration the most recent density or Hamiltonian matrix is used as input to the next step. Usually, however, the Hamiltonian matrix is in each step taken as a linear combination of the most recent and previous Hamiltonian matrices to reduce the number of iterations needed to reach convergence. Popular mixing schemes include simple linear mixing, the direct inversion in the iterative subspace (DIIS) method [49], and Broyden mixing [7, 23].

In Born–Oppenheimer molecular dynamics simulations [33], the self-consistent field procedure is employed in each time step, as the electronic ground state solution is needed in the computation of the interatomic forces. In the context of the present work it does not matter if the sequence of matrices in (2.3) is generated by a self-consistent field optimization or if it is coming from molecular dynamics time stepping. The important point is that there is some procedure generating a sequence of Hamiltonian matrices $H_1, H_2, \dots$ for which the corresponding density matrices need to be computed via (2.1).

2.2. Recursive polynomial expansion of the density matrix. There are several alternatives to the eigenvalue decomposition for the calculation of the matrix step function in (2.1). In the context of linear scaling electronic structure theory some of the first techniques were based on serial Chebyshev expansions [18, 19, 59, 60], where the sparsity of the Hamiltonian can be used efficiently to achieve linear scaling complexity in the calculation of the Chebyshev polynomials through their recurrence relation. The degree of the Chebyshev polynomial required to reach a certain accuracy is proportional to $\kappa = \Delta\epsilon/\xi$ [19]. The computational cost therefore scales as $\mathcal{O}(\kappa)$ or at best $\mathcal{O}(\sqrt{\kappa})$ if the matrix polynomial is explicitly evaluated with the Paterson–Stockmeyer method [32, 48].

A difficulty with the Chebyshev techniques is that the chemical potential in (2.1) needs to be known in advance or has to be adjusted iteratively in order to achieve the correct occupation number. Palser and Manolopoulos proposed a scheme [46] that overcomes this problem through a recursive expansion based on the third-order purification polynomial $3X^2 - 2X^3$ that originally was designed by McWeeny to adjust

the fractional occupation of approximate density matrices to represent pure ensembles [35, 36]. In the Palser–Manolopoulos (PM) scheme (as well as in the original purification method by McWeeny) the density matrix is given through a rapid recursive expansion of the Heaviside step function where

$$(2.4) \quad \theta(\mu I - H) = \lim_{n \rightarrow \infty} f_n(f_{n-1}(\dots f_0(H) \dots)).$$

By adjusting $X_0 = f_0(H)$ such that X_0 has the correct trace, i.e. the desired occupation, N_{occ} , the successive McWeeny polynomials, $f_n(X_{n-1})$, can be modified to preserve the trace. In this trace conserving (or canonical) “purification” scheme the step is automatically formed at the chemical potential and no prior knowledge of the chemical potential is required. By using sparse matrix algebra each polynomial recursion $X_n = f_n(X_{n-1})$ can be calculated with linear scaling effort.

A problem with the Palser–Manolopoulos scheme is that it is slow at low and high occupation, i. e. when N_{occ}/N is close to 0 or 1. A solution to this problem was offered by the recursive SP2 expansion scheme by Niklasson [37]. Instead of using third (or higher) order polynomials that simultaneously project the occupied eigenvalue spectrum of H towards 1 and the unoccupied towards 0, the SP2 scheme uses a combination of two second order spectral projection polynomials,

$$(2.5) \quad f_n(X) = \begin{cases} X^2 \\ 2X - X^2 \end{cases}$$

that either projects the eigenvalues towards 1 for $f_n(X) = 2X - X^2$ or towards 0 for $f_n(X) = X^2$, see Figure 2.2. It is easy to see that any recursive combination of these two polynomials converges to a step function with the step formed in the interval $[0, 1]$. Moreover, since the trace of X^2 is always smaller than the trace of X , given that all eigenvalues of X are in $[0, 1]$, and vice versa for $2X - X^2$, we can choose between the two polynomials to iteratively correct the occupation such that the trace of the converged expansion automatically has the desired value. The occupation correcting SP2 algorithm is described in some detail in Alg. 1. After an initial normalization, where $\lambda_{\text{max/min}}$ are upper and lower bounds of the eigenspectrum of H , all eigenvalues of X are in reverse order in the interval $[0, 1]$. Thereafter we apply the spectral projection polynomials that are chosen to achieve the correct occupation at convergence. Prior knowledge of the chemical potential or the homo-lumo gap is not required.

Linear scaling computational complexity in density matrix methods is usually achieved by neglecting matrix elements deemed not to contribute significantly to the overall accuracy [54]. For recursive polynomial expansions on the form (2.4) the removal of matrix elements can be done in such a way that given a tolerance ε , an accuracy

$$(2.6) \quad \|D - \tilde{D}\| \leq \varepsilon$$

is guaranteed [54, 55]. Here, D and $\tilde{D}$ are the exact and approximate density matrices respectively, and $\|\cdot\|$ is a unitarily invariant norm. The error control is based on an analysis of the perturbation of the occupied subspace due to the neglect of matrix elements. As in the analysis of the condition number, the perturbation can be measured by the canonical angles between the exact and perturbed subspaces. The neglect of small matrix elements is done in such a way that the canonical angles are controlled, which entails control of the error in the density matrix. The error control requires

information about the homo and lumo eigenvalues, or more precisely a lower bound of the homo-lumo gap; compare with the expression for the condition number.

The polynomial expansion order increases very rapidly in the SP2 recursion. After only 30 matrix-matrix multiplications the expansion order is over 1 billion, and the number of required multiplications increases only with the logarithm of the condition number, i. e. as $\mathcal{O}(\log \kappa)$ [37]. Furthermore, the SP2 method performs well wherever the chemical potential may be located. In fact, the SP2 method is at its best at low and high occupation numbers. Also, the SP2 method is very memory-efficient as it requires a total of only two symmetric matrices in memory.

Algorithm 1 Trace-correcting 2nd-order spectral projection expansion (SP2)

```

1:  $X = \frac{\lambda_{\max} I - H}{\lambda_{\max} - \lambda_{\min}}$ 
2: for  $i = 1, 2, \dots, n$  do
3:   if  $|\text{Tr}(X^2) - N_{\text{occ}}| < |\text{Tr}(2X - X^2) - N_{\text{occ}}|$  then
4:      $X = X^2$ 
5:   else
6:      $X = 2X - X^2$ 
7:   end if
8: end for
9: return  $X$ 

```

2.3. Acceleration. The recursive expansion functions in (2.4) are usually taken as polynomials that have fixed points at 0 and 1, vanishing derivatives at 0 and/or 1, and that are monotonically increasing in the $[0, 1]$ interval. Such polynomials are very good for bringing a near-idempotent matrix closer to idempotency. For the SP2 polynomials, for example, the asymptotic convergence rate is quadratic. However, the initial matrix is typically far from being idempotent, and most of the work in the recursive expansion is therefore spent on bringing the matrix near idempotency. For this task polynomials on the form described above are not optimal. By allowing for non-monotonicity in the projection polynomials the convergence can be boosted [52].

The basic idea behind the acceleration technique of the recursive polynomial expansion scheme is that prior knowledge of the eigenvalue distribution allows for a more optimized design of the projection polynomials that accelerates convergence [52]. What is needed, more precisely, are estimates of the homo and lumo eigenvalues λ_{homo} and λ_{lumo} in Figure 2.1. If such estimates are available, the eigenspectrum in each step can be shifted and scaled so that the projection polynomials (X^2 or $2X - X^2$) fold the eigenspectrum, which results in a more rapid convergence. Effectively, the SP2 projection polynomials are shifted and rescaled as illustrated in Figure 2.2. The accelerated trace-correcting expansion is given by Alg. 2.

Too aggressive scaling can lead to the occupied and unoccupied eigenstates being mixed up. Therefore, a lower bound of the homo eigenvalue and an upper bound of the lumo eigenvalue, $\lambda_{\text{homo}}^{\text{out}}$ and $\lambda_{\text{lumo}}^{\text{out}}$ in Figure 2.1, are used in the accelerated algorithm. This makes the algorithm robust with respect to the homo-lumo estimates. If only loose bounds of the homo and lumo eigenvalues are available, the acceleration can still be used. As the algorithm is formulated, loose bounds will never lead to reduced efficiency or accuracy compared to the original SP2 algorithm, but tighter bounds result in a more efficient acceleration.

Although we do not further explore the sparse case in the present work, it is important to note that the schemes for error control of [54, 55] outlined in the previous

FIG. 2.2. The regular SP2 projection polynomials X^2 and $2X - X^2$ (solid lines). The shifted and rescaled polynomials used in the accelerated recursive expansion, $((1-\alpha)I + \alpha X)^2$ and $2\alpha X - (\alpha X)^2$ (dashed lines). Here, the scaling parameter was taken as $\alpha = 1.3$. In practice, this parameter is adjusted during the course of the recursive expansion to give the best possible acceleration, see Alg. 2.

section are applicable also in the accelerated expansion since the algorithm is still in the form (2.4).

The process of bringing the matrix to idempotency can be seen as consisting of two different phases, that we shall refer to as the *conditioning phase* and the *purification phase*. In the conditioning phase the idempotency error $\|X - X^2\|$ is not reduced since in this phase there are still eigenvalues that are not close to 0 or 1 but jump around in the $[0, 1]$ interval. Note that in each recursive expansion step, the problem that remains to be solved can be seen as an independent matrix step function problem with an associated condition number. The recursive expansion procedure may then be seen as a procedure to reduce the condition number to 1. We recall that the condition number is determined by the reciprocal of the homo-lumo gap, which is reduced in the conditioning phase. When the condition number is close to 1, and the conditioning can no longer be substantially improved, we enter the purification phase in which the idempotency error is reduced. This behavior is illustrated in Figure 2.3.

It is clear that the acceleration represents an improvement of the conditioning phase which can be explained by a steeper slope of the projection polynomials at the chemical potential [37], where the occupied states are separated from the unoccupied states. As the condition number comes close to 1, and we enter the purification phase, the acceleration parameter α_i in Alg. 2 comes close to 1 and the method becomes essentially equivalent to the original method without acceleration. Thus, the acceleration is automatically turned off when the purification phase is reached and the quadratic convergence of the SP2 method sets in.

3. Eigenvalue estimates. The drawback with the acceleration technique is that estimates of interior eigenvalues are needed. Standard iterative eigenvalue methods, such as e.g. the Lanczos method, provide rapid convergence to well separated

Algorithm 2 Accelerated expansion (SP2+ACC)

```

1:  $X = \frac{\lambda_{\max} I - H}{\lambda_{\max} - \lambda_{\min}}$ 
2:  $x_1 = \frac{\lambda_{\max} - \lambda_{\text{homo}}^{\text{out}}}{\lambda_{\max} - \lambda_{\min}}$ ,  $x_2 = \frac{\lambda_{\max} - \lambda_{\text{lumo}}^{\text{out}}}{\lambda_{\max} - \lambda_{\min}}$ 
3: for  $i = 1, 2, \dots, n$  do
4:   if  $|\text{Tr}(X^2) - N_{\text{occ}}| < |\text{Tr}(2X - X^2) - N_{\text{occ}}|$  then
5:      $p_i = 1$ 
6:      $\alpha_i = 2/(2 - x_2)$ 
7:      $X = (1 - \alpha_i)I + \alpha_i X$  (shift and scale)
8:      $X = X^2$ 
9:      $x_k = (1 - \alpha_i) + \alpha_i x_k$ ,  $x_k = x_k^2$ ,  $k = 1, 2$ 
10:  else
11:     $p_i = 0$ 
12:     $\alpha_i = 2/(1 + x_1)$ 
13:     $X = \alpha_i X$  (scale)
14:     $X = 2X - X^2$ 
15:     $x_k = \alpha_i x_k$ ,  $x_k = 2x_k - x_k^2$ ,  $k = 1, 2$ 
16:  end if
17:   $v_i = \|X - X^2\|_F$ 
18:   $w_i = \text{Tr}(X - X^2)$ 
19: end for
20: return  $X$ 

```

extremal eigenvalues. When interior eigenvalues are to be computed, spectral transformations are typically employed to move the desired eigenvalues to the ends of the eigenspectrum. Transformations include shift-and-invert operators and various polynomial filters, see e.g. [1, 14]. The computational overhead that is added to the recursive expansion by incorporating iterative eigenvalue methods depends on several factors such as matrix sparsity and hardware and software details. In many cases the computational overhead would be significant. However, in the context of recursive polynomial expansions one may take advantage of the formation of a step at the chemical potential. The recursive expansion leads to extremely good separation between eigenvalues around the chemical potential which can be utilized to drastically reduce the number of Lanczos iterations [56]. Nevertheless, even if it is possible to implement such schemes with good efficiency, the use of *any* iterative method for eigenvalue calculation would add to the complexity of the method.

Eigenvalue estimates can also be cheaply computed directly from the entries of the matrix, for example on the basis of Gershgorin’s circle theorem. Unfortunately, Gershgorin’s theorem rarely provides useful information about interior eigenvalues, at least not in the context of the present work.

In this section, we propose a new simple and efficient approach to calculate estimates of the homo and lumo eigenvalues when the recursive expansion is used in an outer loop such as for example molecular dynamics time stepping. The proposed method consists of two parts. The first is calculation of accurate eigenvalue estimates directly from matrix entries in the recursive expansion of the density matrix. The second is propagation of eigenvalue estimates between time steps (or SCF iterations).

FIG. 2.3. Illustration of the two phases of the recursive expansion. In the conditioning phase, the condition number $\kappa = \Delta\epsilon/\xi$ is lowered but the idempotency error $\|X - X^2\|_2$ is not reduced. When the condition number is close to 1 we have reached the purification phase, where the idempotency error decreases quadratically. The test calculations were performed for a test Hamiltonian for which the condition number $\Delta\epsilon/\xi = 1000$ and with the chemical potential located at $\mu = \lambda_{\min} + 0.3(\lambda_{\max} - \lambda_{\min})$. The tested methods are the trace-conserving or canonical scheme (PM) [46], and the trace-correcting scheme with (SP2+ACC) and without (SP2) acceleration.

3.1. Accurate interior eigenvalue estimates directly from matrix entries. The proposed method is based on the following observations. We first recall that

$$(3.1) \quad \frac{\|A\|_F}{\sqrt{N}} \leq \|A\|_2 \leq \|A\|_F.$$

Let $\{\lambda_j\}$ be the eigenvalues of X_i . Then,

$$(3.2) \quad \|X_i - X_i^2\|_F = \sqrt{\sum_j (\lambda_j - \lambda_j^2)^2}$$

and

$$(3.3) \quad \|X_i - X_i^2\|_2 = \max_j |\lambda_j - \lambda_j^2|.$$

Note that all eigenvalues of X_i are in the $[0, 1]$ interval. The spectral norm in (3.3) is therefore given by the eigenvalue of X_i closest to 0.5. In iterations when the gap between the occupied and unoccupied parts of the eigenspectrum overlaps 0.5, this eigenvalue is either the homo or the lomo eigenvalue. Such overlap of 0.5 will in

FIG. 3.1. *Illustrative example of the purification process and the norm bounds. The SP2 algorithm is here applied to a 1000×1000 Hamiltonian matrix with band gap 0.1 and equidistant eigenvalues in the occupied and unoccupied parts of the spectrum. The upper panel shows the occupied (solid blue lines) and unoccupied (dashed red lines) parts of the eigenspectrum in each iteration. The lower panel shows bounds of the spectral norm $\|X_i - X_i^2\|_2$ that are normalized so that the spectral norm in each iteration corresponds to a value of 1. Blue circles are used for the normalized upper, $\|X_i - X_i^2\|_F / \|X_i - X_i^2\|_2$, and lower, $\|X_i - X_i^2\|_F / (\sqrt{N} \|X_i - X_i^2\|_2)$, bounds given by (3.1). Red crosses are used for the normalized lower bound, $\|X_i - X_i^2\|_F^2 / (\text{Tr}(X_i - X_i^2) \|X_i - X_i^2\|_2)$, given by (3.8).*

particular occur in the last iterations, close to convergence, see Figure 3.1. Estimates or bounds to the norm in (3.3) can then be used to obtain estimates to the homo or the lumo eigenvalue of H [56]. A spectral norm estimate is first translated to an estimate to the homo or lumo eigenvalue of X_i . Thereafter the homo/lumo estimate is inserted in the same recursion sequence that gave X_i , but performed backwards, giving an estimate to the homo or lumo eigenvalue of H .

Different ways of approximating the spectral norm value exist. Here we want to make use of (3.1) to obtain upper and lower bounds of the homo and lumo eigenvalues directly from the entries of the matrix. In the last recursive expansion iterations, most of the eigenvalues of X_i are very close to zero and one, and the sum over all eigenvalues in (3.2) is dominated by the few eigenvalues furthest from convergence. As we get closer to convergence, the Frobenius norm therefore becomes a better estimate of the spectral norm, and the upper bound given by (3.1) becomes tighter, see Figure 3.1. The lower bound given by (3.1) instead becomes looser as we approach convergence. However, in the present case with $0 \leq \lambda_j \leq 1$, $j = 1, 2, \dots, N$ we can devise a tighter lower bound.

Let $\eta_j = \lambda_j - \lambda_j^2$, $j = 1, 2, \dots, N$ be the eigenvalues of $X_i - X_i^2$. Since $\eta_j \geq 0$, $j =$

$1, 2, \dots, N$, we have that

$$(3.4) \quad \sum_j \eta_j^2 \leq \max_j \eta_j \sum_k \eta_k$$

and since

$$(3.5) \quad \sum_j \eta_j^2 = \|X_i - X_i^2\|_F^2,$$

$$(3.6) \quad \sum_j \eta_j = \text{Tr}(X_i - X_i^2)$$

and

$$(3.7) \quad \|X_i - X_i^2\|_2 = \max_j \eta_j$$

we have that

$$(3.8) \quad \frac{\|X_i - X_i^2\|_F^2}{\text{Tr}(X_i - X_i^2)} \leq \|X_i - X_i^2\|_2.$$

Both the lower bound in (3.8) and the lower bound

$$(3.9) \quad \frac{\|X_i - X_i^2\|_F}{\sqrt{N}} \leq \|X_i - X_i^2\|_2$$

given by (3.1) are sharp in the sense that if all eigenvalues of $X_i - X_i^2$ are equal, then the left hand side equals the right hand side. However, the bound in (3.8) is sharp also in the sense that if $X_i - X_i^2$ has only one nonzero eigenvalue, then the left hand side equals the right hand side as well. Many eigenvalues of $X_i - X_i^2$ are expected to be close to zero in the last recursive expansion iterations, and we then expect the bound in (3.8) to be much better. See the illustrative example in Figure 3.1.

During the recursion, the Frobenius norm $v_i = \|X_i - X_i^2\|_F$ and trace $w_i = \text{Tr}(X_i - X_i^2)$ are stored in each iteration i , see Alg. 2. By (3.1) and (3.8) we have that

$$(3.10) \quad \frac{v_i^2}{w_i} \leq \lambda - \lambda^2 \leq v_i$$

where λ is the eigenvalue of X_i that is closest to 0.5. Assume that $v_i < 1/4$, then we have that λ is either in

$$(3.11) \quad \left[\frac{1}{2}(1 - \sqrt{1 - 4v_i^2/w_i}), \frac{1}{2}(1 - \sqrt{1 - 4v_i}) \right]$$

or in

$$(3.12) \quad \left[\frac{1}{2}(1 + \sqrt{1 - 4v_i}), \frac{1}{2}(1 + \sqrt{1 - 4v_i^2/w_i}) \right].$$

The bounds $\frac{1}{2}(1 - \sqrt{1 - 4v_i^2/w_i})$ and $\frac{1}{2}(1 + \sqrt{1 - 4v_i^2/w_i})$, herein the *outer bounds*, are given by the inequality on the left hand side in (3.10). The bounds $\frac{1}{2}(1 - \sqrt{1 - 4v_i})$ and $\frac{1}{2}(1 + \sqrt{1 - 4v_i})$, herein the *inner bounds*, are given by the inequality on the right hand side in (3.10).

We want to make use of (3.11) and (3.12) to obtain bounds of the lumo and homo eigenvalues, respectively. However, before we can formulate an algorithm we need to address a few issues. The first question is how to identify iterations when the gap between occupied and unoccupied parts of the eigenspectrum overlaps 0.5, so that we can be sure that our eigenvalue bounds are bounds for the homo or the lumo eigenvalue and not some other eigenvalue.

3.1.1. Identification of eligible iterations. The present development does not depend on a particular choice of stopping criterion for the recursive expansion. Regardless of the stopping criterion used, if the recursive expansion has been deemed successful, it is reasonable to assume that the gap overlaps 0.5 in the last iteration n . It is also reasonable to assume that the homo and lumo eigenvalues are close to 1 and 0, respectively. Since the polynomials employed in the recursion set a limit on how far the eigenvalues can move from one iteration to the next we can find a number of iterations in the end of the recursion in which the gap overlaps 0.5. Frobenius norm and trace values from all those iterations can be used to obtain bounds for the homo and lumo eigenvalues.

Assume now that in iteration $i + 1$, $\lambda_{\text{homo}} > 0.5$ and that $\lambda_{\text{lumo}} < 0.5$, and that the interval $[\gamma, 1 - \gamma]$ is free from eigenvalues, where γ is a parameter between 0 and 0.5. Furthermore, assume that the interval $[\gamma, 1 - \gamma]$ is free from eigenvalues also in iteration i . Assume also that it is not possible for an eigenvalue to jump from a value larger than $1 - \gamma$ to a value smaller than γ and vice versa in one iteration. Then, we know that $\lambda_{\text{homo}} > 0.5$ and that $\lambda_{\text{lumo}} < 0.5$ also in iteration i .

The interval $[\gamma, 1 - \gamma]$ is free from eigenvalues if and only if $\|X_i - X_i^2\|_2 < \gamma - \gamma^2$. By (3.1) we have that $v_i = \|X_i - X_i^2\|_F \geq \|X_i - X_i^2\|_2$ and therefore, if $v_i < \gamma - \gamma^2$ then, the interval $[\gamma, 1 - \gamma]$ is free from eigenvalues in iteration i .

With the most extreme scaling an eigenvalue can jump from a value larger than $1 - \gamma$ to a value smaller than γ and vice versa in one iteration only if $\gamma \geq 6 - 4\sqrt{2} \approx 0.343146$. For the iteration $((1 - \alpha)I + \alpha X)^2$ the γ -value is given by solving

$$(3.13) \quad \alpha^2 \lambda^2 + 2\alpha(1 - \alpha)\lambda + (1 - \alpha)^2 = \gamma$$

where $\lambda = 1 - \gamma$ and $\alpha = \frac{2}{2 - \gamma}$, for $\gamma \in [0, 0.5]$. For the iteration $2\alpha X - (\alpha X)^2$ the γ -value is given by solving

$$(3.14) \quad 2\alpha\lambda - (\alpha\lambda)^2 = 1 - \gamma$$

where $\lambda = \gamma$ and $\alpha = \frac{2}{1 + (1 - \gamma)}$, for $\gamma \in [0, 0.5]$.

We are now ready to state the criterion for eligible iterations. Assume that in the last iteration n , $\lambda_{\text{homo}} > 0.5$ and that $\lambda_{\text{lumo}} < 0.5$. Then, the Frobenius norm and trace value from iteration i can be used for eigenvalue estimation provided that

$$(3.15) \quad v_j < \gamma - \gamma^2, j = i, i + 1, \dots, n$$

with $\gamma = 6 - 4\sqrt{2}$. We will use norm and trace values only for iterations that fulfill this criterion.

3.1.2. Knowing what eigenvalue is computed. It would now be natural to ask the question how to know if it is the homo or the lumo eigenvalue that is closest to 0.5. Our response to this question is that we do not need to know. Regardless of whether the homo or the lumo eigenvalue is closest, $\lambda - \lambda^2 \leq v_i$ holds for all eigenvalues λ of X_i . In iterations when the γ -criterion (3.15) is fulfilled, we also

know that $\lambda_{\text{lumo}} < 0.5$ and that $\lambda_{\text{homo}} > 0.5$. For those iterations, the inner bounds $\lambda_{\text{lumo}} \leq \frac{1}{2}(1 - \sqrt{1 - 4v_i})$ and $\lambda_{\text{homo}} \geq \frac{1}{2}(1 + \sqrt{1 - 4v_i})$ always hold.

The lower bound $v_i^2/w_i \leq \lambda - \lambda^2$ used to derive outer bounds, on the other hand, does not necessarily hold for all eigenvalues. However, in iterations when the γ -criterion (3.15) is fulfilled, the spectral norm $\|X_i - X_i^2\|_2$ is either given by the homo or the lumo eigenvalue. Therefore, at least one of the outer bounds $\lambda_{\text{lumo}} \geq \frac{1}{2}(1 - \sqrt{1 - 4v_i^2/w_i})$ and $\lambda_{\text{homo}} \leq \frac{1}{2}(1 + \sqrt{1 - 4v_i^2/w_i})$ holds in those iterations. This means that provided that each of the homo and lumo eigenvalues is closer than the other to 0.5 in at least one iteration for which the γ -criterion holds, the most conservative outer bounds hold.

When the polynomial in each iteration is chosen based on the trace of the matrix, as in Algs. 1 and 2 it can in principle happen that one of the homo or lumo eigenvalues dominates the norm value in all the final iterations. This could happen if there are many eigenvalues clustered very close to either the homo or the lumo eigenvalue. This problem can be avoided by choosing polynomials based on the homo and lumo eigenvalue estimates, as in [55], which would also improve convergence.

3.1.3. Effects of numerical errors. Numerical errors coming from removal of small matrix elements or floating point rounding cause eigenvalues to be disturbed. At convergence, eigenvalues should therefore be expected to deviate from their desired values of 0 and 1 with a magnitude at the level of numerical errors. When all eigenvalues are that close to their desired values of 0 or 1, the Frobenius norm and trace values do not provide particularly meaningful information about the whereabouts of the homo and lumo eigenvalues, and the recursive expansion should probably already have been terminated. However, it is important to consider how such a situation would affect our eigenvalue bounds. In general, the deviation of eigenvalues from 0 and 1 close to convergence leads to artificially large spectral (3.3) and Frobenius (3.2) norm values. Such large norm values will by our algorithm (incorrectly) be interpreted as eigenvalues lying within the homo-lumo gap, and our algorithm will produce candidates for eigenvalue bounds on the basis of those norm values. However, since the artificial eigenvalues lie within the homo-lumo gap, the inner bounds for those artificial eigenvalues also hold for the homo and lumo eigenvalues. The outer bounds, on the other hand, may not hold. In our algorithm we use the least conservative inner bounds and the most conservative outer bounds. Therefore, special treatment of iterations that may give rise to artificial bounds is generally not needed.

3.2. Algorithm for homo and lumo eigenvalues. The algorithm to calculate homo and lumo eigenvalue bounds is given in Alg. 3. For each iteration i in which the γ -criterion holds, candidates for upper and lower bounds of the homo and lumo eigenvalues of X_i are computed. This is done on lines 4-7 in Alg. 3 by making use of (3.11) and (3.12). For those values (y_1, y_2, y_3, y_4) , the recursive expansion is on lines 8-14 carried out backwards so that potential bounds for the homo and lumo eigenvalues of X_{i-1} , X_{i-2} , ..., and eventually X_0 are obtained. Then, on lines 15-16, an attempt is made to improve the inner bounds to make them tighter and the outer bounds to make them more reliable, see the discussion above. Finally, when all potential bounds have been tested, the selected bounds are on lines 19-22 translated to bounds for the homo and lumo eigenvalues of H .

The algorithm requires some information that should be collected during the course of the recursive expansion:

- The parameters $\lambda_{\min}$ and $\lambda_{\max}$ used for the initial normalization.

- A boolean vector $[p_1, p_2, \dots, p_n]$ specifying which polynomial was used in each iteration.
- A vector $[\alpha_1, \alpha_2, \dots, \alpha_n]$ containing the acceleration or scaling parameters used in each iteration.
- A vector $[v_1, v_2, \dots, v_n]$ containing the Frobenius norms of $X - X^2$ for each iteration.
- A vector $[w_1, w_2, \dots, w_n]$ containing the traces of $X - X^2$ for each iteration.

All this information can be extracted from the recursive expansion iterations at practically no extra cost, see Alg. 2.

Algorithm 3 Eigenvalue computation

```

1:  $i = n$ 
2:  $x_1 = 1, \quad x_2 = 1, \quad x_3 = 0, \quad x_4 = 0$ 
3: while  $v_i < \gamma - \gamma^2$  do
4:    $y_1 = \frac{1}{2}(1 - \sqrt{1 - 4v_i^2/w_i})$ 
5:    $y_2 = \frac{1}{2}(1 - \sqrt{1 - 4v_i})$ 
6:    $y_3 = \frac{1}{2}(1 + \sqrt{1 - 4v_i})$ 
7:    $y_4 = \frac{1}{2}(1 + \sqrt{1 - 4v_i^2/w_i})$ 
8:   for  $j = i, i-1, \dots, 2, 1$  do
9:     if  $p_j$  then
10:       $y_k = \sqrt{y_k},$ 
11:       $y_k = (y_k - 1 + \alpha_j)/\alpha_j, \quad k = 1, 2, 3, 4$ 
12:     else
13:       $y_k = 1 - \sqrt{1 - y_k},$ 
14:       $y_k = y_k/\alpha_j, \quad k = 1, 2, 3, 4$ 
15:     end if
16:   end for
17:    $x_k = \min(x_k, y_k), \quad k = 1, 2$ 
18:    $x_k = \max(x_k, y_k), \quad k = 3, 4$ 
19:    $i = i - 1$ 
20: end while
21:  $\lambda_{\text{homo}}^{\text{out}} = \lambda_{\text{max}} - (\lambda_{\text{max}} - \lambda_{\text{min}})x_4$ 
22:  $\lambda_{\text{homo}}^{\text{in}} = \lambda_{\text{max}} - (\lambda_{\text{max}} - \lambda_{\text{min}})x_3$ 
23:  $\lambda_{\text{lumo}}^{\text{in}} = \lambda_{\text{max}} - (\lambda_{\text{max}} - \lambda_{\text{min}})x_2$ 
24:  $\lambda_{\text{lumo}}^{\text{out}} = \lambda_{\text{max}} - (\lambda_{\text{max}} - \lambda_{\text{min}})x_1$ 

```

3.3. Sensitivity to system size. Since our eigenvalue bounds are based on the Frobenius norm one may be concerned that the bounds become looser when the system size is increased. Note for example that while the spectral norm of an $N \times N$ identity matrix equals 1, the Frobenius norm equals $\sqrt{N}$.

In our case we rely on the Frobenius norm of the matrix $X_i - X_i^2$ in iterations i when eigenvalues of $X_i - X_i^2$ away from the homo and lumo eigenvalues are expected to be close to zero, and thus do not contribute to the Frobenius norm (3.2). Whether the norm bounds become looser or not therefore depends on the distribution of eigenvalues. Distributions where eigenvalues are mainly added away from the homo and lumo eigenvalues when the system size is increased are most favorable. The Frobenius norm then remains a good estimate to the spectral norm, since the additional eigenvalues do not significantly contribute to the Frobenius norm. The worst scenario is that eigenvalues are only added close to the homo and lumo eigenvalues. However, if

FIG. 3.2. An example of the dependency of norm bounds and eigenvalue bounds on the system size. The SP2 algorithm is here applied to test Hamiltonians of varying size with band gap 0.1 and equidistant eigenvalues in the occupied and unoccupied parts of the spectrum. The upper panel shows the upper and lower bounds of the spectral norm given by (3.1) (blue circles) and the lower bound given by (3.8) (red crosses) for varying system size in the last SP2 iteration. As in Figure 3.1, all values have been normalized so that the spectral norm in each iteration corresponds to a value of 1, see the caption of Figure 3.1. The lower panel shows the computed intervals containing the homo and lumo eigenvalues. The homo and lumo eigenvalues are located at 0.45 and 0.55 (dashed lines).

the system size is increased by increasing the number of atoms in the system, then the density of states, i.e. the distribution of eigenvalues, is expected to remain essentially unchanged for sufficiently large systems. This scenario lies somewhere between the most favorable and worst scenarios. Figure 3.2 illustrates how the quality of the norm and eigenvalue bounds varies with increasing system size for such a scenario. In this example, the upper bound of the spectral norm becomes looser when the system size increases. The lower bound, on the other hand, is less sensitive to the increase in system size. However, the eigenvalue bounds depend only weakly on the quality of the norm bounds. The reason is that the mapping of an eigenvalue estimate for X_i back to the corresponding estimate for X_0 and H is an extremely well conditioned operation since we are taking the inverse of a function with steep slope in the region of interest. The eigenvalue bounds therefore remain accurate and stable across 8 orders of magnitude in systems size as illustrated in the example in Figure 3.2.

3.4. Propagation of eigenvalue estimates. In the previous subsections, we developed a method to extract bounds for the homo and lumo eigenvalues from the recursive polynomial expansion. This means that the homo and lumo bounds are

available only when the expansion has already completed. In order to accelerate the recursive expansion the bounds are needed in advance. However, as discussed in Section 2.1.1, the recursive expansions are often used repeatedly for a sequence of effective Hamiltonian matrices $H_0, H_1, \dots$. It is then possible to translate the eigenvalue bounds from one iteration to the next [55].

Let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$ be the eigenvalues of H_{i-1} and $\tilde{\lambda}_1 \leq \tilde{\lambda}_2 \leq \dots \leq \tilde{\lambda}_N$ the eigenvalues of H_i . Then, by Weyl's theorem on eigenvalue movement [62],

$$(3.16) \quad |\lambda_j - \tilde{\lambda}_j| \leq \|H_i - H_{i-1}\|_2, \quad j = 1, 2, \dots, N.$$

Therefore, by expanding the intervals containing the homo and lumo eigenvalues of H_{i-1} with the norm $\|H_i - H_{i-1}\|_2$, we obtain intervals containing the homo and lumo eigenvalues of H_i . If the spectral norm cannot be readily obtained, the Frobenius norm can be used, see (3.1). If the norm is small (as is typically the case between SCF cycles), the acceleration of the following recursive expansion is hardly affected by this translation. If the norm is large (as is typically the case between time steps in a molecular dynamics simulation), the acceleration will not be optimal, but we have ensured the robustness of the expansion with respect to the homo and lumo estimates.

If two successive Hamiltonian matrices differ a lot the eigenvalue bounds may become so loose that the acceleration is without effect. Also, in the initial iteration, eigenvalue bounds for H_0 may not be available. However, if no better bounds or if no bounds at all are available one may run the algorithm with $\lambda_{\text{homo}}^{\text{out}} = \lambda_{\text{min}}$ and $\lambda_{\text{lumo}}^{\text{out}} = \lambda_{\text{max}}$ which makes Alg. 2 essentially equivalent to Alg. 1. Thus, our method works also in case of loose outer bounds, but less efficiently. In a molecular dynamics simulation, no time step size restriction is therefore needed to make the eigenvalue bounds tight enough. At worst, Alg. 2 falls back to the original SP2 algorithm without acceleration.

4. Quantum mechanical molecular dynamics.

4.1. Born–Oppenheimer molecular dynamics. Among computationally feasible methods for molecular dynamics simulations, Born–Oppenheimer molecular dynamics based on self-consistent field methods, such as density functional theory [12, 21, 29, 47], is often considered a gold standard. In Born–Oppenheimer molecular dynamics the forces acting on the atoms are explicitly quantum mechanical and calculated at the relaxed electronic ground state [33]. The electronic ground state is given through the self-consistent field optimization procedure, discussed in Section 2.1.1, which involves iterative solutions of the density matrix and Hamiltonian in (2.1). The self-consistent field optimization accounts for details in the charge distribution and is important for the accuracy of the interatomic forces.

In regular Born–Oppenheimer molecular dynamics the computational cost, which is dominated by the iterative self-consistent field optimization, is significantly reduced by using a good initial guess to the optimization procedure, which is naturally provided through an extrapolation of the ground state solutions from previous time steps. Since the iterative self-consistent field optimization in practice never is complete and always approximate, this extrapolation technique generates a time-dependent fictitious dynamics of the underlying electronic degrees of freedom. Because of the non-linearity of the self-consistent field optimization this process is irreversible, which typically is manifested in an unphysical systematic drift in the energy and phase space [42, 50], where the electronic degrees of freedom act like a heat sink or source, constantly removing or adding energy to the system. Only by increasing the number of self-consistent field

iterations can the energy drift be reduced, though it never fully disappears. Recently an extended Lagrangian formulation of Born–Oppenheimer molecular dynamics (XL-BOMD) was introduced that overcomes these fundamental problems by restoring time-reversibility and symplecticity to the dynamics [8, 38, 39, 40, 43, 61, 67].

Instead of using an extrapolation of the electronic ground state from previous time steps, as in regular Born–Oppenheimer molecular dynamics, an auxiliary electron density $n(\mathbf{r})$ is introduced in the extended Lagrangian formulation as an extended dynamical variable that is centered around the ground state density $\rho(\mathbf{r})$ through a harmonic well. The equations of motion for the nuclear coordinates $\mathbf{R} = \{R_I\}$ and the extended density are given by

$$(4.1) \quad M_I \ddot{R}_I = -\frac{\partial U(\mathbf{R}; \rho)}{\partial R_I},$$

$$(4.2) \quad \ddot{n}(\mathbf{r}) = \omega^2 (\rho(\mathbf{r}) - n(\mathbf{r})),$$

where M_I are the nuclear masses and ω is a frequency parameter that determines the curvature of the harmonic well. $U(\mathbf{R}; \rho)$ is the potential energy term that describes the electronic and ionic energy at the self-consistent electronic ground state. The dots denote time-derivatives.

The extended Lagrangian molecular dynamics, (4.1) and (4.2), has several important features: *i)* $n(\mathbf{r})$ moves in a harmonic well centered around the ground state density $\rho(\mathbf{r})$ and is therefore a good initial guess to the iterative self-consistent field optimization procedure, *ii)* since $n(\mathbf{r})$ is a dynamical variable, in contrast to $\rho(\mathbf{r})$, it can be integrated with a geometric reversible scheme in the same way as the nuclear coordinates [30, 38, 43], *iii)* in this way reversibility in the underlying propagation of the electronic degrees of freedom is not broken, and *iv)* since the equations of motion for the nuclear coordinate is identical to “exact” Born–Oppenheimer molecular dynamics, the total Born–Oppenheimer energy,

$$(4.3) \quad E_{\text{tot}} = \frac{1}{2} \sum_I M_I \dot{R}_I^2 + U(\mathbf{R}; \rho),$$

is still a constant of motion that should exhibit long-term stability [38].

Since our estimates of the homo-lumo eigenvalues, which are used for the scale-and-fold acceleration in the construction of the density matrix, are given from previous time steps, it is important to show that this additional artificial propagation does not have any effect on the reversibility and long-term energy conservation. Here we will demonstrate how the accelerated SP2 density matrix expansion using estimates of the homo-lumo eigenvalues from the previous time step is fully compatible with extended Lagrangian Born–Oppenheimer molecular dynamics, without causing any systematic drift in the energy. Our accelerated expansion algorithm therefore provides a significant reduction of the computational cost, while keeping the accuracy, for an already highly efficient scheme.

4.2. Numerical experiments. In our numerical experiments the interaction potential $U(\mathbf{R}; \rho)$ in (4.1) is based on the self-consistent density functional tight-binding (SCC-DFTB) formulation of density functional theory [13, 15, 16, 17, 66] as implemented in the program LATTE [8, 9, 39, 44, 58]. The equations of motion are integrated with the Verlet scheme, in its symplectic velocity formulation for the

FIG. 4.1. The conservation of the total energy, (4.3), for regular Born–Oppenheimer molecular dynamics (BOMD) and extended Lagrangian Born–Oppenheimer molecular dynamics (XL-BOMD), with and without the acceleration scheme (ACC) for the SP2 algorithm. The two lower curves (red and blue) are virtually identical.

nuclear degrees of freedom and using a regular Verlet scheme for the electronic degrees of freedom with a dissipative force term that removes any accumulation of numerical noise [40, 61]. In all cases the SP2 algorithm, with (Alg. 2) or without (Alg. 1) acceleration, was used with the recursions terminated when the occupation change, as measured by the difference between using X^2 or $2X - X^2$, was less than 10^{-10} . Two self-consistent field iterations were used in each time step, which together with the force calculation sums up to totally three density matrix constructions per time step. The homo and lumo eigenvalue bounds were translated from one time step or self-consistent field iteration to the next as described in Section 3.4, using the Frobenius matrix norm in (3.16). All calculations were performed using regular dense (non-sparse) matrix algebra.

Figure 4.1 shows the stability of the total Born–Oppenheimer energy, E_{tot} in (4.3), either with regular Born–Oppenheimer molecular dynamics using the electronic ground state solution from the previous time step as an initial guess to the self-consistent field optimization or with extended Lagrangian Born–Oppenheimer molecular dynamics, (4.1) and (4.2), with or without the acceleration scheme. The systematic drift occurring in regular Born–Oppenheimer molecular dynamics vanish in the extended Lagrangian formulation. Moreover, we find virtually no effect in the total energy from using the accelerated SP2 scheme with estimates of the homo and lumo eigenvalues obtained from the previous iteration. The two graphs are practically identical even after thousands of time steps.

FIG. 4.2. The average number of matrix-matrix multiplications of the SP2 algorithm during molecular dynamics simulations with and without acceleration.

The gain in efficiency using the acceleration technique is significant. Figure 4.2 shows the average number of matrix-matrix multiplications required in the density matrix expansion in molecular dynamics simulations of various systems. For the simulation of the polyethene molecule we chose a 100 carbon atom chain ($C_{100}H_{202}$) and the methane example consisted of liquid methane $(CH_4)_{100}$. All the simulations were performed with a time step of 0.5 fs using a structurally optimized ground state configuration with an initial Gaussian velocity distribution corresponding to 550 K. The efficiency gain is governed by the size of the homo-lumo gap [52]. For small gap systems, like naphthalene, the scale-and-fold acceleration gives a significantly higher reduction of the cost compared to large gap systems, like methane. In other words, the acceleration works best for ill-conditioned systems, as defined in Section 2.1. This also means that there is a much smaller variation in the number of matrix multiplications required for the accelerated scheme. The work load is therefore not only smaller but also less sensitive to the particular choice of materials system. Furthermore, the computational cost is more evenly distributed over time in a molecular dynamics simulation when there are large variations in the homo-lumo gap.

Figure 4.3 (panel a-d) illustrates the effect of the acceleration as a function of simulation time for a system of liquid isocyanic acid using XL-BOMD. As seen in the two upper panels (a and b) the energetics between the regular and the accelerated scheme, as measured by the total energy or the temperature, are indistinguishable. Figure 4.3 (panel d) also shows how the number of matrix multiplications for the density matrix construction (that is required in the last self-consistent field iteration in each time step) varies over time. Variations in the work load may be sensitive to the choice of convergence criterion used to terminate the SP2 recursions, which here is set fairly tight. The accuracy of the homo-lumo gap estimate, shown in panel (c) is surprisingly good, in particular as measured by the inner (Inner) $\lambda_{\text{homo}}^{\text{in}}$ and $\lambda_{\text{lumo}}^{\text{in}}$ eigenvalue bounds. The upper bound of the gap, given by the outer (Outer)

FIG. 4.3. The fluctuations during a *XL-BOMD* simulation of liquid isocyanic acid $(\text{HCNO})_{56}$ of the total energy (*Etot* in panel a), the temperature (*Temp* in panel b), the gap estimates (panel c), and the variation of the number of matrix-matrix multiplications (panel d) of the *SP2* algorithm with and without acceleration (*ACC*).

eigenvalue bounds, is good but not as tight. The reason is that for the outer bounds we choose the most conservative values to make sure that we get valid bounds. It is in principle possible to devise a scheme that keeps track of which eigenvalue (the homo or luma eigenvalue) that corresponds to the spectral norm of $X - X^2$. This information can be used to obtain more accurate outer homo-luma bounds. However, such an algorithm would be more complicated with small gains in efficiency.

5. Discussion. In this article we have devised a surprisingly simple technique to extract information about the homo-luma gap that enables a cost-efficient practical application of the scale-and-fold acceleration technique in molecular dynamics simulations. The technique is quite general and should be applicable to a broad class of recursive density matrix expansion algorithms besides the *SP2* scheme described here [52, 63]. As a stand-alone technique, our method for calculating the homo-luma gap estimates may also be useful in studies of materials properties.

Scale-and-fold acceleration is not the only way to improve or design efficient recursive expansion polynomials. It is in principle possible to construct a number of “optimal” recursive expansion schemes where the projection polynomials, $\{f_n\}$, in the recursive expansion (2.4) are defined, for example, by minimizing the error in some given norm for a certain number of recursion steps, N , and with a maximum number,

$M_{\max}$, of necessary matrix multiplications, M , e.g.

$$(5.1) \{f_n\}_{n=0}^N = \arg \min_{\{f_n\}} \{\|\theta(\mu I - H) - f_N(f_{N-1}(\dots f_0(H)\dots))\|; M \leq M_{\max}\}.$$

With a detailed knowledge about the particular eigenvalue spectrum of H this problem can be solved using a number of various non-linear optimization techniques. The key issue with such an optimized recursion scheme, (5.1), is that we in general do not have prior knowledge of the interior eigenvalue distribution. Moreover, in a molecular dynamics simulation or a self-consistent field optimization the optimization would have to be performed on-the-fly in each iteration. Alternatively, optimizations for different classes of distributions could be made, e.g. for uniform eigenvalue distributions with various homo-lumo gaps and spectral bounds, that could be estimated with our scheme for interior eigenvalues. Our accelerated SP2 recursion scheme may very well belong to such a hypothetical class of optimized expansion methods.

In this article we have taken a cautious standpoint with respect to the homo and lumino estimates used in the acceleration. Although the inner bounds of the homo and lumino eigenvalues have proven to be extremely tight, we have used the more conservative outer bounds for the acceleration. Moreover, those outer bounds have been moved further away from the gap as a result of the propagation between time steps or self-consistent field iterations by (3.16). All those measures are not strictly necessary in order to avoid mixing up occupied and unoccupied states and lead to a less efficient acceleration. However, we also have to avoid folding extremal eigenvalues into the gap since this could lead to invalid homo and lumino eigenvalue bounds in the following iterations. With the conservative measures taken in this work this will never happen, though in practice it is often possible to use less restrictive homo-lumo estimates leading to a more efficient acceleration.

The experiments presented in this work were all carried out using dense matrix algebra. The use of sparse matrix algebra to achieve a linear scaling algorithm involves additional issues regarding ways to bring about matrix sparsity without introducing uncontrolled errors [54]. Recursive expansions with rigorous error control requires inner homo and lumino eigenvalue bounds [55]. Thus, the eigenvalue estimates of the present work are also expected to be useful in schemes to select small matrix elements for removal to achieve error control in the linear scaling construction of the density matrix. Such use of the eigenvalue estimates has not been explored in the present work.

6. Conclusions. Interior eigenvalues of the effective Hamiltonian matrix are needed for acceleration and error control of density matrix expansions in electronic structure calculations. Unfortunately, interior eigenvalues are often difficult and expensive to come by. In this work, we have developed a simple extension to recursive density matrix expansions to extract accurate interior (homo-lumo) eigenvalue estimates at practically no extra computational cost.

The new eigenvalue estimates were here used for the acceleration of the recursive polynomial expansion. The behavior and performance of the eigenvalue estimates in combination with the the scale-and-fold acceleration method were illustrated with density functional tight-binding molecular dynamics simulations, where the computational effort is dominated by the density matrix construction. Our scheme was found to be fully compatible with extended Lagrangian Born–Oppenheimer molecular dynamics and long-term energy stability.

We identified two different phases in recursive density matrix expansions, the conditioning phase and the purification phase. With this point of view it was clear that

the acceleration technique represents an improvement of the conditioning phase. In this way, the acceleration is analogous to preconditioners used with iterative solvers for linear systems. By using prior knowledge of the system, the condition number of the problem is more rapidly reduced in the (pre-)conditioning phase, which gives a significant reduction of the computational cost, as demonstrated in quantum mechanical molecular dynamics simulations.

REFERENCES

- [1] Z. BAI, J. DEMMEL, J. DONGARRA, A. RUHE, AND H. VAN DER VORST, eds., *Templates for the solution of algebraic eigenvalue problems*, SIAM, Philadelphia, 2000.
- [2] M. BARRAULT, E. CANCÈS, W. W. HAGER, AND C. LE BRIS, *Multilevel domain decomposition for electronic structure calculations*, J. Comput. Phys., 222 (2007), pp. 86–109.
- [3] M. BENZI, P. BOITO, AND N. RAZOUK, *Decay properties of spectral projectors with applications to electronic structure*, SIAM Review, 55 (2013), pp. 3–64.
- [4] E. G. BIRGIN, J. M. MARTÍNEZ, L. MARTÍNEZ, AND G. B. ROCHA, *Sparse projected-gradient method as a linear-scaling low-memory alternative to diagonalization in self-consistent field electronic structure calculations*, J. Chem. Theory Comput., 9 (2013), pp. 1043–1051.
- [5] D. R. BOWLER AND T. MIYAZAKI, *Calculations for millions of atoms with density functional theory: linear scaling shows its potential*, J. Phys.: Condens. Matter, 22 (2010), p. 074207.
- [6] ———, *$O(n)$ methods in electronic structure calculations*, Rep. Prog. Phys., 75 (2012), pp. 036503–036546.
- [7] C. G. BROYDEN, *A class of methods for solving nonlinear simultaneous equations*, Math. Comp., 19 (1965), p. 557.
- [8] M. J. CAWKWELL AND A. M. N. NIKLASSON, *Energy conserving, linear scaling Born–Oppenheimer molecular dynamics*, J. Chem. Phys., 137 (2012), p. 134105.
- [9] M. J. CAWKWELL, E. J. SANVILLE, S. M. MNISZEWSKI, AND A. M. N. NIKLASSON, *Computing the density matrix in electronic structure theory on graphics processing units*, J. Chem. Theory Comput., 8 (2012), p. 4094.
- [10] M. CHALLACOMBE, *A simplified density matrix minimization for linear scaling self-consistent field theory*, J. Chem. Phys., 110 (1999), pp. 2332–2342.
- [11] M. S. DAW, *Model for energetics of solids based on the density matrix*, Phys. Rev. B, 47 (1993), pp. 10895–10898.
- [12] R. M. DREIZLER AND K. U. GROSS, *Density-functional theory*, Springer Verlag, Berlin Heidelberg, 1990.
- [13] M. ELSTNER, D. PORESAG, G. JUNGnickel, J. ELSTNER, M. HAUGK, T. FRAUENHEIM, S. SUHAI, AND G. SEIFERT, *Self-consistent-charge density-functional tight-binding method for simulations of complex materials properties*, Phys. Rev. B, 58 (1998), p. 7260.
- [14] H. FANG AND Y. SAAD, *A filtered Lanczos procedure for extreme and interior eigenvalue problems*, SIAM J. Sci. Comput., 34 (2012), pp. A2220–A2246.
- [15] M. FINNIS, *Interatomic forces in condensed matter*, Oxford University Press, 2003.
- [16] M. FINNIS, A. T. PAXTON, M. METHFESSEL, AND M. VAN SCHILFGAARDE, *Crystal structures of zirconia from first principles and self-consistent tight binding*, Phys. Rev. Lett., 81 (1998), p. 5149.
- [17] T. FRAUENHEIM, G. SEIFERT, M. ELSTNER, Z. HAJNAL, G. JUNGnickel, D. PORESAG, S. SUHAI, AND R. SCHOLZ, *A self-consistent charge density-functional based tight-binding method for predictive materials simulations in physics, chemistry and biology*, Phys. stat. sol. (b), 217 (2000), p. 41.
- [18] S. GOEDECKER, *Linear scaling electronic structure methods*, Rev. Mod. Phys., 71 (1999), pp. 1085–1123.
- [19] S. GOEDECKER AND L. COLOMBO, *Efficient linear scaling algorithm for tight-binding molecular-dynamics*, Phys. Rev. Lett., 73 (1994), pp. 122–125.
- [20] N. J. HIGHAM, *Functions of matrices: theory and computation*, SIAM, Philadelphia, 2008.
- [21] P. HOHENBERG AND W. KOHN, *Inhomogeneous electron gas*, Phys. Rev., 136 (1964), pp. B:864–B871.
- [22] A. HOLAS, *Transforms for idempotency purification of density matrices in linear-scaling electronic-structure calculations*, Chem. Phys. Lett., 340 (2001), pp. 552–558.
- [23] D. D. JOHNSON, *Modified Broyden’s method for accelerating convergence in self-consistent calculations*, Phys. Rev. B, 38 (1988), p. 12807.
- [24] J. KIM AND Y. JUNG, *Accelerated purification using generalized nonpurifying intermediate*

- functions for large-scale self-consistent field calculations*, J. Chem. Theory Comput., 7 (2011), pp. 3853–3858.
- [25] J. KIM, F. MAURI, AND G. GALLI, *Total-energy global optimizations using nonorthogonal localized orbitals*, Phys. Rev. B, 52 (1995), pp. 1640–1648.
 - [26] B. KIRCHNER, J. DI DIO PHILIPP, AND J. HUTTER, *Real-world predictions from ab initio molecular dynamics simulations*, Top. Curr. Chem., 307 (2012), p. 109.
 - [27] W. KOHN, *Analytic properties of Bloch waves and Wannier functions*, Phys. Rev., 115 (1959), pp. 809–821.
 - [28] ———, *Density functional and density matrix method scaling linearly with the number of atoms*, Phys. Rev. Lett., 76 (1996), p. 3168.
 - [29] W. KOHN AND L. J. SHAM, *Self-consistent equations including exchange and correlation effects*, Phys. Rev., 140 (1965), p. 1133.
 - [30] B. LEIMKUHLER AND S. REICH, *Simulating Hamiltonian Dynamics*, Cambridge University Press, 2004.
 - [31] X.-P. LI, R. W. NUNES, AND D. VANDERBILT, *Density-matrix electronic-structure method with linear system-size scaling*, Phys. Rev. B, 47 (1993), pp. 10891–10894.
 - [32] W. LIANG, C. SARAVANAN, Y. SHAO, R. BAER, A. T. BELL, AND M. HEAD-GORDON, *Improved Fermi operator expansion methods for fast electronic structure calculations*, J. Chem. Phys., 119 (2003), pp. 4117–4125.
 - [33] D. MARX AND J. HUTTER, *Modern Methods and Algorithms of Quantum Chemistry*, ed. J. Grotendorst, John von Neumann Institute for Computing, Jülich, Germany, second ed., 2000.
 - [34] D. A. MAZZIOTTI, *Towards idempotent reduced density matrices via particle-hole duality: Mcweeny’s purification and beyond*, Phys. Rev. E, 68 (2003), p. 066701.
 - [35] R. MCWEENY, *The density matrix in self-consistent field theory. i. iterative construction of the density matrix*, Proc. R. Soc. London Ser. A, 235 (1956), pp. 496–509.
 - [36] ———, *Some recent advances in density matrix theory*, Rev. Mod. Phys., 32 (1960), p. 335.
 - [37] A. M. N. NIKLASSON, *Expansion algorithm for the density matrix*, Phys. Rev. B, 66 (2002), p. 155115.
 - [38] A. M. N. NIKLASSON, *Extended Born–Oppenheimer molecular dynamics*, Phys. Rev. Lett., 100 (2008), p. 123004.
 - [39] A. M. N. NIKLASSON AND M. J. CAWKWELL, *Fast method for quantum mechanical molecular dynamics*, Phys. Rev. B, 86 (2012), p. 174308.
 - [40] A. M. N. NIKLASSON, P. STENETEG, A. ODELL, N. BOCK, M. CHALLACOMBE, C. J. TYMCZAK, E. HOLMSTRÖM, G. ZHENG, AND V. WEBER, *Time-reversible Born–Oppenheimer molecular dynamics with dissipation*, J. Chem. Phys., 130 (2009), p. 214109.
 - [41] A. M. N. NIKLASSON, C. J. TYMCZAK, AND M. CHALLACOMBE, *Trace resetting density matrix purification in $O(N)$ self-consistent-field theory*, J. Chem. Phys., 118 (2003), p. 8611.
 - [42] A. M. N. NIKLASSON, C. J. TYMCZAK, AND M. CHALLACOMBE, *Time-reversible Born–Oppenheimer molecular dynamics*, Phys. Rev. Lett., 97 (2006), p. 123001.
 - [43] A. ODELL, A. DELIN, B. JOHANSSON, N. BOCK, M. CHALLACOMBE, AND A. M. N. NIKLASSON, *Higher-order symplectic integration in Born–Oppenheimer molecular dynamics*, J. Chem. Phys., 131 (2009), p. 244106.
 - [44] A. ODELL, A. DELIN, B. JOHANSSON, M. J. CAWKWELL, AND A. M. N. NIKLASSON, *Geometric integration in Born–Oppenheimer molecular dynamics*, J. Chem. Phys., 135 (2011), p. 224105.
 - [45] P. ORDEJÓN, D. A. DRABOLD, M. P. GRUMBACH, AND R. M. MARTIN, *Unconstrained minimization approach for electronic computations that scales linearly with system size*, Phys. Rev. B, 48 (1993), pp. 14646–14649.
 - [46] A. H. R. PALSER AND D. E. MANOLOPOULOS, *Canonical purification of the density matrix in electronic-structure theory*, Phys. Rev. B, 58 (1998), p. 12704.
 - [47] R. G. PARR AND W. YANG, *Density-functional theory of atoms and molecules*, Oxford University Press, Oxford, 1989.
 - [48] M. S. PATERSON AND L. STOCKMEYER, *On the number of nonscalar multiplications necessary to evaluate matrix polynomials*, SIAM J. Comput., 2 (1973), pp. 60–66.
 - [49] P. PULAY, *Convergence acceleration of iterative sequences. the case of scf iteration*, Chem. Phys. Lett., 73 (1980), pp. 393–398.
 - [50] P. PULAY AND G. FOGARASI, *Fock matrix dynamics*, Chem. Phys. Lett., 386 (2004), p. 272.
 - [51] C. C. J. Roothaan, *New developments in molecular orbital theory*, Rev. Mod. Phys., 23 (1951), pp. 69–89.
 - [52] E. H. RUBENSSON, *Nonmonotonic recursive polynomial expansions for linear scaling calculation of the density matrix*, J. Chem. Theory Comput., 7 (2011), pp. 1233–1236.

- [53] E. H. RUBENSSON, *Controlling errors in recursive Fermi–Dirac operator expansions with applications in electronic structure theory*, SIAM J. Sci. Comput., 34 (2012), pp. B1–B23.
- [54] E. H. RUBENSSON AND E. RUDBERG, *Bringing about matrix sparsity in linear scaling electronic structure calculations*, J. Comput. Chem., 32 (2011), pp. 1411–1423.
- [55] E. H. RUBENSSON, E. RUDBERG, AND P. SALEK, *Density matrix purification with rigorous error control*, J. Chem. Phys., 128 (2008), p. 074106.
- [56] E. H. RUBENSSON AND S. ZAHEDI, *Computation of interior eigenvalues in electronic structure calculations facilitated by density matrix purification*, J. Chem. Phys., 128 (2008), p. 176101.
- [57] E. RUDBERG AND E. H. RUBENSSON, *Assessment of density matrix methods for linear scaling electronic structure calculations*, J. Phys.: Condens. Matter, 23 (2011), p. 075502.
- [58] E. SANVILLE, N. BOCK, W. M. CHALLACOMBE, A. M. N. NIKLASSON, M. J. CAWKWELL, D. M. DATTELBAUM, AND S. SHEFFIELD, *Proceedings of the Fourteenth International Detonation Symposium*, Office of Naval Research, Arlington VA, ONR-351-10-185, 2010, pp. 91–101.
- [59] R. N. SILVER AND H. RÖDER, *Densities of states of mega-dimensional Hamiltonian matrices*, Int. J. Mod. Phys. C, 5 (1994), pp. 735–753.
- [60] R. N. SILVER, H. RÖDER, A. F. VOTER, AND J. D. KRESS, *Kernel polynomial approximations for densities of states and spectral functions*, Int. J. Comput. Phys., 124 (1996), p. 115.
- [61] P. STENETEG, I. A. ABRİKOSOV, V. WEBER, AND A. M. N. NIKLASSON, *Wave function extended lagrangian Born–Oppenheimer molecular dynamics*, Phys. Rev. B, 82 (2010), p. 075110.
- [62] G. W. STEWART AND J. SUN, *Matrix perturbation theory*, Academic Press, San Diego, 1990.
- [63] P. SURYANARAYANA, *Optimized purification for density matrix calculation*, Chem. Phys. Lett., 555 (2013), pp. 291–295.
- [64] M. TUCKERMAN, *Ab initio molecular dynamics: basic concepts, current trends and novel applications*, J. Phys.: Condens. Matter, 50 (2002), p. 1297.
- [65] J. VANDEVONDELE, U. BORSTNIK, AND J. HUTTER, *Linear scaling self-consistent field calculations with millions of atoms in the condensed phase*, J. Chem. Theory Comput., 8 (2012), pp. 3565–3573.
- [66] G. ZHENG, S. IRLE, AND K. MOROKUMA, *Performance of the DFTB method in comparison to DFT and semiempirical methods for geometries and energies of c20c86 fullerene isomers*, Chem. Phys. Lett., 412 (2005), p. 210.
- [67] G. ZHENG, A. M. N. NIKLASSON, AND M. KARPLUS, *Lagrangian formulation with dissipation of Born–Oppenheimer molecular dynamics using the density-functional tight-binding method*, J. Chem. Phys., 135 (2011), p. 044122.