

# SELF-CONSISTENT-FIELD CALCULATIONS USING CHEBYSHEV-FILTERED SUBSPACE ITERATION \*

YUNKAI ZHOU <sup>†</sup>, YOUSEF SAAD <sup>†</sup>, MURILO L. TIAGO <sup>‡</sup>, AND JAMES R.  
CHELIKOWSKY <sup>‡</sup>

**Abstract.** The power of Density Functional Theory is often limited by the high computational demand in solving an eigenvalue problem at each Self-Consistent-Field (SCF) iteration. The method presented in this paper replaces the explicit eigenvalue calculations by an approximation of the wanted invariant subspace, obtained with the help of well-selected Chebyshev polynomial filters. In this approach, only the initial SCF iteration requires solving an eigenvalue problem, in order to provide a good initial subspace. In the remaining SCF iterations, no iterative eigensolvers are involved. Instead, Chebyshev polynomials are used to refine the subspace. The subspace iteration at each step is easily five to ten times faster than solving a corresponding eigenproblem by the most efficient eigen-algorithms. Moreover, the subspace iteration reaches self-consistency within roughly the same number of steps as an eigensolver-based approach. This results in a significantly faster SCF iteration.

**Key words.** Density functional theory, Self-consistent-field, Chebyshev Polynomial filter, subspace iteration, eigenproblem, real-space pseudopotential.

**1. Introduction.** Since its formulation, density-functional theory (DFT) [18, 21] has been recognized as a major achievement in the development of quantum many-body theories. Its basis lies in describing the ground state of a many-electron system solely in terms of the charge density function. As a result, the task of solving the *Schrödinger equation* for a many-electron system is replaced by the immensely simpler one of solving a non-linear, self-consistent eigenvalue problem: the *Kohn-Sham equations*. During the last several decades, DFT has been successfully applied to a range of problems in condensed matter physics, material sciences, chemistry and biology [29, 13]. In typical numerical implementations of DFT, the most time-consuming part is spent in computing the self-consistent solution of the Kohn-Sham equations. Because of the high computational demand of matrix diagonalizations when the number of wanted eigenvalues as well as the matrix dimension are large, applying DFT to very large system (e.g., molecules containing thousands of atoms) still remains a highly challenging problem.

Many researchers have considered avoiding matrix diagonalization or reducing its cost. Typical examples of alternatives that have been explored include the use of the Conjugate Gradient (CG) method which directly minimizes the total energy (or Rayleigh-quotient) [30], the Car-Parrinello molecular dynamics method [7], and the DIIS variants [31, 32, 42, 22] which minimize the residual vectors instead of Rayleigh-quotients.

It has long been realized that full diagonalization is too expensive for large problems and so iterative eigensolvers which only compute the wanted eigenpairs have been utilized, with varying degree of success. In the early days of the development of DFT computational codes, the linear *Subspace Iteration* algorithm was used. For example, the 1988 paper by Martins and Cohen [26], see also [27], mentioned the

---

\*Work supported by DOE under grants DE-FG02-03ER25585, DE-FG02-03ER15491, by NSF grant ITR-0428774, and by the Minnesota Supercomputing Institute.

<sup>†</sup>Department of Computer Science & Engineering, University of Minnesota, Minneapolis, MN 55455, USA. (zhou@msi.umn.edu, saad@cs.umn.edu).

<sup>‡</sup>Institute for Computational Engineering and Sciences, University Station, University of Texas at Austin, Austin, Texas 78712, USA. (mtiago@ices.utexas.edu, jrc@ices.utexas.edu).

use of a code called Ritzit, initially published in Algol [33], which constituted an ideally suited technique for diagonalization in planewave codes at the time. In [27], the Chebyshev acceleration in Ritzit was replaced by a Jacobi preconditioner. In contrast with [27], the classical Ritzit code (with Chebyshev acceleration) was found to be an excellent partial diagonalization tool in the solution of self-consistent Schrödinger equations which arise in the simulation of electrons in quantum wires [20]. The main difference with our approach is that we do not use subspace iteration for computing eigenvalues and eigenvectors. Instead, our approach can be viewed as a nonlinear subspace iteration in that after each subspace filtering the Hamiltonian is updated.

More recent examples of the use of iterative diagonalization, include the use of multigrid approaches [5, 6, 11, 14, 17, 24] and preconditioned Davidson method [35, 39]. We note that the complexity of the minimization-based methods is similar to that of the iterative eigensolver-based methods. The computational cost of an eigensolver approach is usually dominated by the cost of the matrix-vector products and that of maintaining orthogonality of the basis vectors. In [4] partial reorthogonalized Lanczos without restart is used in order to reduce the orthogonalization cost, at the expense of a higher memory requirement.

In this paper we present a method which avoids solving large eigenvalue problems explicitly. The method utilizes Chebyshev polynomial filtered subspace iteration. In this approach only the initial SCF iteration requires solving an eigenvalue problem, by means of any available efficient eigensolver. This step is used to provide a good initial subspace. Because the subspace dimension is slightly larger than the number of wanted eigenvalues (denoted by  $k_{want}$ ), the method does not require as much memory as standard restarted eigensolvers such as ARPACK (see [38, 25]) and TRLan (Thick - Restart Lanczos) [43, 44]. Moreover, the cost of orthogonalization is much reduced. This is because the new approach only requires a subspace of dimension around  $k_{want}$ , and the orthogonalization is done only once per SCF iteration. In contrast, standard eigensolvers using restart usually require a subspace of dimension  $2k_{want}$  in order to compute  $k_{want}$  eigenpairs efficiently, the orthogonalization cost approximately amounts to orthogonalizing  $2k_{want}k_{restart}$  number of vectors, where  $k_{restart}$  is the number of restarts needed in order to converge the eigenvectors.

Chebyshev polynomial filtering has been utilized in electronic structure calculations (see e.g. [36, 40, 16, 2, 19]), where the focus is on approximating the Fermi-Dirac operator, i.e., Chebyshev polynomials only over interval  $[-1, 1]$  is considered. The approach we take here is different from the existing Chebyshev methods in electronic structure calculations. The fundamental difference is that we exploit the exponential growth property of Chebyshev polynomial outside the  $[-1, 1]$  interval, hence the polynomial degree required is much lower. In our approach, we never map the full spectrum of the Hamiltonian into  $[-1, 1]$ ; instead, we adaptively decide on the unwanted part of the spectrum and map only this part into  $[-1, 1]$  for damping it. Our Chebyshev filtering is based on [45]. Note that the focus in [45] is on eigenvalue computations while here we only use Chebyshev polynomial to filter subspaces. More details are presented in Section 4.

The current approach has been implemented in PARSEC, our real-space DFT code, and it has been observed to be significantly faster than iterative eigensolver-based approaches. As an example, it took less than 2 hours for the sequential implementation of the new method to reach self-consistency for the Silicon cluster  $Si_{525}H_{276}$  on a single SGI 1.3 GHz Madison processor. In contrast, a parallel calculation done with the preceding version of PARSEC around the year 1997, on the CRAY T3E,

took a total of 20 hours [39] *using 48 processors*. Back in 1997 the problem could not be solved in fewer than 48 processors due to memory requirements. The remarkable gains do not come only from a reduced cost of diagonalization but also from exploiting symmetry [41] and from gains made in other parts of the code. Section 5 presents comparisons with methods based on two of the most efficient eigensolvers currently available (ARPACK and TRLan). All three methods utilize the same symmetry operations, therefore the reported speedup is from the Chebyshev filtering approach.

**2. Basics of Self-consistent-field calculation.** In this section we briefly review the self-consistent-field calculation. Here we focus discussion on SCF in DFT calculations, but we note that SCF is also used in Hartree-Fock and other approximations. As mentioned in the introduction, the most time-consuming part of an SCF calculation is in matrix diagonalization, which consists of computing the self-consistent solutions of the following Kohn-Sham equation:

$$\left[ \frac{-\hbar^2}{2M} \nabla^2 + V_{total}(\rho(r), r) \right] \Psi_i(r) = E_i \Psi_i(r), \quad (2.1)$$

where  $\Psi_i(r)$  is a wave function,  $E_i$  is a Kohn-Sham eigenvalue,  $\hbar$  is the Planck constant,  $M$  is the electron mass. The total potential

$$V_{total}(\rho(r), r) = V_{ion}(r) + V_H(r) + V_{XC}(r) \quad (2.2)$$

includes the ionic potential  $V_{ion}$ , the Hartree potential  $V_H$  and the exchange-correlation potential  $V_{XC}$ . In DFT the total potential depends only on  $\rho(r)$ —the *charge density*. The *charge density* is given by

$$\rho(r) = 2 \sum_{i=1}^{n_{occ}} |\Psi_i(r)|^2, \quad (2.3)$$

where  $n_{occ}$  is the number of occupied states (half the number of valence electrons in the system) and the factor of two comes from spin multiplicity. Equation (2.3) can be easily extended to situations where the highest occupied states have fractional occupancy or when there is an imbalance in the number of electrons for each spin component.

Because the potential (2.2) depends on the charge density (2.3), which in turn depends on eigenfunctions of the the Hamiltonian in Equation (2.1), the eigenvalue problem (2.1) is in fact nonlinear. Self-consistent iterations for solving this problem consist of starting with an initial guess of the charge density  $\rho_0(r)$ , then obtaining a guess for  $V_{total}$  and solving (2.1) for  $\Psi_i(r)$ 's to update  $\rho(r)$  and  $V_{total}$ . Then (2.1) is solved again for the new  $\Psi_i(r)$ 's and the process is carried on until the difference between two consecutive  $V_{total}$ 's is below a certain tolerance (equivalently, the wave functions are close to stationary). Algorithm 2.1 contains a pseudo code for this SCF loop.

We note that, since the charge density does not depend on eigenstates beyond  $n_{occ}$ , the number of eigenvectors needed in *Step 2* of Algorithm 2.1 is limited. Nevertheless, the eigenvalue problem is still very challenging in complex systems (i.e., systems with a very large number of electrons), when the Hamiltonian has large dimension and  $n_{occ}$  is also large.

Our computational code uses a real-space implementation of the above SCF method [10, 35, 39]. In this implementation, wave functions are expressed directly as

ALGORITHM 2.1. *Self-Consistent Iteration:*

- 
1. *Initial guess for  $\rho(r)$ , get  $V_{total}(\rho(r), r)$*
  2. *Solve  $\left[\frac{-\hbar^2}{2M}\nabla^2 + V_{total}(\rho(r), r)\right]\Psi_i(r) = E_i\Psi_i(r)$  for  $\Psi_i(r)$ ,  $i = 1, 2, \dots, n_{occ}$*
  3. *Compute new charge density  $\rho(r) = 2\sum_{i=1}^{n_{occ}} |\Psi_i(r)|^2$*
  4. *Solve for new Hartree potential  $V_H$  from  $\frac{-\hbar^2}{2M}\nabla^2 V_H = 2\pi\rho(r)$*
  5. *Update  $V_{XC}$  and  $V_{ion}$ ; get new  $\tilde{V}_{total}(\rho(r), r) = V_{ion}(r) + V_H(r) + V_{XC}(r)$   
(often followed by a potential-mixing step)*
  6. *If  $\|\tilde{V}_{total} - V_{total}\| < tol$ , stop; Else,  $V_{total} \leftarrow \tilde{V}_{total}$ , goto 2.*
- 

functions of position, and they are required to vanish outside a specified boundary that encloses the physical system (alternatively, periodic boundary conditions can be imposed [1]). The region inside this boundary is discretized by using a regular grid with adjustable spacing between neighbouring points. We use pseudopotentials to describe the interaction between valence electrons and ionic cores (core electrons + nuclei) and solve the SCF problem for valence electrons only. In addition, we make use of symmetry operations in the arrangement of atoms and reduce the sampled region to a smaller one: the ‘‘irreducible wedge’’. Appropriate boundary conditions and the existing symmetry operations are used to expand wave functions from the irreducible wedge to the full volume. For highly symmetric systems, such as the atom clusters analyzed in Section 5, reducing the volume of interest to an irreducible wedge can easily lead to a 10-fold reduction in the computational load [41] than without exploiting symmetry.

**3. Self-consistent-field calculation without explicit eigenvectors.** While the charge density  $\rho(r)$  depends explicitly on the wave functions  $\Psi_i(r)$ , see (2.3), it has been observed that any basis for the subspace spanned by these eigenvectors can be used. After discretization each  $\Psi_i(r)$  becomes an eigenvector (denoted as  $\psi_i$ ) of the discretized Hamiltonian. The charge density is then simply the diagonal of the density matrix

$$P = \Phi\Phi^T, \quad (3.1)$$

in which  $\Phi$  is the matrix with column vectors the  $\{\psi_i\}$ ’s. Entry  $(j, j)$  of this projector  $P$  is equal to the charge density at the mesh-point  $r_j$ . Notice that for any orthonormal matrix  $Q$  of dimension  $s$ ,  $P = (\Phi Q)(\Phi Q)^T$ . Explicit eigenvectors are therefore not needed to calculate the charge density in Algorithm 2.1. Any orthonormal basis of the eigensubspace corresponding to occupied states will give the correct charge density.

Techniques based on this observation have appeared in, e.g., [40, 37, 19, 4, 3]. In particular, the recent article [4] stresses the importance of de-emphasizing eigenvectors in favor of the underlying eigenspace. This observation has also been exploited in the linear scaling approaches to DFT, see e.g. [15] for a survey.

In our approach we progressively refine a subspace by rather low degree Chebyshev polynomials. After self-consistency is reached, this subspace includes the eigensubspace corresponding to occupied states. Explicit eigenvectors can be easily obtained by a *Rayleigh-Ritz refinement* step [28] (called a *subspace rotation* in materials science terminology).

**4. SCF subspace iteration with Chebyshev filtering.** The main idea of the proposed approach is to start with a good initial eigen-basis  $V$  corresponding to occupied states of the initial Hamiltonian  $H_0$ , and then to adaptively improve the subspace by polynomial filtering. That is, at a given self-consistent step, a degree- $m$  polynomial filter  $p_m(t)$  is constructed for the current Hamiltonian  $H$ . Note that the polynomial will be different at each SCF step since  $H$  will change. The goal of the filter is to make the subspace spanned by  $p_m(H)V$  approximate the eigensubspace corresponding to the occupied states of  $H$ . There is no need to make  $p_m(H)V$  approximate the wanted eigensubspace of  $H$  to high accuracy at the intermediate steps. Instead, the filtering is designed so that each new subspace obtained at the end of the self-consistent loop will progressively approximate the wanted eigensubspace of the final Hamiltonian when self-consistency is reached. This can be efficiently achieved by exploiting the Chebyshev polynomials, specifically the fast growth property outside the  $[-1, 1]$  interval. All that is required to obtain a good filter at a given SCF step, is to provide a lower bound and an upper bound of an interval of the spectrum of the current Hamiltonian  $H$  in which we want  $p_m(t)$  to be small. Moreover, the lower bound can be readily obtained from the Ritz values computed from the previous step, and the upper bound can be inexpensively obtained by a very small number of (say, 4 or 5) Lanczos steps [23]. Hence the main cost of the filtering at the  $i$ -th step is in computing  $p_m(H)V$ . This computation is accomplished by exploiting the convenient three-term recurrence property of Chebyshev polynomials.

**4.1. The main algorithm.** The pseudo code of the Chebyshev polynomial filtered subspace iteration is presented in Algorithm 4.1. The purpose of Algorithm 4.1 is to replace the eigenvalue problem at *Step 2* in Algorithm 2.1 by a single Chebyshev filtering step. At the first step of the SCF loop, *Step 2* in Algorithm 2.1 is carried out by an iterative eigensolver like ARPACK or TRLan to provide an initial orthonormal basis  $\Phi = [\psi_1, \dots, \psi_s]$ . In order not to miss occupied eigenstates, a standard practice is to choose  $s > n_{occ}$ , here we follow this practice by fixing an integer  $n_{state}$  which is slightly larger than  $n_{occ}$ , then set  $s = n_{state} + n_{add}$  (where  $n_{add} \leq 10$ ).

ALGORITHM 4.1. *Chebyshev polynomial Filtered Subspace Iteration:*

- 
1. Get the lower bound  $b_{low}$  from previous Ritz values (use the largest one).
  2. Compute the upper bound  $b_{up}$  of the spectrum of the current discretized Hamiltonian  $H$  (call Algorithm 4.3).
  3. Perform Chebyshev filtering to the previous basis  $\Phi$ : (call Algorithm 4.2)  
 $\Phi = \text{Chebyshev\_filter}(\Phi, m, b_{low}, b_{up})$ .
  4. Perform ortho-normalization to  $\Phi$  to make it ortho-normal.
  5. Perform the Rayleigh-Ritz step: Compute  $\hat{H} = \Phi^T H \Phi$ ;  
 Compute the eigendecomposition of  $\hat{H}$ :  $\hat{H}Q = QD$ ;  
 (where  $Q$  contains eigenvectors of the  $\hat{H}$ ,  $D$  contains Ritz values of  $H$ )  
 Refine the basis as:  $\Phi = \Phi Q$ .
- 

For the ortho-normalization step in Algorithm 4.1 (*step 4*) we use the DGKS method [12] which uses iterated classical Gram-Schmidt.

Note that Algorithm 4.1 does not compute eigen-basis of the current Hamiltonian  $H$ , but it computes eigen-decomposition of the projected  $\hat{H}$  of size  $s$ . Even though  $s$  is normally much smaller than the size of  $H$ , the Rayleigh-Ritz refinement (*step 5*) may still be expensive when  $s$  becomes very large (e.g., large complex systems without physical symmetry), since the eigen-decomposition step is of complexity  $O(s^3)$ .

A important feature of the Chebyshev filtered subspace iteration is that the whole *step 5* in Algorithm 4.1 can be waived for huge systems where  $s$  is very large. The theoretic foundation is that  $\Phi Q$  spans the same subspace as  $\Phi$ . In this case we can construct another suitable filter so that the components of  $\Phi$  corresponding to unoccupied states will be filtered out. Rayleigh-Ritz refinement is performed when  $s$  is moderate because it is not expensive, and this refinement makes columns in  $\Phi Q$  approximate the eigenvectors of  $H$  better than  $\Phi$  does. That is, the  $n_{occ}$  columns in  $\Phi Q$  that correspond to the  $n_{occ}$  smallest Ritz values can be readily returned to the main program for the calculation of  $\rho(r)$  to continue the SCF loop. Results in the literature (e.g. [22, 11]) for different approaches also show that *subspace rotation* improves stability and convergence rate. In our approach, the Rayleigh-Ritz step is also helpful because it provides a convenient lower bound for the next Chebyshev filtering.

If the Rayleigh-Ritz step is not performed, the wanted lower bound can still be estimated from the largest Rayleigh-quotient among  $\psi_j^T H \psi_j$ , where  $\psi_j$  ( $j = 1, \dots, s$ ) is the column vector of  $\Phi$ . Thanks to the Courant-Fisher min-max theorem [28, p.206] this will still be a good enough approximation for the filtering to work well.

The next section will show useful properties of Chebyshev polynomials and discuss how to adaptively change these filters.

**4.2. Controlling the Chebyshev polynomial filters.** The well-known Chebyshev polynomials of the first kind are defined by ([28, p.371] [34, p.142])

$$C_k(t) = \begin{cases} \cos(k \cos^{-1}(t)), & -1 \leq t \leq 1, \\ \cosh(k \cosh^{-1}(t)), & |t| > 1. \end{cases}$$

Note that  $C_0(t) = 1, C_1(t) = t$ . The following important 3-term recurrence is easy to derive from properties of  $\cos(t)$  and  $\cosh(t)$ ,

$$C_{k+1}(t) = 2t C_k(t) - C_{k-1}(t), \quad t \in \mathbb{R}. \quad (4.1)$$

The rapid growth property of the Chebyshev polynomial outside  $[-1, 1]$  is discussed in [28]. Figure 4.1 illustrates this property. Here we only plot the  $[-2, 1]$  interval, but note that the farther away from -1 or 1, the larger the magnitude of the real part of the Chebyshev polynomials.

Assume that the full spectrum of  $H$  (denoted as  $\sigma(H)$ ) is contained in  $[a_0, b]$ . As is well-known [33, 28, 34, 45], in order to approximate the eigensubspace associated with the lower end of the spectrum, say  $[a_0, a]$  with  $a_0 < a < b$ , essentially one only needs to map  $[a, b]$  into  $[-1, 1]$ . This can be easily realized by a linear mapping defined as

$$l(t) := \frac{t - c}{e}; \quad c = \frac{a + b}{2}, \quad e = \frac{b - a}{2}$$

where  $c$  denotes the center and  $e$  the half-width.

The Chebyshev iteration, which utilizes the three-term recurrence (4.1) with the goal to dampen values on a given  $[a, b]$  is presented in Algorithm 4.2. Here the formula




FIG. 4.1. Rapid increase outside  $[-1, 1]$  of the  $m$ -th degree Chebyshev polynomial.

derived in [34, p. 223] for the complex Chebyshev iteration is adapted to the real case. The iteration of the algorithm is equivalent to computing

$$Y = p_m(H)X \quad \text{where} \quad p_m(t) = C_m \left[ \frac{t-c}{e} \right]. \quad (4.2)$$

Although Algorithm 4.2 only explicitly filters the  $[a, b]$  interval, we note that by the property of the Chebyshev polynomial, the filter values on the interval to the left of  $[a, b]$  will be magnified—which is what is needed to approximate the eigensubspace associated with the lower end of  $\sigma(H)$ .

As seen from Algorithm 4.2, a desired filter can be easily controlled by providing two endpoints of the higher end of  $\sigma(H)$ . The higher endpoint can be estimated by a few steps of standard Lanczos (which is presented in Algorithm 4.3 below). A lower bound of the full  $\sigma(H)$  is not needed. Instead, the wanted lower bound is any value which is larger than the Fermi-level but smaller than the higher endpoint. As presented in Algorithm 4.1, this lower bound is readily available as the largest Rayleigh-quotient of the previous iteration. Hence there is negligible extra work associated with computing bounds for the Chebyshev filtering, the major cost being in the three-term recurrences in Algorithm 4.2 which involve matrix-vector products. The polynomial degree  $m$  is left as a free parameter. Our experiences indicates that an  $m$  within 8 to 20 is good enough to achieve overall fast convergence in the SCF loop.

ALGORITHM 4.2.  $[Y] = \text{Chebyshev\_filter}(X, m, a, b)$ .

---

*Purpose:* Filter vectors in  $X$  by an  $m$  degree Chebyshev polynomial that dampens on the interval  $[a, b]$ , and output the filtered vectors in  $Y$ .

1.  $e = (b - a)/2$ ;  $c = (b + a)/2$ ;
  2.  $\sigma = e/(a - c)$ ;
  3.  $\sigma_1 = \sigma$ ;
  4.  $Y = (HX - cX)\sigma_1/e$ ;
  5. For  $i = 2 : m$ 
    6.  $\sigma_2 = 1/(2/\sigma_1 - \sigma)$ ;
    7.  $Y_{new} = 2(HY - cY)\sigma_2/e - \sigma\sigma_2X$ ;
    8.  $X = Y$ ;
    9.  $Y = Y_{new}$ ;
    10.  $\sigma = \sigma_2$ ;
  11. End For
- 

**4.3. Estimating an upper bound of the spectrum.** Now we present an inexpensive way to estimate an upper bound of  $\sigma(H)$ . As pointed out in [45], the upper bound has to bound the full spectrum of  $H$ . This is because the Chebyshev polynomial also grows fast on the right of  $[-1, 1]$ . So if  $[a, b]$  with  $b < \sigma_{max}(H)$  is mapped into  $[-1, 1]$ , then the  $[b, \sigma_{max}(H)]$  portion of the spectrum will also be magnified, which will cause the whole procedure to fail. We need  $b$  to be larger than  $\sigma_{max}(H)$  but it cannot be too large as this will affect convergence. There are several ways to estimate this type of upper bound, for example, by using the one-norm of  $H$ , or by applying Gerschgorin's Circle Theorem if possible. Bounds obtained this way can, however, overestimate  $\sigma_{max}(H)$ .

We found that a few steps of the standard Lanczos procedure are sufficient to provide an effective upper bound. For example, a  $k$ -step Lanczos leads to a Lanczos decomposition

$$HV_k = V_k T_k + f_k e_k^T$$

where  $V_k$  contains the  $k$  Lanczos basis,  $T_k$  is a size- $k$  tridiagonal matrix,  $f_k$  is a residual vector and  $e_k$  is a length  $k$  unit vector with only the first element nonzero.

Notice that Lanczos iteration often quickly approximate the outermost eigenvalues, and that

$$\|HV_k\|_2 = \|V_k T_k + f_k e_k^T\|_2 \leq \|T_k\|_2 + \|f_k\|_2.$$

We can start with a random unit vector, carry out  $k$  steps of the Lanczos procedure, and use  $\|T_k\|_2 + \|f_k\|_2$  as an upper bound for  $\sigma(H)$ . This is presented in Algorithm 4.3. For simplicity we skip the safeguards for  $\beta = 0$ , as this does not happen for small values of  $k$  in general.

In practice, we found that  $k = 4$  or  $k = 5$  is sufficient to yield a proper upper bound of  $\sigma(H)$ . In fact we found that it is often counter-productive to take the bounds obtained from larger values of  $k$  (say,  $k > 10$ ).

**5. Numerical Results.** The numerical experiments are performed using our own DFT package called PARSEC which is written in Fortran 95. PARSEC is based



ALGORITHM 4.3. *Estimating an upper bound of  $\sigma(H)$  by  $k$ -step Lanczos:*

- 
1. Generate a random vector  $v$ , set  $v \leftarrow v/\|v\|_2$ ;
  2. Compute  $f = Hv$ ;  $\alpha = f^T v$ ;  $f \leftarrow f - \alpha v$ ;  $T(1,1) = \alpha$ ;
  3. Do  $j = 2$  to  $\min(k,10)$
  4.      $\beta = \|f\|_2$ ;
  5.      $v_0 \leftarrow v$ ;  $v \leftarrow f/\beta$ ;
  6.      $f = Hv$ ;  $f \leftarrow f - \beta v_0$ ;
  7.      $\alpha = f^T v$ ;  $f \leftarrow f - \alpha v$ ;
  8.      $T(j,j-1) = \beta$ ;  $T(j-1,j) = \beta$ ;  $T(j,j) = \alpha$ ;
  9. End Do
  10. Return  $\|T\|_2 + \|f\|_2$  as the upper bound.
- 




FIG. 5.1. *37-stencil of a 12-th order centered finite difference.*

on the real-space pseudopotential method [9, 10], where higher order finite differences are used for the discretization of the Kohn-Sham equation on a uniform grid, as discussed in Section 2. We use a 12th order centered finite difference scheme. At each point  $(x, y, z)$  the local stencil involves 37-points. Figure 5.1 illustrates this stencil.

We compare three different methods in PARSEC, two of which are based on solving eigenvalue problems (2.1) at each SCF iteration. The solvers used are ARPACK [25] and TRLan [43, 44] which represent two of the most efficient publicly available eigenvalue packages. The third method is our Chebyshev filtered subspace iteration, with the initial SCF step solved using TRLan. Note that each method is applied in the same package to the same problems, that is, each method can exploit the same physical symmetry operations if they exist, hence the performance difference is due solely to the three different methods.

Note that we only use the solver for standard symmetric eigenproblems in ARPACK. We mention that ARPACK is currently one of the most efficient general purpose eigensolvers, especially for nonsymmetric eigenproblems. It is by far the best known package for general eigenvalue calculations.

TRLan is a Fortran 90 package for standard symmetric eigenproblems. It uses reduced full orthogonalization and several *thick restart* strategies [43, 44], hence it can be more efficient than the symmetric eigensolver in ARPACK. We call TRLan

for the first step SCF iteration for our Chebyshev filtered subspace iteration method. We denote the new method as *ChebSI*.




FIG. 5.2. Atomic structure of the quantum dot  $Si_{525}H_{276}$ . The red and white balls represent Si and H atoms respectively. The quantum dot contains 25 shells of Si atoms and is 27.2 Å in diameter [8].

The test problems are a Silicon clusters  $Si_{525}H_{276}$ , a Silicon Germanium cluster  $Si_{65}Ge_{65}H_{98}$ , a Gallium Arsenide cluster  $Ga_{41}As_{41}H_{72}$ , and two iron clusters  $Fe_{27}$  and  $Fe_{51}$ . These problems constitute four typical materials in electronic structure calculations. We note that metallic systems are difficult for SCF calculation because of the *charge sloshing* effect [22].

Figure 5.2 shows the atomic structure of  $Si_{525}H_{276}$ . Table 5.1 contains some numerical parameters related to these test problems. The number of symmetry operations used to construct the irreducible wedge is denoted  $n_{symm}$ , and “reduced *H* size” is the number of grid points in the wedge (equal to the dimension of the discretized reduced Hamiltonian). At each SCF iteration, ARPACK and TRLan compute approximately  $n_{state}$  eigenpairs from  $n_{symm}$  reduced Hamiltonians at *Step 2* of Algorithm 2.1, while *ChebSI* replaces *Step 2* of Algorithm 2.1 by Algorithm 4.1 and returns the same number of wave vectors as ARPACK or TRLan does from *Step 2* to *Step 3* in Algorithm 2.1. The number of states of  $Fe_{27}$  and  $Fe_{51}$  is doubled by 2

because of spin degrees of freedom, i.e., for each spin channel at least 520 eigenpairs are computed.

In each table, the `total_eV/atom` counts the total energy per atom, given in electron-volts. This is a control parameter used to assess the accuracy of the final result. The `# SCF steps` is the iteration steps used to reach self-consistency; the `# MV products` counts the number of matrix-vector products. Clearly this is not the only factor that determines CPU time, the reduced orthogonalization can also have a crucial effect in CPU time.

model	size of $H$	$n_{state}$	$n_{symm}$	reduced $H$ size
$Si_{525}H_{276}$	292584	1194	4	73146
$Si_{65}Ge_{65}H_{98}$	185368	313	2	92684
$Ga_{41}As_{41}H_{72}$	268096	210	1	268096
$Fe_{27}$	697504	$520 \times 2$	8	87188
$Fe_{51}$	874976	$520 \times 2$	8	109372

TABLE 5.1  
Relevant data of the test problems

method	# MV products	# SCF steps	total_eV /atom	CPU(secs)
ChebSI	124761	11	-77.316873	5946.69
ARPACK	142047	10	-77.316873	62026.37
TRLan	145909	10	-77.316873	26852.84

TABLE 5.2  
 $Si_{525}H_{276}$ , Polynomial degree used is 8.

method	# MV products	# SCF steps	total_eV /atom	CPU (secs)
ChebSI	42919	13	-140.076118	2344.06
ARPACK	51752	9	-140.076118	12770.81
TRLan	53892	9	-140.076118	6056.11

TABLE 5.3  
 $Si_{65}Ge_{65}H_{98}$ , Polynomial degree used is 8.

All the numerical runs are performed on the SGI Altix 3700 cluster of the Minnesota supercomputing Institute. The CPU type is a 1.3 GHz Intel Madison processor. The Operating System is a 64bit Linux with kernel version 2.4.21. The compiler used is the Intel Fortran compiler `ifort`, with optimization flag `-O3` for all codes.

As seen from Tables 5.2–5.6, the ChebSI method is usually five to ten times faster than the eigenvector-based methods represented by two of the best available iterative sparse eigensolvers ARPACK and TRLan. Although eigenspaces are not explicitly and accurately computed at each SCF step, ChebSI only requires a few more SCF steps to reach self-consistency. However, each SCF step using ChebSI is much cheaper than an SCF step based on eigenvectors. About ten other tests on smaller systems (Hamiltonian sizes around 100K–160K) not reported here showed similar results. We also mention that there are cases where the Chebyshev filtering may become less effective. We observed these in Gallium Arsenide clusters, as seen from Table 5.4.

method	# MV products	# SCF steps	totalLeV /atom	CPU (secs)
ChebSI	138672	37	-89.634940	12923.27
ARPACK	58506	10	-89.634940	44305.97
TRLan	58794	10	-89.634940	16733.68

TABLE 5.4

$Ga_{41}As_{41}H_{72}$ . Polynomial degree used is 16. The spectrum of each reduced Hamiltonian spans a large interval, making the Chebyshev filtering not as effective as other examples.

method	# MV products	# SCF steps	totalLeV /atom	CPU (secs)
ChebSI	363728	30	-776.575290	15408.16
ARPACK	750883	21	-776.586420	118693.64
TRLan	807652	21	-776.586422	83726.20

TABLE 5.5

$Fe_{27}$ , Polynomial degree used is 9.

The cases of reduced efficiency correspond to situations where the full spectrum of a given Hamiltonian spreads over a very large interval. In such situations, high degree Chebyshev polynomials must be used. However, we see that even in these unfavorable cases, ChebSI is still reasonably faster than eigenvector-based methods.

Finally, we should mention that ChebSI is as robust as eigenvector-based methods. In fact, it has been used successfully for unreported periodic systems, with performance gains similar to the ones obtained in clusters. On the other hand, an eigenvector-based method using the preconditioned Davidson method failed to converge for the same periodic systems.

**6. Conclusions.** An algorithm based on Chebyshev filtered subspace iteration has been developed for performing SCF calculations in Density Function Theory, which has the advantage of not explicitly relying on eigenvectors, except at the first SCF step. This leads to significant gains in computational time relative to traditional eigenvector-based methods which are inevitably constrained by the high computational cost of diagonalization at each SCF step. Numerical results show that the new method is five to ten times faster than eigenvector-based methods using two of the best iterative eigensolvers. Even though implemented sequentially at present, the ChebSI method can solve realistic systems of moderate size within a reasonable time frame. Parallel implementation of the ChebSI method will be reported in a forthcoming paper, where we study larger and more complex material systems.

**Acknowledgements:** We thank Dr. Manuel Alemany for his help in using PARSEC and Dr. Shen Li for providing us with the physical data of the iron clusters. Calculations were performed at the Minnesota Supercomputing Institute.

## REFERENCES

- [1] M. M. G. ALEMANY, M. JAIN, L. KRONIK, AND J. R. CHELIKOWSKY, *Real-space pseudopotential method for computing the electronic properties of periodic systems*, Phys. Rev. B, 69 (2004), pp. 075101–1–6.
- [2] R. BAER AND M. HEAD-GORDON, *Chebyshev expansion methods for electronic structure calculations on large molecular systems*, J. Chem. Phys., 107 (1997), pp. 10003–10013.
- [3] S. BARONI AND P. GIANNOZZI, *Towards very large scale electronic structure calculations*, Europhys. Lett., 17 (1992), pp. 547–552.

method	# MV products	# SCF steps	total_eV /atom	CPU (secs)
ChebSI	474773	37	-777.019294	37701.54
ARPACK	1272441	34	-777.038081	235662.96
TRLan	1241744	32	-777.038086	184580.33

TABLE 5.6  
 $Fe_{51}$ , Polynomial degree used is 9.

- [4] C. BEKAS, Y. SAAD, M. L. TIAGO, AND J. R. CHELIKOWSKY, *Computing charge densities with partially reorthogonalized lanczos*, tech. rep., CS Dept., Univ. of Minnesota, 2005.
- [5] A. BRANDT, *Multiscale scientific computation: review 2001*, in *Multiscale and multiresolution methods*, T. Barth, T. Chan, and R. Haimes, eds., vol. 20 of *Lect. Notes Comput. Sci. Eng.*, Springer-Verlag, 2002, pp. 3–95.
- [6] E. L. BRIGGS, D. J. SULLIVAN, AND J. BERNHOLC, *Large-scale electronic-structure calculations with multigrid acceleration*, *Phys. Rev. B*, 52 (1995), pp. R5471–5474.
- [7] R. CAR AND M. PARRINELLO, *Unified approach for molecular dynamics and density-functional theory*, *Phys. Rev. Lett.*, 55 (1985), pp. 2471–2474.
- [8] J. CHELIKOWSKY, *The pseudopotential-density functional method applied to nanostructures*, *J. Phys. D: Appl. Phys.*, 33 (2000), pp. R33–R50.
- [9] J. R. CHELIKOWSKY, N. TROULLIER, AND Y. SAAD, *Finite-difference-pseudopotential method: Electronic structure calculations without a basis*, *Phys. Rev. Lett.*, 72 (1994), pp. 1240–1243.
- [10] J. R. CHELIKOWSKY, N. TROULLIER, K. WU, AND Y. SAAD, *Higher-order finite-difference pseudopotential method: An application to diatomic molecules*, *Phys. Rev. B*, 50 (1994), pp. 11355–11364.
- [11] S. COSTINER AND S. TÁASAN, *Adaptive multigrid techniques for large-scale eigenvalue problems: Solutions of the schrödinger problem in two and three dimensions*, *Phys. Rev. E*, 51 (1995), pp. 3704–3717.
- [12] J. DANIEL, W. B. GRAGG, L. KAUFMAN, AND G. W. STEWART, *Reorthogonalization and stable algorithms for updating the Gram–Schmidt QR factorization*, *Math. Comp.*, 30 (1976), pp. 772–795.
- [13] R. M. DREIZLER AND E. K. U. GROSS, *Density functional theory: an approach to the quantum many-body problem*, Springer-Verlag, Berlin, 1990.
- [14] J.-L. FATTEBERT AND J. BERNHOLC, *Towards grid-based  $O(N)$  density-functional theory methods: Optimized nonorthogonal orbitals and multigrid acceleration*, *Phys. Rev. B*, 63 (2000), pp. 1713–1722.
- [15] S. GOEDECKER, *Linear scaling electronic structure methods*, *Rev. Mod. Phys.*, 71 (1999), pp. 1085–1123.
- [16] S. GOEDECKER AND L. COLOMBO, *Efficient linear scaling algorithm for tight-binding molecular dynamics*, *Phys. Rev. Lett.*, 73 (1994), pp. 122–125.
- [17] M. HEISKANEN, T. TORSTI, M. J. PUSKA, AND R. M. NIEMINEN, *Multigrid method for electronic structure calculations*, *Phys. Rev. B*, 63 (2001), p. 245106.
- [18] P. HOHENBERG AND W. KOHN, *Inhomogeneous electron gas*, *Phys. Rev.*, 136 (1964), pp. B864–B871.
- [19] L. O. JAY, H. KIM, Y. SAAD, AND J. R. CHELIKOWSKY, *Electronic structure calculations using plane wave codes without diagonalization*, *Comput. Phys. Comm.*, 118 (1999), pp. 21–30.
- [20] T. KERKHOVEN, A. GALICK, U. RAVAIOLI, J. ARENDS, AND Y. SAAD, *Efficient numerical simulation of electron states in quantum wires*, *J. Applied Physics*, 68 (1990), pp. 3461–3469.
- [21] W. KOHN AND L. J. SHAM, *Self-consistent equations including exchange and correlation effects*, *Phys. Rev.*, 140 (1965), pp. A1133–A1138.
- [22] G. KRESSE AND J. FURTHMÜLLER, *Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set*, *Phys. Rev. B*, 54 (1996), pp. 11169–11186.
- [23] C. LANCZOS, *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*, *J. Research Nat. Bur. Standards*, 45 (1950), pp. 255–282.
- [24] I. LEE, Y. KIM, AND R. M. MARTIN, *One-way multigrid method in electronic-structure calculations*, *Phys. Rev. B*, 61 (2000), pp. 4397–4400.
- [25] R. B. LEHOUCQ, D. C. SORENSEN, AND C. YANG, *ARPACK USERS GUIDE: Solution of Large Scale Eigenvalue Problems by Implicitly Restarted Arnoldi Methods*, SIAM, Philadelphia, 1998. Available at <http://www.caam.rice.edu/software/ARPACK/>.
- [26] J. L. MARTINS AND M. COHEN, *Diagonalization of large matrices in pseudopotential band-*

- structure calculations: Dual-space formalism*, Phys. Rev. B, 37 (1988), pp. 6134–6138.
- [27] J. L. MARTINS, N. TROULLIER, AND S.-H. WEI, *Pseudopotential plane-wave calculations for ZnS*, Phys. Rev. B, 43 (1991), pp. 2213–2217.
- [28] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, no. 20 in Classics in Applied Mathematics, SIAM, Philadelphia, PA, 1998.
- [29] R. PARR AND W. YANG, *Density Functional Theory of Atoms and Molecules*, Oxford Univ. Press, 1989.
- [30] M. C. PAYNE, M. P. TETER, D. C. ALLAN, T. A. ARIAS, AND J. D. JOANNOPOULOS, *Iterative minimization techniques for ab initio total-energy calculations: molecular dynamics and conjugate gradients*, Rev. Mod. Phys., 64 (1992), pp. 1045–1097.
- [31] P. PULAY, *Convergence acceleration of iterative sequences: The case of SCF iteration*, Chem. Phys. Lett., 73 (1980), pp. 393–398.
- [32] ———, *Improved SCF convergence acceleration*, J. Comput. Chem., 3 (1982), pp. 556–560.
- [33] H. RUTISHAUSER, *Simultaneous iteration method for symmetric matrices*, in Handbook for Automatic Computation (Linear Algebra), J. H. Wilkinson and C. Reinsh, eds., vol. II, Springer-Verlag, 1971, pp. 284–302.
- [34] Y. SAAD, *Numerical Methods for Large Eigenvalue Problems*, John Wiley, New York, 1992. Available at <http://www.cs.umn.edu/~saad/books.html>.
- [35] Y. SAAD, A. STATHOPOULOS, J. CHELIKOWSKY, K. WU, , AND S. ÖĞÜT, *Solution of large eigenvalue problems in electronic structure calculations*, BIT, 36 (1996), pp. 563–578.
- [36] O. F. SANKEY, D. A. DRABOLD, AND A. GIBSON, *Projected random vectors and the recursion method in the electronic-structure problem*, Phys. Rev. B, 50 (1994), pp. 1376–1381.
- [37] A. P. SEITSONEN, M. J. PUSKA, AND R. M. NIEMINEN, *Real-space electronic-structure calculations: Combination of the finite-difference and conjugate-gradient methods*, Phys. Rev. B, 51 (1995), pp. 14057–14061.
- [38] D. C. SORENSEN, *Implicit application of polynomial filters in a k-step Arnoldi method*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 357–385.
- [39] A. STATHOPOULOS, S. ÖĞÜT, Y. SAAD, J. CHELIKOWSKY, AND H. KIM, *Parallel methods and tools for predicting materials properties*, IEEE Computing in Science and Engineering, 2 (2000), pp. 9–18.
- [40] U. STEPHAN, D. A. DRABOLD, AND R. M. MARTIN, *Improved accuracy and acceleration of variational order-n electronic structure computations by projection techniques*, Phys. Rev. B, 58 (1998), pp. 13472–13481.
- [41] M. L. TIAGO AND J. R. CHELIKOWSKY, tech. rep., Univ. Texas, Austin, (in preparation).
- [42] D. WOOD AND A. ZUNGER, *A new method for diagonalising large matrices*, J. Phys. A: Math. Gen., 18 (1985), pp. 1343–1359.
- [43] K. WU, A. CANNING, H. D. SIMON, AND L.-W. WANG, *Thick-restart Lanczos method for electronic structure calculations*, J. Comput. Phys., 154 (1999), pp. 156–173.
- [44] K. WU AND H. SIMON, *Thick-restart Lanczos method for large symmetric eigenvalue problems*, SIAM J. Matrix Anal. Appl., 22 (2000), pp. 602–616.
- [45] Y. ZHOU AND Y. SAAD, *A Chebyshev-Davidson algorithm for large symmetric eigenvalue problems*, tech. rep., Minnesota Supercomputing Institute, University of Minnesota, 2005 (Submitted to SIAM J. Matrix Anal. Appl.).