

Robust determination of the chemical potential in the pole expansion and selected inversion method for solving Kohn-Sham density functional theory

Weile Jia, and Lin Lin

Citation: *The Journal of Chemical Physics* **147**, 144107 (2017); doi: 10.1063/1.5000255

View online: <http://dx.doi.org/10.1063/1.5000255>

View Table of Contents: <http://aip.scitation.org/toc/jcp/147/14>

Published by the [American Institute of Physics](#)



Scilight

Sharp, quick summaries **illuminating**
the latest physics research

Sign up for **FREE!**

AIP
Publishing

Robust determination of the chemical potential in the pole expansion and selected inversion method for solving Kohn-Sham density functional theory

Weile Jia^{1,a)} and Lin Lin^{1,2,b)}

¹Department of Mathematics, University of California, Berkeley, California 94720, USA

²Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

(Received 14 August 2017; accepted 24 September 2017; published online 11 October 2017)

Fermi operator expansion (FOE) methods are powerful alternatives to diagonalization type methods for solving Kohn-Sham density functional theory (KSDFT). One example is the pole expansion and selected inversion (PEXSI) method, which approximates the Fermi operator by rational matrix functions and reduces the computational complexity to at most quadratic scaling for solving KSDFT. Unlike diagonalization type methods, the chemical potential often cannot be directly read off from the result of a single step of evaluation of the Fermi operator. Hence multiple evaluations are needed to be sequentially performed to compute the chemical potential to ensure the correct number of electrons within a given tolerance. This hinders the performance of FOE methods in practice. In this paper, we develop an efficient and robust strategy to determine the chemical potential in the context of the PEXSI method. The main idea of the new method is not to find the exact chemical potential at each self-consistent-field (SCF) iteration but to dynamically and rigorously update the upper and lower bounds for the true chemical potential, so that the chemical potential reaches its convergence along the SCF iteration. Instead of evaluating the Fermi operator for multiple times sequentially, our method uses a two-level strategy that evaluates the Fermi operators in parallel. In the regime of full parallelization, the wall clock time of each SCF iteration is always close to the time for one single evaluation of the Fermi operator, even when the initial guess is far away from the converged solution. We demonstrate the effectiveness of the new method using examples with metallic and insulating characters, as well as results from *ab initio* molecular dynamics. *Published by AIP Publishing.* <https://doi.org/10.1063/1.5000255>

I. INTRODUCTION

Kohn-Sham density functional theory (KSDFT) is the most widely used theory for electronic-structure calculations. In the framework of the self-consistent field (SCF) iteration, the computational cost for solving KSDFT is mainly determined by the cost associated with the evaluation of the electron density for a given Kohn-Sham potential during each iteration. The most widely used method to perform such an evaluation is to partially or fully diagonalize the Kohn-Sham Hamiltonian, by means of computing a set of eigenvectors corresponding to the algebraically smallest eigenvalues of the Hamiltonian. The complexity of this approach is $\mathcal{O}(N_e^3)$, where N_e is the number of electrons in the atomistic system of interest. As the number of atoms or electrons in the system increases, the cost of this diagonalization step becomes prohibitively expensive.

In the past two decades, various numerical algorithms have been developed for solving KSDFT without invoking the diagonalization procedure. One particular class of algorithms are the linear scaling algorithms,^{1–8} which relies on the near-sightedness principle for insulating systems with large gaps^{9,10} to truncate elements of the density matrix away

from the diagonal. Among such methods, the Fermi operator expansion (FOE) method¹¹ was also originally proposed as a linear scaling method for insulating systems. Recently, the pole expansion and selected inversion (PEXSI) method, which can be viewed as a FOE method by approximating the Fermi operator using rational matrix functions, first achieved computational complexity that is at most $\mathcal{O}(N_e^2)$ for both insulating and metallic systems.^{12–14} More specifically, the computational complexity of PEXSI depends on the dimensionality of the system: the cost for quasi-1D systems such as nanotubes is $\mathcal{O}(N_e)$, i.e., linear scaling; for quasi-2D systems such as graphene and surfaces (slabs) it is $\mathcal{O}(N_e^{1.5})$; and for general 3D bulk systems, it is $\mathcal{O}(N_e^2)$. PEXSI can be accurately applied to general material systems including small gapped systems and metallic systems and remains accurate at low temperatures. The PEXSI method has a two-level parallelism structure and is by design highly scalable using 10 000 ~ 100 000 processors on high performance machines. The PEXSI software package¹⁵ has been integrated into a number of electronic structure software packages such as BigDFT,¹⁶ CP2K,¹⁷ SIESTA,^{18,19} DGDFT,^{20,21} FHI-aims,²² and QuantumWise ATK²³ and has been used for accelerating materials simulation with more than 10 000 atoms.^{24,25}

One challenge for the PEXSI method, and for FOE methods using rational approximations in general, is to determine

^{a)}jiaweile@berkeley.edu

^{b)}linlin@math.berkeley.edu

the chemical potential μ so that the computed number of electrons at each SCF iteration is N_e within some given tolerance. This amounts to solving a scalar equation

$$N_\beta(\mu) = N_e. \quad (1)$$

Here $N_\beta(\mu)$ is the number of electrons evaluated using PEXSI at a given chemical potential and $\beta = 1/(k_B T)$ is the inverse temperature.

Note that $N_\beta(\mu)$ is a monotonically non-decreasing function with respect to μ , and the simplest strategy to solve Eq. (1) is the bisection method. However, starting from a reasonably large search interval, the bisection method can require tens of iterations to converge. This makes it more difficult to reach the crossover point compared to diagonalization type methods and hinders the effectiveness of the method. It also introduces potentially large fluctuation in terms of the running time among different SCF iterations. One option to accelerate the convergence of the bisection method is to use Newton's method, which takes the derivative information into account and is expected to converge within a few iterations. However, the effectiveness of Newton's method relies on the assumption that the derivative of $N_\beta(\mu)$ with respect to μ does not vanish. This assumption fails whenever μ is inside a bandgap. A possible remedy is to use a regularized derivative, but this may instead reduce the convergence rate and increase the number of PEXSI evaluations per SCF iteration. Furthermore, when the initial guess is far away from the true chemical potential, the derivative information is not very useful in general. A robust algorithm needs to handle all cases efficiently and find the solution within at most a handful of evaluations starting from a possible wild initial guess. Hence the seemingly innocent scalar equation (1) turns out to be not so easy in practice.

In Ref. 18, we have proposed a hybrid Newton type method for determining the chemical potential in the context of the PEXSI method. This method uses an inertia counting strategy to rapidly reduce the size of the search interval for the chemical potential, starting from a large search interval. When the search interval becomes sufficiently small, a Newton type method is then used. With the help of the inertia counting strategy, the number of Newton steps is in general small (usually no more than 5), and each Newton step amounts to one step of evaluation of the Fermi operator. As discussed above, the Newton step may still occasionally over-correct the chemical potential due to the small but not vanishing derivative information, and the correction needs to be discarded when it exceeds a certain threshold value. In such a case, the inertia counting procedure needs to be invoked again with some updated searching criterion. In this unfavorable regime, the effectiveness of the PEXSI method is hindered. This procedure also inevitably introduces extra tunable parameters, of which the values may be difficult to predict *a priori*.

In this work, we develop a new method for determining the chemical potential that simultaneously improves the robustness and the efficiency of the PEXSI method. Our main idea is to relax the requirement of satisfying Eq. (1) for each SCF iteration, so that the error of the chemical potential is comparable to the residual error in the SCF iteration, and the number of evaluations of the Fermi operator can therefore be reduced.

In particular, when the SCF converges, the chemical potential converges as well and there is no loss of accuracy. The key to achieve this goal is to maintain rigorous lower and upper bounds for the true chemical potential at each SCF step and to dynamically update the interval along the progress of the SCF iteration. Due to the availability of these rigorous bounds, the new method does not suffer from the over-correction problem as in Newton type methods. In the regime of full parallelization, the wall clock time of the new method during each SCF iteration is always approximately the same as that for one single evaluation of the Fermi operator, even when the initial guess of the chemical potential is far away from the true solution.

The simplicity provided by the rigorous bounds of the chemical potential also reduces the number of tunable parameters in the practical implementation of the PEXSI method. We find that the remaining parameters are much less system dependent, and the default values are already robust for both insulating and metallic systems. This facilitates the usage of the PEXSI package as a black box software package for general systems and is therefore more user-friendly. We also compare the pole expansion obtained by an contour integral approach¹² and that obtained from an optimization procedure by Moussa recently.²⁶ We find that Moussa's method can reduce the number of poles by a factor of $2 \sim 3$ compared to the contour integral approach, and hence we adopt this method as the default option for pole expansion. Our new method is implemented in the PEXSI software package.¹⁵ For electronic structure calculations, the PEXSI method can be accessed more easily from the recently developed "Electronic Structure Infrastructure" (ELSI) software package.²⁷ We demonstrate the performance of the new method using the discontinuous Galerkin DFT (DGDFT) software package,^{20,21} using graphene and phosphorene as examples for metallic and insulating systems, respectively. Even though the chemical potential is not fully accurate in each step of the SCF iteration in the new method, we find that the number of SCF iterations required to converge is almost the same as that needed by full diagonalization methods, as well as the method in Ref. 18 which requires multiple PEXSI evaluations per SCF iteration. The accuracy of the new method is further confirmed by *ab initio* molecular dynamics, where we observe negligible energy drift in an NVE simulation for a graphene system.

The rest of the paper is organized as follows. After reviewing the PEXSI method in Sec. II, we describe the new method for robust determination of the chemical potential in Sec. III. The numerical results are given in Sec. IV, followed by conclusion and discussion in Sec. V. Some estimates related to the update of the bounds among consecutive SCF steps is given in the Appendix.

II. POLE EXPANSION AND SELECTED INVERSION METHOD

Assume that the Kohn-Sham orbitals are expanded with a set of basis functions $\Phi = [\varphi_1, \dots, \varphi_N]$ and denote by H , S the discretized Kohn-Sham Hamiltonian matrix and overlap matrices, respectively. Without loss of generality, we consider

isolated systems or solids with Γ point sampling strategy of the Brillouin zone, and H, S are real symmetric matrices. The standard method for solving the discretized system is to solve the generalized eigenvalue problem

$$HC = SC\Lambda \quad (2)$$

via direct or iterative methods. Here $\Lambda = \text{diag}[\lambda_1, \dots, \lambda_N]$ is a diagonal matrix containing the Kohn-Sham eigenvalues. The single particle density matrix is then defined as

$$\Gamma = C f_\beta(\Lambda - \mu I) C^T. \quad (3)$$

Here $\beta = 1/k_B T$ is the inverse temperature, and

$$f_\beta(x) = \frac{1}{1 + e^{\beta x}} \quad (4)$$

is the Fermi-Dirac distribution (spin degeneracy is omitted). The chemical potential μ is chosen to ensure that

$$N_\beta(\mu) = \text{Tr}[f_\beta(\Lambda - \mu I)] = \text{Tr}[S\Gamma] = N_e, \quad (5)$$

where N_e is the number of electrons. The computational complexity for the solution of the generalized eigenvalue problem is typically $\mathcal{O}(N_e^3)$.

Since $f_\beta(\cdot)$ is a smooth function when β is finite, the Fermi operator expansion (FOE) method expands $f_\beta(\cdot)$ using a linear combination of simple functions such as polynomials or rational functions, so that the density matrix can be evaluated by matrix-matrix multiplication or matrix inversion, without computing any eigenvalue or eigenfunction. The recently developed pole expansion and selected inversion (PEXSI) method^{12,13} is one type of FOE method, which expands Γ using a rational approximation as

$$\Gamma \approx \sum_{l=1}^P \Im m \left(\omega_l^p (H - (z_l + \mu)S)^{-1} \right). \quad (6)$$

Here $G_l := (H - (z_l + \mu)S)^{-1}$ defines an inverse matrix (or Green's function) corresponding to the complex shift z_l . The pole expansion is not unique. The expansion in Ref. 12 uses a discretized Cauchy contour integral technique and can accurately approximate the density matrix with only $\mathcal{O}(\log \beta \Delta E)$ terms. Here $\Delta E := \max_{1 \leq i \leq N} |\mu - \lambda_i|$ is the spectral radius of the matrix pencil (H, S) . The number of poles needed in practice is typically 40 ~ 80.

Recently, a new expansion for the Fermi-Dirac function has been developed by Moussa,²⁶ which has the same form as in Eq. (6) but with a different choice of complex shifts $\{z_l\}$ and weights $\{\omega_l\}$. This approach is based on modifying the Zolotarev expansion for the sign function through numerical optimization. We find that this new approach further reduces the number of poles to 10 ~ 30 for approximating the Fermi-Dirac function for a wide range of $\beta \Delta E$. Furthermore, the number of poles only approximately depends on $\beta \Delta E_o$. Here $\Delta E_o := \max_{1 \leq i \leq N} (\mu - \lambda_i)$ is the spectral radius corresponding to the occupied eigenvalues, which can be significantly smaller than the spectral radius of the matrix pencil (H, S) . Due to these advantages, we adopt Moussa's method as the default option for the pole expansion for the density matrix.

However, the optimization based pole expansion approach has two minor drawbacks compared to the contour integral approach. First, the validity of the contour integral approach

only depends on the analytical structure of the function to be approximated in the complex plane, rather than the detailed form of the function. This fact allows us to use exactly the same set of poles to approximate multiple matrix functions simultaneously, such as the energy density matrix and the free energy density matrix used for computing the Pulay force and the electronic entropy, respectively.^{13,18} This property does not hold for optimization-based pole expansion method. Second, the contour integral approach is a semi-analytic approach and can achieve very high accuracy (e.g., the error of the force can be as small as 10^{-9} hartree/bohr when compared to results from the diagonalization method) without suffering from numerical problems. On the other hand, the optimization problem for finding the pole expansion can become increasingly ill-conditioned as the requirement of accuracy increases.

In order to obtain the electron density in the real space, it is not necessary to evaluate the entire density matrix Γ . When the local or semilocal exchange-correlation functionals are used, only the *selected elements* $\{\Gamma_{ij} | H_{ij} \neq 0\}$ are in general needed, even if the off-diagonal elements of Γ decay slowly as in the case of metallic systems. We remark that when a matrix element H_{ij} is accidentally zero (often due to symmetry conditions), such zero elements should also be treated as nonzero elements. According to Eq. (6), we only need to evaluate these selected elements of each Green's function. This can be achieved via the selected inversion method.^{12,14,28} For a (complex) symmetric matrix of the form $A = H - zS$, the selected inversion algorithm first constructs an LDL^T or LU factorization of A . The computational scaling of the selected inversion algorithm is only related to the number of nonzero elements in the L, U factors. The computational complexity is $\mathcal{O}(N)$ for quasi-1D systems, $\mathcal{O}(N^{1.5})$ for quasi-2D systems, and $\mathcal{O}(N^2)$ for 3D bulk systems, thus achieving universal asymptotic improvement over the diagonalization method for systems of all dimensions. It should be noted that the selected inversion algorithm is an *exact* method for computing selected elements of A^{-1} if exact arithmetic is employed, and in practice the only source of error originates from the roundoff error. In particular, the selected inversion algorithm does not rely on any localization property of A^{-1} .

In addition to its favorable asymptotic complexity, the PEXSI method is also inherently more scalable than the standard approach based on matrix diagonalization when it is implemented on a parallel computer. The parallelism in PEXSI exists at two levels. First, the LU factorization and the selected inversion processes associated with different poles are completely independent. Second, each LU factorization and selected inversion can be parallelized. We use the SuperLU_DIST²⁹ package for parallel LU factorization and PSELInv^{14,30} for parallel selected inversion, respectively. In practice, the PEXSI method can harness over 100 000 processors on high performance computers.

III. ROBUST DETERMINATION OF THE CHEMICAL POTENTIAL

Unlike diagonalization type methods, in general the chemical potential μ cannot be read off directly from one evaluation

of the Fermi operator. Typically several iterations are needed to identify the chemical potential in order to satisfy Eq. (5). As discussed in the Introduction, the behavior of the function N_β can depend on whether the system is insulating or metallic and whether the temperature is low or high compared to the magnitude of the gap. Furthermore, when the initial guess is far away from the true chemical potential, the derivative information $N'_\beta(\mu)$ is in general not very useful. Figure 1 illustrates the behavior of $N_\beta(\mu)$ for an insulating and a metallic system. The details of the setup of the systems are given in Sec. IV.

Instead of converging the chemical potential in each SCF iteration, the main idea of this paper is to compute rigorous upper and lower bounds for the chemical potential. The chemical potential is obtained using such bounds through an interpolation procedure and may not necessarily be fully accurate in each SCF iteration. However, as the SCF iteration converges, the size of the search interval characterized by the bounds decreases to zero, and hence the chemical potential also converges to the true solution.

We obtain such bounds through a coarse level and a fine level procedure as follows. At the coarse level, we use an inertia counting procedure previously developed in Ref. 18, which is an inexpensive procedure to compute the zero temperature limit $N_{\beta=\infty}(\mu)$ for a set of values of μ . Since the temperature effect is usually on the order of 100 K, which is orders of magnitude smaller than the size of the search interval which is on the order of Hartree (1 Hartree \approx 315 774 K), such zero temperature information can provide estimates of the upper and lower bounds until the finite temperature effects become non-negligible. Then at the fine level, we use PEXSI to evaluate $N_\beta(\mu)$ for a smaller set of values of μ (the size of this set is denoted by N_{point}), which properly takes the finite temperature effect into account and refines the bounds. The evaluation of the Fermi operator at multiple values of μ also allows us to interpolate the chemical potential as well as the density matrix, so that Eq. (1) is satisfied up to the error of the interpolation procedure. The procedure above is all that is involved in a single SCF iteration, and no further iteration is necessary within the SCF iteration.

At the coarse level and the fine level, the multiple evaluations of $N_\infty(\mu)$ and $N_\beta(\mu)$ can be evaluated in an embarrassingly parallel fashion. Compared to Ref. 18, this adds a third layer of parallelism. This is the key to achieve high efficiency using the PEXSI method. We assume that the total number of processors is denoted by

$$N_{\text{proc}} = N_{\text{sparse}} N_{\text{pole}} N_{\text{point}}, \quad (7)$$

where N_{sparse} is the number of processors for operations on each matrix, such as LU factorization or selected inversion. $N_{\text{pole}} = P$ is the number of poles in the pole expansion. When the number of processors is a multiple of $N_{\text{pole}} N_{\text{point}}$, all poles can be evaluated fully in parallel, and we refer to this case as the full parallelization regime. In this case, the wall clock time of the new method during each SCF iteration is approximately the same as that for one single evaluation of the Fermi operator.

In order to guarantee that μ can converge to the true value of the chemical potential when SCF converges, the upper and lower bounds of the chemical potential must also reduce proportionally with respect to the residual error in the SCF iteration. Our method dynamically updates the interval between consecutive SCF steps, which is rigorously controlled by the magnitude of the change of the Kohn-Sham potential. In particular, when the SCF iteration is close to its convergence, the size of the search interval characterized by the upper and lower bounds is smaller than the finite temperature effect. In such a case, the coarse level inertia counting procedure can be safely skipped, and the wall clock time is precisely the same as that for one single PEXSI evaluation.

A. Coarse level: Inertia counting

While the computation of $N_\beta(\mu)$ requires the evaluation of the Fermi operator, it turns out to be much easier to compute $N_\infty(\mu)$ without diagonalizing the matrix pencil (H, S) . Here the subscript $\beta = \infty$ refers to the zero temperature limit. The method of Ref. 18 uses Sylvester's law of inertia,³¹ which states that the inertia (the number of negative, zero, and positive eigenvalues, respectively) of a real symmetric matrix does not change under a congruent transform. Our strategy is to perform a matrix decomposition of the shifted matrix

$$H - \mu S = LDL^T,$$

where L is unit lower triangular and D is diagonal. Since D is congruent to $H - \mu S$, D has the same inertia as that of $H - \mu S$. Hence, we can obtain $N_\infty(\mu)$ by simply counting the number of negative entries in D . The matrix decomposition can be computed efficiently by using a sparse LDL^T or LU factorization with a symmetric permutation strategy. The same conclusion holds when H is Hermitian, and in this case, one then replaces the LDL^T factorization by the LDL^* factorization, where L^* is the Hermitian conjugate of L . Compared to the evaluation of the Fermi operator using PEXSI, the inertia counting step is fast for a number of reasons: (1) PEXSI requires both the sparse factorization and selected inversion,




FIG. 1. Fermi-Dirac distribution for phosphorene nanoribbon (PNR) 180 atoms (left) and graphene nanoribbon (GRN) 180 atoms (right). The blue dashed line shows the exact number of electrons (900 for PNR and 720 for GRN).

and inertia counting only requires a sparse factorization. (2) PEXSI requires evaluations of P Green's functions to obtain one value of $N_\beta(\mu)$ and inertia counting obtains $N_\infty(\mu)$ with one factorization. (3) PEXSI requires complex arithmetic, and for real matrices, the inertia counting procedure only requires real arithmetic and thus fewer floating point operations. Hence the inertia counting step takes only a fraction of the time by each PEXSI evaluation.

The inertia counting strategy is naturally suited for parallelization. The N_{proc} processors in Eq. (7) can be partitioned into $N_{\text{pole}}N_{\text{point}}$ groups. Starting from an initial guess interval $(\mu_{\text{min}}, \mu_{\text{max}})$, a set of values $\{N_\infty(\mu_g)\}$ can be simultaneously evaluated on a uniform grid

$$\mu_g = \mu_{\text{min}} + \frac{g-1}{N_{\text{pole}}N_{\text{point}}-1}(\mu_{\text{max}} - \mu_{\text{min}}), \quad (8)$$

$$g = 1, \dots, N_{\text{pole}}N_{\text{point}}.$$

The inertia counting procedure can determine that only one of the intervals should contain the chemical potential, and the same procedure can be applied to this refined interval. It is easy to see that the size of the search interval shrinks rapidly as $(N_{\text{pole}}N_{\text{point}}-1)^{-k}$ with respect to the number of iterations k . For example when $N_{\text{pole}}N_{\text{point}} = 40$, 3 iterations reduces the search interval size from 10 hartree to 4 meV.

The effectiveness of the inertia searching procedure relies on that $N_\beta(\mu)$ can be well approximated by $N_\infty(\mu)$. This approximation is clearly valid when the search interval is much larger than $1/\beta = k_B T$. When the search interval is comparable to $k_B T$, the difference of the two quantities becomes noticeable, and care must be taken so that the search interval does not become too small to leave the true chemical potential outside. In Ref. 18, we employ an interpolation procedure to estimate $N_\beta(\mu)$ from $N_\infty(\mu)$. The potential drawback of this procedure is that the true chemical potential may not be always included in the search interval. In this case, the subsequent PEXSI step

can fail, and one must go back to the previous inertia counting stage with an expanded search interval.

In this work, we use the information $N_\infty(\mu_g)$ only to calculate the upper and lower bounds for $N_\beta(\mu_g)$, which guarantees that the true chemical potential is always contained in the search interval, up to the error controlled by a single parameter τ_β . More specifically, since the Fermi-Dirac function $f_\beta(x)$ is a non-increasing function and rapidly approaches 1 when $x < 0$ and 0 when $x > 0$, we can select a number τ_β so that we can approximate

$$f_\beta(x) \approx \begin{cases} 1, & x \leq -\tau_\beta, \\ 0, & x \geq \tau_\beta. \end{cases} \quad (9)$$

τ_β is a tunable parameter but is not system dependent. In practice we find that setting $\tau_\beta = 3/\beta = 3k_B T$ is a sufficiently conservative value for the robustness of our method. With this controlled approximation, we have

$$N_\beta(\mu - \tau_\beta) \leq N_\infty(\mu) \leq N_\beta(\mu + \tau_\beta). \quad (10)$$

Hence each evaluation of N_∞ provides an upper and a lower bound for N_β at two other energy points according to (10). This provides an estimate of the chemical potential on each grid point of the uniform grid in the inertia counting. If $\mu - \tau_\beta$ exceeds the lower bound or $\mu + \tau_\beta$ exceeds the upper bound, we can just take the lower bound to be 0 or the upper bound to be the matrix size, respectively. We set μ_{min} to be the largest μ so that the upper bound is below N_e , and μ_{max} is the smallest μ so that the lower bound is above N_e . Figure 2 illustrates the refinement procedure of the bounds for the chemical potential for the PNR 180 system.

The inertia counting stops when $\mu_{\text{max}} - \mu_{\text{min}}$ is below certain tolerance denoted by $\tau_{\text{inertia}}^\mu$ or when the lower and upper bounds cannot be further refined, whichever is satisfied first. In particular, if the size of the initial search interval is smaller than the given tolerance, the inertia counting procedure should be skipped directly. As will be discussed later,

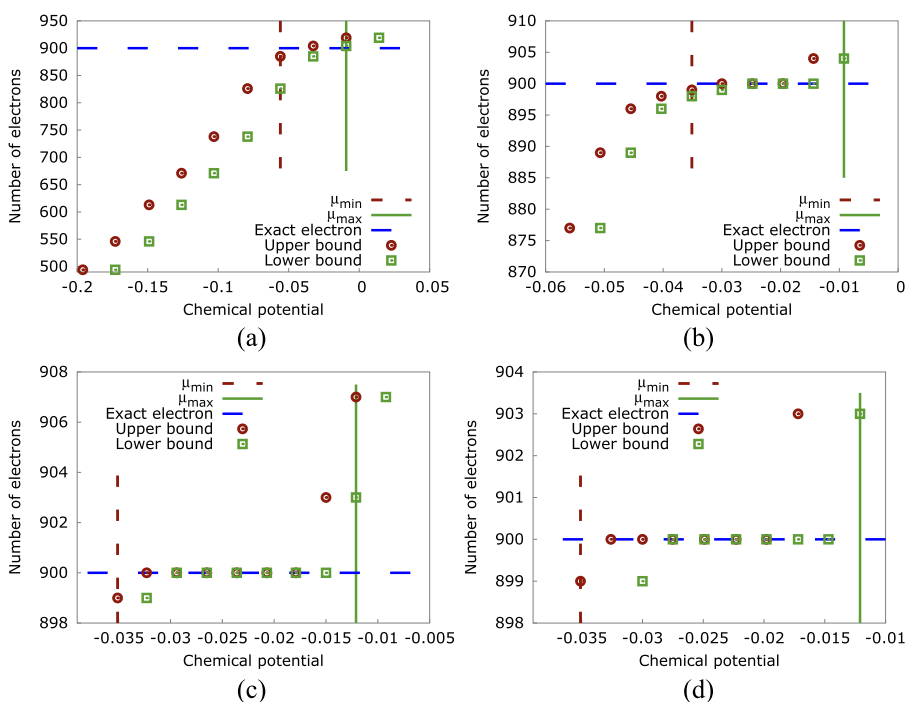


FIG. 2. Refinement of the bounds for the chemical potential in 4 inertia counting steps for the PNR 180 system. For illustration purpose, we evaluate the inertia counting on 10 points. In step 4, the inertia counting procedure stops because the upper and lower bounds cannot be further refined. (a) Inertia counting step 1. (b) Inertia counting step 2. (c) Inertia counting step 3. (d) Inertia counting step 4.

the capability of skipping the inertia counting procedure is important for the self-correction of the search interval for the chemical potential. The output of the inertia counting procedure is an updated search interval, still denoted by (μ_{\min}, μ_{\max}) ready for the evaluation of the Fermi operator.

B. Fine level: Multiple point evaluation of the Fermi operator

After the inertia counting step converges, we refine the upper and lower bounds of the chemical potential by evaluating the Fermi operator on multiple points using PEXSI. This step also gives the density matrix Γ . Our target is to minimize the *wall clock time* with the help of parallelization. Hence motivated from the inertia counting procedure that performs multiple matrix factorizations simultaneously, we can also perform N_{point} PEXSI evaluations simultaneously. In the current context, the method in Ref. 18 can be regarded as choosing $N_{\text{point}} = 1$.

Starting from the search interval (μ_{\min}, μ_{\max}) from the output of the inertia counting procedure, this interval is uniformly divided into $N_{\text{point}} - 1$ intervals as

$$\mu_g = \mu_{\min} + \frac{g-1}{N_{\text{point}}-1}(\mu_{\max} - \mu_{\min}), \quad g = 1, \dots, N_{\text{point}}, \quad (11)$$

and the density matrix $\Gamma(\mu_g)$ and hence number of electrons $N_{\beta}(\mu_g)$ can be evaluated simultaneously for all points. Unlike inertia counting, $N_{\beta}(\mu_g)$ is evaluated accurately with the finite temperature effect correctly taken into account, and the only error is from the pole expansion. This naturally updates the upper and lower bounds of the chemical potential μ_{\max}, μ_{\min} , respectively. If $|N_{\beta}(\mu_g) - N_e|$ is already smaller than the given

tolerance τ^{N_e} for some g , then the PEXSI step converges. We set $\mu = \mu_g$, $\Gamma = \Gamma(\mu_g)$, and we may proceed to the next SCF iteration. If the condition is satisfied for multiple g , we simply choose the first μ_g that satisfies the convergence criterion. If the convergence criterion is not met for any μ_g , we construct a piecewise interpolation polynomial $\tilde{N}_{\beta}(\mu)$ that is monotonically non-decreasing and satisfies

$$\tilde{N}_{\beta}(\mu_g) = N_{\beta}(\mu_g), \quad g = 1, \dots, N_{\text{point}}. \quad (12)$$

Then the chemical potential is identified by solving $\tilde{N}_{\beta}(\mu) = N_e$. Such an interpolation can be constructed using, e.g., the monotone cubic spline interpolation.³² In practice, we find that the following linear interpolation procedure is simpler and works almost equally efficiently: We identify an interval (μ_{g^*}, μ_{g^*+1}) so that $N_{\beta}(\mu_{g^*}) < N_e < N_{\beta}(\mu_{g^*+1})$. Then the chemical potential is found by solving the linear equation

$$N_e = N_{\beta}(\mu_{g^*}^*) + \frac{\mu - \mu_{g^*}^*}{\mu_{g^*+1} - \mu_{g^*}^*} (N_{\beta}(\mu_{g^*+1}) - N_{\beta}(\mu_{g^*}^*)). \quad (13)$$

Once μ is obtained, the density matrix is linearly mixed similarly as

$$\Gamma = \Gamma(\mu_{g^*}^*) + \frac{\mu - \mu_{g^*}^*}{\mu_{g^*+1} - \mu_{g^*}^*} (\Gamma(\mu_{g^*+1}) - \Gamma(\mu_{g^*}^*)). \quad (14)$$

Note that the interpolation procedure does not guarantee that μ satisfies the condition (5). However, it will ensure that the search interval is reduced at least by a factor $(N_{\text{point}} - 1)^{-1}$, and hence with k steps of iteration, the convergence rate of the chemical potential is at least $(N_{\text{point}} - 1)^{-k}$. At each step of the iteration, the true chemical potential is always contained in the search interval. Hence the refinement is more robust than Newton type methods, and it does not need to fold back

Algorithm 1. Robust determination of the chemical potential.

Input:

$\tau_{\text{inertia}}^{\mu}, \tau^{N_e}, N_{\text{proc}} = N_{\text{sparse}} N_{\text{point}} N_{\text{pole}}$ (assuming full parallelization),
 (μ_{\min}, μ_{\max}) as the initial search interval.

Output:

Converged density matrix Γ and chemical potential μ .

- 1: **while** SCF has not converged **do**
 - 2: Construct the projected Hamiltonian matrix H and overlap matrix S
 - 3: **while** Inertia counting has not converged **do**
 - 4: **for** $g = 1$ to $N_{\text{pole}} N_{\text{point}}$ **do**
 - 5: Evaluate $N_{\infty}(\mu_g)$ for the processor group associated with μ_g from Eq. (8).
 - 6: Construct the upper and lower bounds for $N_{\beta}(\mu_g)$ for each point μ_g .
 - 7: Update the search interval (μ_{\min}, μ_{\max}) .
 - 8: **for** $g = 1$ to N_{point} **do**
 - 9: **for** $l = 1$ to N_{pole} **do**
 - 10: Evaluate the pole $(H - (z_l + \mu)S)^{-1}$ for the processor group associated with l
 - 11: Evaluate μ and Γ for processor group associated with μ_g from Eq. (11) if needed.
 - 12: Perform linear interpolation for μ and Γ according to Eqs. (13) and (14).
 - 13: Evaluate the density and the new potential. Compute the difference of the potential ΔV .
 - 14: Update the search interval $(\mu_{\min}, \mu_{\max}) \leftarrow (\mu_{\min} + \Delta V_{\min}, \mu_{\max} + \Delta V_{\max})$.
-
-

to the inertia counting stage. The extra flexibility in choosing N_{point} means that faster convergence can be achieved when a larger amount of computational resource is available. We also remark that the two end points of the search interval in Eq. (11) are always the lower and upper bounds for μ and hence can often be discarded in practical calculations. This increases the efficiency especially when N_{point} is small. In practice, we find that even the extreme case with $N_{\text{point}} = 2$ with dropped end points is already very robust, and this is the default choice in our implementation.

C. Dynamical update of the search interval

After the inertia counting and the multiple point evaluation step, the search interval (μ_{\min}, μ_{\max}) already provides very accurate information of the upper and lower bounds for the chemical potential. Here we demonstrate how to reuse this information in the following SCF iteration, while maintaining rigorous bounds for the chemical potential of the updated Hamiltonian matrix.

For KSDFT calculations with local and semi-local exchange-correlation functionals, the change of the Kohn-Sham Hamiltonian during the SCF iterations is always given by a local potential denoted by $\Delta V_{\text{SCF}}(\mathbf{r})$. Define

$$\Delta V_{\min} = \inf_{\mathbf{r}} \Delta V_{\text{SCF}}(\mathbf{r}), \quad \Delta V_{\max} = \sup_{\mathbf{r}} \Delta V_{\text{eff}}(\mathbf{r}). \quad (15)$$

Then following the derivation in the Appendix, the chemical potential must be contained in the interval $(\mu_{\min} + \Delta V_{\min}, \mu_{\max} + \Delta V_{\max})$. This interval provides an upper and a lower bound for the chemical potential for the new matrix pencil (H, S) and can be used as the initial search interval in the next SCF iteration.

Note that when $(\mu_{\max} - \mu_{\min}) + (\Delta V_{\max} - \Delta V_{\min})$ is smaller than the stopping criterion of the inertia counting τ_{β} , the inertia counting step will be skipped automatically and only the PEXSI step will be executed. When the SCF iteration is close to its convergence, $\Delta V_{\max} - \Delta V_{\min}$ becomes small, and the search interval will be systematically reduced. A pseudocode for our method is summarized in Algorithm 1.

IV. NUMERICAL RESULTS

We have implemented the new method in the PEXSI package,¹⁵ which is a standalone software package for evaluating the density matrix for a given matrix pencil (H, S) . For electronic structure calculations, the PEXSI package is also integrated into the recently developed ‘‘Electronic Structure

Infrastructure’’ (ELSI) software package.²⁷ In order to demonstrate that the performance of the new method in the context of SCF iterations for real materials, we test its performance in the DGDFT (Discontinuous Galerkin Density Functional Theory) software package,^{20,21} which is a massively parallel electronic structure software package for large scale DFT calculations using adaptive local basis (ALB) functions.

Our test systems include a semiconducting phosphorene nanoribbon (PNR) with 180 atoms and two metallic graphene nanoribbon (GRN) systems with 180 and 6480 atoms, respectively. As shown in Fig. 3, the GRN system is a metallic system and PNR is a semi-conductor system with a bandgap of 0.48 eV. We set the kinetic energy cutoff to be 40 hartree and use 15 adaptive local basis functions per atom. The size of the Hamiltonian matrix represented in the adaptive local basis set for PNR 180 and GRN 180 is 2700 and is 92 700 for GRN 6480, respectively. All calculations are carried out on the Edison systems at the National Energy Research Scientific Computing Center (NERSC). Each node consists of two Intel ‘‘Ivy Bridge’’ processors with 24 cores in total and 64 gigabyte (GB) of memory. Our implementation only uses MPI. The number of cores is equal to the number of MPI ranks used in the simulation.

We first demonstrate the accuracy of the pole expansion using two approaches: the contour integral approach¹² and the optimization approach by Moussa.²⁶ The test is performed using a fully converged Hamiltonian matrix from the GRN 180 system, benchmarked against results from diagonalization methods using the pdsyevd routine in ScaLAPACK. The accuracy of the pole expansion is measured by the error of the energy denoted by ΔE and the maximum error of the force denoted by ΔF , respectively. As shown in Fig. 4, in the contour integral approach, both errors decay exponentially with respect to the number of poles. When 80 poles are used, the pole expansion is highly accurate, as $\Delta E \approx 10^{-7}$ hartree and $\Delta F \approx 10^{-9}$ hartree/bohr, respectively. However, in practical KSDFT calculations, the accuracy reached by expansion with 60 poles is already sufficiently accurate. Such accuracy can be achieved by using 20 poles with the optimization method, which reduces the number of poles by a factor of 3. Nonetheless, when a larger number of poles are requested, the optimization method solves an increasing, more ill-conditioned problem, and the error stops decreasing after 30 poles are used.

As discussed in Sec. III, the method in Ref. 18 uses a hybrid Newton type method that requires multiple evaluations of the Fermi operator performed sequentially to guarantee




FIG. 3. The total densities of states (DOS) of (a) PNR and (b) GRN 180 atoms, respectively. The fermi levels are marked by the green dashed line.

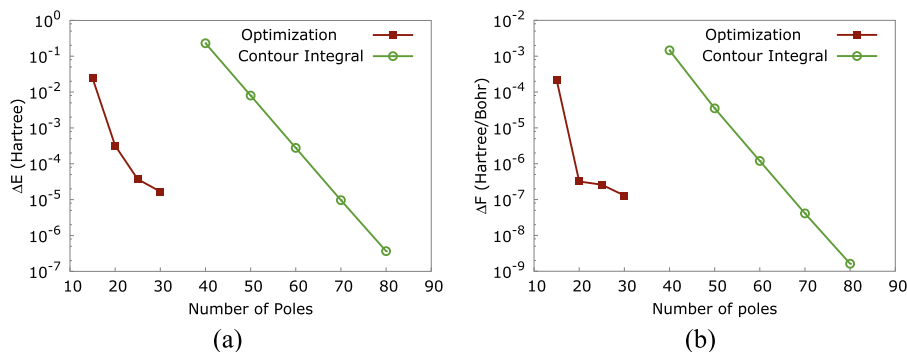


FIG. 4. Error of the (a) energy and (b) force of the PEXSI method with respect to the number of poles for the graphene 180 system.

convergence of the chemical potential at each step of the SCF iteration. This is referred to as the “PEXSI-old” method, as opposed to the “PEXSI-new” method in this paper. We then examine the convergence of the SCF iterations using the PEXSI-new, PEXSI-old, and ScaLAPACK methods, respectively. The residual error of the potential is defined as $\|V^{\text{out}} - V^{\text{in}}\|/\|V^{\text{in}}\|$, where V^{in} and V^{out} correspond to the input and output Kohn-Sham potential in each SCF iteration, respectively. According to Fig. 5, the decay rate of the residual error is nearly the same for all three methods. In our tests, the convergence criterion for the residual error is set to 10^{-7} for GRN 180 and 10^{-6} for PNR 180. When the SCF iteration reaches convergence, we find that the error of the total energy per atom between PEXSI-new and ScaLAPACK is 2.8×10^{-6} hartree/atom for PNR and is 5.5×10^{-8} hartree/atom for GRN, respectively. Similarly, the maximum error of the force between PEXSI-new and ScaLAPACK is 1.48×10^{-5} hartree/bohrs for PNR and is 2.31×10^{-6} hartree/bohrs for GRN, respectively.

In order to demonstrate the impact of the choice of the parameter N_{point} , we study the performance for the GRN180 system using $N_{\text{point}} = 1, 2, 4$, respectively. The convergence criteria are set to 10^{-7} , and the same number of cores (576) is used in all cases. It takes 91, 33, 33 SCF steps to reach convergence and each PEXSI iteration takes 0.96, 1.71, 3.21 seconds for $N_{\text{point}} = 1, 2, 4$, respectively. Setting N_{point} to 1 is not efficient in terms of the number of SCF steps nor the total wall clock time. The choices $N_{\text{point}} = 2$ and $N_{\text{point}} = 4$ have the same convergence rate in terms of SCF steps. Hence in this test, setting $N_{\text{point}} = 2$ is about $2\times$ faster than $N_{\text{point}} = 1$ and 4. We find that this is a typical behavior and recommend $N_{\text{point}} = 2$ to be used as the default choice.

In order to demonstrate the performance of PEXSI for a relatively large metallic system, we further consider the GRN

6480 atom system. Figure 6(a) shows the number of iterations at the coarse level (inertia counting) and at the fine level (evaluation of the Fermi operator) in the PEXSI-new and PEXSI-old methods, respectively. The PEXSI-new method uses only one iteration at the fine level by construction, while the PEXSI-old method uses 2 – 3 fine level steps during the SCF iteration. We also find that PEXSI-new involves on average less number of coarse level inertia counting steps as well. Figure 6(b) reports the average wall clock time per SCF iteration for PEXSI-new, PEXSI-old, and diagonalization using ScaLAPACK, all using 5184 cores. We also report the timing for LU factorization and selected inversion separately for the evaluation of the Fermi operator at the fine level. We can clearly observe that despite the inertia counting steps at the coarse level may require multiple iterations, it is much less costly compared to the fine level evaluation step. Compared to PEXSI-old, PEXSI-new reduces the wall clock time by around a factor of 2 due to the reduced number of fine level evaluations. We also remark that while the wall clock time using ScaLAPACK cannot be improved with further increase of the number of cores, PEXSI can scale up to 51 840 cores, where the average wall clock time per SCF from PEXSI-new and PEXSI-old becomes 31 s and 57 s, respectively.

In order to further demonstrate the effectiveness of the strategy of on-the-fly convergence of the chemical potential, we apply the PEXSI-new method to perform *ab initio* molecular dynamics simulation for the PNR 180 system in the NVE ensemble. We use the Verlet method to produce a trajectory of length 250 fs. We set the initial temperature to be 300 K. Figure 7 shows the potential energy E_{pot} and the total energy E_{tot} , as well as the drift of the total energy along the simulation, which is defined as $E_{\text{drift}}(t) = \frac{E_{\text{tot}}(t) - E_{\text{tot}}(0)}{E_{\text{tot}}(0)}$. We find that the use of the PEXSI-new method leads to a very small drift less than 1.7×10^{-6} .

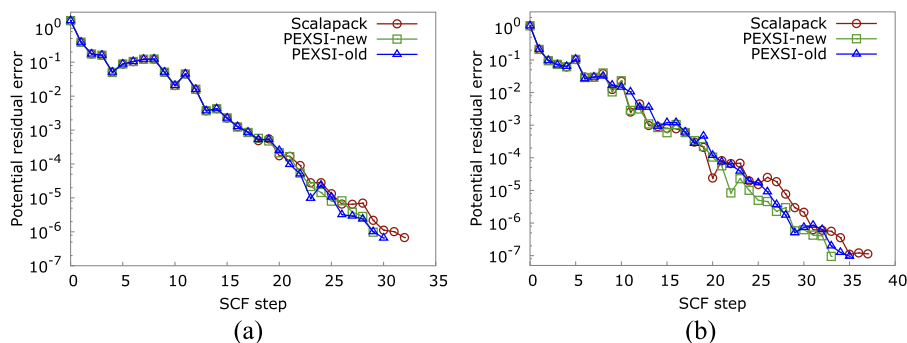


FIG. 5. SCF convergence with ScaLAPACK and PEXSI-old and PEXSI-new along the SCF steps for (a) PNR 180 and (b) GRN 180 systems.

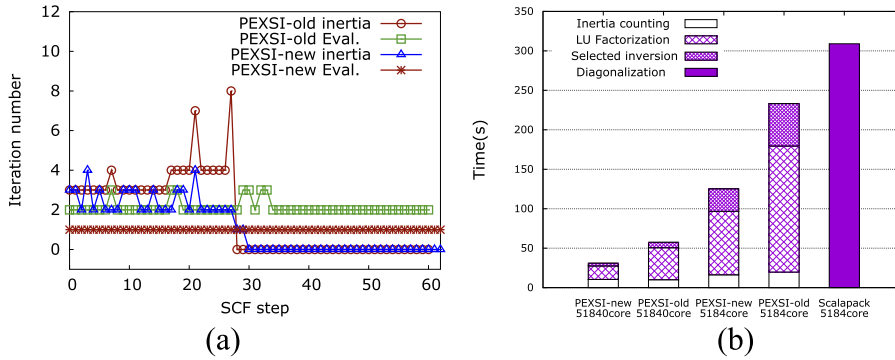


FIG. 6. Comparison of the performance of PEXSI-new and PEXSI-old, and diagonalization for the GRN 6480 system. (a) Number of inertia counting step and evaluation of the Fermi operator steps. (b) The timing for PEXSI-new, PEXSI-old, and diagonalization per SCF step.




FIG. 7. (a) The potential energy, total energy, and (b) drift of the total energy using the PEXSI-new method.

V. CONCLUSION

In this paper, we developed a robust and efficient approach for finding the chemical potential in the pole expansion and selected inversion (PEXSI) approach for solving KSDFT. The main idea of the new method is not to find the exact chemical potential at each SCF iteration but to dynamically update the upper and lower bounds for the true chemical potential, so that the accuracy of the chemical potential is comparable to that of the residual error in the SCF iteration. In particular, when the SCF converges, the chemical potential converges as well. The new method reduces the number of tunable parameters and is more robust. In the regime of full parallelization, the wall clock time in each SCF iteration is always approximately the same as only one evaluation of the Fermi operator. This reduces the absolute value as well as the uncertainty of the wall clock time when the PEXSI method is used for solving KSDFT. We demonstrated the efficiency of the new method using insulating and metallic systems as examples, and the new method is available in the PEXSI software package. Finally, our approach is not restricted to the PEXSI method and can be applied to other Fermi operator expansion (FOE) type of methods as well.

ACKNOWLEDGMENTS

This work was partially supported by the National Science Foundation under Grant No. 1450372 (W.J. and L.L.), by the National Science Foundation under Grant No. DMS-1652330, the Scientific Discovery through Advanced Computing (SciDAC) program, and the Center for Applied Mathematics for Energy Research Applications (CAMERA) funded by U.S. Department of Energy (L.L.). This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported

by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. We thank Jianfeng Lu, Meiyue Shao, Chao Yang, and Victor Yu for useful discussions.

APPENDIX: EIGENVALUE ESTIMATES

Let $\Delta V_{\text{SCF}}(\mathbf{r})$ be the difference of the Kohn-Sham potential in the real space between two consecutive SCF iterations, and we assume that the basis set denoted by Φ does not change during the SCF iterations. The corresponding matrix pencils are (H, S) and (\tilde{H}, S) , respectively. The difference of the Hamiltonian matrices is

$$\Delta H = \tilde{H} - H = \Phi^T \Delta V_{\text{SCF}} \Phi, \quad (\text{A1})$$

and the overlap matrix $S = \Phi^T \Phi$ remains the same. Denote by λ_k (respectively, $\tilde{\lambda}_k$) the k th smallest eigenvalue of H (respectively, \tilde{H}). Then the Courant-Fisher minimax theorem³³ indicates that

$$\tilde{\lambda}_k = \min_{\dim \mathcal{S}=k} \max_{0 \neq u \in \mathcal{S}} \left(\frac{u^T \tilde{H} u}{u^T S u} \right), \quad (\text{A2})$$

where the minimization is over all subspaces \mathcal{S} with dimension k . Hence

$$\begin{aligned} \tilde{\lambda}_k &\leq \min_{\dim \mathcal{S}=k} \left[\max_{0 \neq u \in \mathcal{S}} \left(\frac{u^T H u}{u^T S u} \right) + \max_{0 \neq u \in \mathcal{S}} \left(\frac{u^T \Delta H u}{u^T S u} \right) \right] \\ &\leq \min_{\dim \mathcal{S}=k} \max_{0 \neq u \in \mathcal{S}} \left(\frac{u^T H u}{u^T S u} \right) + \max_{0 \neq u} \left(\frac{u^T \Delta H u}{u^T S u} \right) \\ &\leq \lambda_k + \Delta V_{\text{max}}. \end{aligned} \quad (\text{A3})$$

Here we have used the Courant-Fisher minimax theorem for λ_k as well as the estimate of the largest eigenvalue of the

matrix pencil $(\Delta H, S)$ in terms of ΔV_{SCF} . Similarly calculation shows

$$\begin{aligned} \tilde{\lambda}_k &\geq \min_{\dim S=k} \max_{0 \neq u \in S} \left(\frac{u^T H u}{u^T S u} \right) + \min_{0 \neq u} \left(\frac{u^T \Delta H u}{u^T S u} \right) \\ &\geq \lambda_k + \Delta V_{\min}. \end{aligned} \quad (\text{A4})$$

Again use the fact that the Fermi-Dirac function $f_\beta(\epsilon)$ is a non-increasing function, we have

$$\begin{aligned} \tilde{N}_\beta(\mu_{\min} + \Delta V_{\min}) &:= \text{Tr}[f(\tilde{\Lambda} - (\mu_{\min} + \Delta V_{\min})I)] \\ &\leq \text{Tr}[f(\Lambda - \mu_{\min}I)] = N_\beta(\mu_{\min}) \leq N_e, \end{aligned} \quad (\text{A5})$$

where the last inequality comes from that μ_{\min} is a lower bound for the chemical potential associated with the matrix pencil (H, S) . Similarly

$$\tilde{N}_\beta(\mu_{\max} + \Delta V_{\max}) \geq N_\beta(\mu_{\max}) \geq N_e. \quad (\text{A6})$$

The inequalities (A5) and (A6) establish that the chemical potential associated with the matrix pencil (\tilde{H}, S) is contained in the updated interval $(\mu_{\min} + \Delta V_{\min}, \mu_{\max} + \Delta V_{\max})$.

- ¹D. R. Bowler, T. Miyazaki, and M. J. Gillan, "Recent progress in linear scaling *ab initio* electronic structure techniques," *J. Phys.: Condens. Matter* **14**, 2781–2799 (2002).
- ²J. L. Fattebert and J. Bernholc, "Towards grid-based $O(N)$ density-functional theory methods: Optimized nonorthogonal orbitals and multigrid acceleration," *Phys. Rev. B* **62**, 1713–1722 (2000).
- ³N. D. Hine, P. D. Haynes, A. A. Mostofi, C. K. Skylaris, and M. C. Payne, "Linear-scaling density-functional theory with tens of thousands of atoms: Expanding the scope and scale of calculations with ONETEP," *Comput. Phys. Commun.* **180**, 1041–1053 (2009).
- ⁴W. Yang, "Direct calculation of electron density in density-functional theory," *Phys. Rev. Lett.* **66**, 1438–1441 (1991).
- ⁵X.-P. Li, R. W. Nunes, and D. Vanderbilt, "Density-matrix electronic-structure method with linear system-size scaling," *Phys. Rev. B* **47**, 10891–10894 (1993).
- ⁶R. McWeeny, "Some recent advances in density matrix theory," *Rev. Mod. Phys.* **32**, 335–369 (1960).
- ⁷S. Goedecker, "Linear scaling electronic structure methods," *Rev. Mod. Phys.* **71**, 1085–1123 (1999).
- ⁸D. R. Bowler and T. Miyazaki, " $O(N)$ methods in electronic structure calculations," *Rep. Prog. Phys.* **75**, 036503 (2012).
- ⁹W. Kohn, "Density functional and density matrix method scaling linearly with the number of atoms," *Phys. Rev. Lett.* **76**, 3168–3171 (1996).
- ¹⁰E. Prodan and W. Kohn, "Nearsightedness of electronic matter," *Proc. Natl. Acad. Sci. U. S. A.* **102**, 11635–11638 (2005).
- ¹¹S. Goedecker, "Integral representation of the Fermi distribution and its applications in electronic-structure calculations," *Phys. Rev. B* **48**, 17573 (1993).
- ¹²L. Lin, J. Lu, L. Ying, and W. E., "Pole-based approximation of the Fermi-Dirac function," *Chin. Ann. Math., Ser. B* **30**, 729 (2009).

- ¹³L. Lin, M. Chen, C. Yang, and L. He, "Accelerating atomic orbital-based electronic structure calculation via pole expansion and selected inversion," *J. Phys.: Condens. Matter* **25**, 295501 (2013).
- ¹⁴M. Jacquelin, L. Lin, and C. Yang, "PSELInv—A distributed memory parallel algorithm for selected inversion: The symmetric case," *ACM Trans. Math. Software* **43**, 21 (2016).
- ¹⁵See <http://www.pexsi.org> for distributed in bsd license.
- ¹⁶S. Mohr, L. E. Ratcliff, P. Boulanger, L. Genovese, D. Caliste, T. Deutsch, and S. Goedecker, "Daubechies wavelets for linear scaling density functional theory," *J. Chem. Phys.* **140**, 204110 (2014).
- ¹⁷J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, "QUICKSTEP: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach," *Comput. Phys. Commun.* **167**, 103 (2005).
- ¹⁸L. Lin, A. García, G. Huhs, and C. Yang, "SIESTA-PEXSI: Massively parallel method for efficient and accurate *ab initio* materials simulation without matrix diagonalization," *J. Phys.: Condens. Matter* **26**, 305503 (2014).
- ¹⁹J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal, "The SIESTA method for *ab initio* order-N materials simulation," *J. Phys.: Condens. Matter* **14**, 2745–2779 (2002).
- ²⁰L. Lin, J. Lu, L. Ying, and W. E., "Adaptive local basis set for Kohn-Sham density functional theory in a discontinuous Galerkin framework I: Total energy calculation," *J. Comput. Phys.* **231**, 2140–2154 (2012).
- ²¹W. Hu, L. Lin, and C. Yang, "DGDF: A massively parallel method for large scale density functional theory calculations," *J. Chem. Phys.* **143**, 124110 (2015).
- ²²V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, "*Ab initio* molecular simulations with numeric atom-centered orbitals," *Comput. Phys. Commun.* **180**, 2175–2196 (2009).
- ²³See (www.quantumwise.com) for Atomistix toolkit, quantumwise a/s.
- ²⁴W. Hu, L. Lin, C. Yang, and J. Yang, "Electronic structure of large-scale graphene nanoflakes," *J. Chem. Phys.* **141**, 214704 (2014).
- ²⁵W. Hu, L. Lin, C. Yang, J. Dai, and J. Yang, "Edge-modified phosphorene nanoflake heterojunctions as highly efficient solar cells," *Nano Lett.* **16**, 1675 (2016).
- ²⁶J. E. Moussa, "Minimax rational approximation of the Fermi-Dirac distribution," *J. Chem. Phys.* **145**, 164108 (2016).
- ²⁷See <http://www.elsi-interchange.org/> for information about ELSI.
- ²⁸L. Lin, C. Yang, J. Meza, J. Lu, L. Ying, and W. E., "SelInv—An algorithm for selected inversion of a sparse symmetric matrix," *ACM Trans. Math. Software* **37**, 1 (2011).
- ²⁹X. Li and J. Demmel, "SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems," *ACM Trans. Math. Software* **29**, 110 (2003).
- ³⁰M. Jacquelin, L. Lin, N. Wichmann, and C. Yang, "Enhancing the scalability and load balancing of the parallel selected inversion algorithm via tree-based asynchronous communication," in *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)* (IEEE, 2016), pp. 192–201.
- ³¹J. J. Sylvester, "A demonstration of the theorem that every homogeneous quadratic polynomial is reducible by real orthogonal substitutions to the form of a sum of positive and negative squares," *Philos. Mag.* **4**, 138–142 (1852).
- ³²G. Wolberg and I. Alfy, "Monotonic cubic spline interpolation," in *1999 Proceedings of the International Conference on Computer Graphics* (IEEE, 1999), pp. 188–195.
- ³³G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed. (Johns Hopkins University Press, Baltimore, 2013).