Linear-scaling subspace-iteration algorithm with optimally localized nonorthogonal wave functions for Kohn-Sham density functional theory

C. J. García-Cervera*
Department of Mathematics, University of California, Santa Barbara, California 93106, USA

Jianfeng Lu† and Yulin Xuan‡
PACM, Princeton University, Princeton, New Jersey 08544, USA

Weinan E§
Department of Mathematics and PACM, Princeton University, Princeton, New Jersey 08544, USA
(Received 4 August 2008; revised manuscript received 3 February 2009; published 13 March 2009)

We present a linear-scaling method for electronic structure computations in the context of Kohn-Sham density functional theory (DFT). The method is based on a subspace iteration, and takes advantage of the nonorthogonal formulation of the Kohn-Sham functional, and the improved localization properties of nonorthogonal wave functions. A one-dimensional linear problem is presented as a benchmark for the analysis of linear-scaling algorithms for Kohn-Sham DFT. Using this one-dimensional model, we study the convergence properties of the localized subspace-iteration algorithm presented. We demonstrate the efficiency of the algorithm for practical applications by performing fully three-dimensional computations of the electronic density of alkane chains.

DOI: 10.1103/PhysRevB.79.115110 PACS number(s): 31.15.E−, 31.15.xr, 31.15.xf, 71.15.−m

I. INTRODUCTION

Kohn-Sham density functional theory (DFT) (Refs. 1–3) is a very popular tool for electronic structure analysis, with applications in materials science, chemistry, and other areas. Compared to the quantum many-body problem, which describes a system with N electrons by a $3N$-dimensional antisymmetric wave function (ignoring spin degeneracy), the Kohn-Sham density functional theory describes such a system by N one-particle wave functions. By suitably approximating the functionals, Kohn-Sham DFT can become a very effective alternative to the quantum many-body problem, with satisfactory accuracy and much improved efficiency.[4]

The Kohn-Sham equations are a system of nonlinear eigenvalue problems. The traditional self-consistent approach[1] for the solution of this eigenvalue problem consists of two nested iterations: in the inner iteration, the orbitals $\{\psi_j\}_{j=1}^N$ are obtained by a process of diagonalization and orthogonalization; in the outer iteration, the electron density is updated until self-consistency is reached. The diagonalization and/or orthogonalization procedure scales typically as $O(N^3)$, which is prohibitively expensive for relatively small problems.

In the past 20 years, a number of new methodologies appeared in the literature, which attempt to exploit the locality of the problem in order to reduce the computational complexity.[5] Locality, in quantum mechanics, refers to the property that a small disturbance in a molecule only has a local effect in the electron density, a phenomenon coined by Kohn[6] as *nearsightedness*. In the divide-and-conquer method of Yang,[7,8] the electron density is divided into a set of subsystems that are coupled through the chemical potential and each subsystem is solved separately by diagonalization. In the density-matrix minimization,[9,10] the energy is rewritten in terms of the density matrix, combined with the McWeeny[11] purification transformation. The density matrix is then obtained using nonlinear conjugate gradient minimization. Linear scaling is achieved by truncating the density matrix beyond a predetermined cutoff radius. Galli and Parrinello[12] introduced a plane-wave-based algorithm using localized nonorthogonal wave functions. In their approach, the localization of the wave functions was imposed by adding an additional potential to the Hamiltonian. An unconstrained variational approach is described in Ref. 13. It is observed that the direct truncation of the wave functions implemented in Ref. 13 can introduce artificial minima in the energy.[14,15] Various approaches have been raised to ameliorate the problem, such as increasing the number of localized orbitals, using the pseudoinverse of the overlap matrix,[16,17] and introducing a localization procedure before the truncation.[14]

In this paper, we describe in detail the linear-scaling algorithm introduced in Ref. 18. The algorithm is based on a subspace-iteration procedure with localized nonorthogonal wave functions. The general ideas of such a localized subspace iteration (LSI) were first introduced in Ref. 19. In the context of electronic structure analysis, the subspace-iteration approach has been used in Ref. 20. Our algorithm replaces the diagonalization and orthogonalization step in the subspace-iteration algorithm introduced in Ref. 20 by a localization procedure (described in detail in Sec. II B). This localization procedure differs from the one used in Ref. 12 in that it does not introduce additional terms in the Hamiltonian. Linear scaling is achieved in our algorithm by exploiting the localization properties of nonorthogonal wave functions.

Kohn-Sham DFT and the Kohn-Sham equations are fairly complicated problems. As a variational problem, the Kohn-Sham functional is nonlinear, nonlocal, and nonconvex. The Kohn-Sham equations are a set of nonlinear and nonlocal eigenvalue problems. For practical applications, there is the added complexity coming from the exchange and correlation

functional and the pseudopotential. These factors together make Kohn-Sham DFT a rather formidable problem from a mathematical and numerical viewpoints. However, some of these complexities are inessential for the purpose of developing and understanding numerical algorithms. For these reasons, we introduce here a one-dimensional linear model and use it as a benchmark for clarifying the essential features of the localized subspace-iteration algorithm, its performance and its limitations. By *linear* we mean that the Hamiltonian is linear, i.e., the potential does not depend on the wave functions. This one-dimensional problem, as well as its nonlinear version, which will be discussed in subsequent publications, is very simple and it has many of the essential features of the Kohn-Sham DFT. In subsequent papers, we will present, using this model, a systematic analysis of other linear-scaling algorithms and self-consistent iterations for Kohn-Sham DFT.

One interesting feature of the LSI is that it does not converge to a fixed point. Under favorable conditions (reasonable gap size, large enough localization region, etc), the numerical errors will decrease to a rather small value and then start to fluctuate. This is due to the conflicting effect of the filtering and truncation steps in the LSI and should be a general feature of linear-scaling algorithms involving truncation. Using the one-dimensional model, we are able to understand the convergence properties of the LSI in considerable detail.

This paper is organized as follows. In the remainder of section, we will review Kohn-Sham DFT and some general numerical issues. Hartree atomic units are adopted throughout the text ($\hbar=e=m_e=1$). The linear-scaling algorithm and its main components are described in considerable detail in Sec. II. We introduce the aforementioned one-dimensional model problem, and analyze how the different components of the algorithm impact its performance and accuracy, in Sec. III. The fully three-dimensional implementation of the algorithm for Kohn-Sham DFT is described in Sec. IV, where the results are validated and linear scaling is demonstrated for practical examples.

A. Kohn-Sham density functional theory

Consider a system consisting of N_a atoms and $2N$ electrons. In Kohn-Sham DFT, the state of the system is described by a set of N wave functions $\{\psi_j\}_{j=1}^N$, representing the interacting electrons. We consider the wave functions to be real throughout this article. In the Born-Oppenheimer approximation, the spinless Kohn-Sham energy functional can be written as

$$E_{\mathrm{KS}}[\{\psi_j\}] = 2\sum_{j=1}^{N}\int_{\mathbb{R}^3}\psi_j\left(-\frac{1}{2}\Delta\psi_j\right)d\boldsymbol{x} + E_{\mathrm{XC}}[\rho] + \frac{1}{2}\int_{\mathbb{R}^3}\int_{\mathbb{R}^3}\frac{(\rho-m)(\boldsymbol{x})(\rho-m)(\boldsymbol{y})}{|\boldsymbol{x}-\boldsymbol{y}|}d\boldsymbol{x}d\boldsymbol{y} + E_{\mathrm{PS}}[\{\psi_i\}], \tag{1}$$

where the electron density is defined as

$$\rho(\boldsymbol{x}) = 2\sum_{j=1}^{N}|\psi_j(\boldsymbol{x})|^2,$$

and the wave functions $\{\psi_i\}_{j=1}^N$ are required to be orthonormal

$$\int_{\mathbb{R}^3}\psi_i(\boldsymbol{x})\psi_j(\boldsymbol{x})d\boldsymbol{x} = \delta_{ij}, \quad i,j=1,\ldots,N. \tag{2}$$

In Eq. (1), the exact form of the exchange and correlation functional $E_{\mathrm{XC}}[\rho]$ is unknown and needs to be approximated. We adopt here the local-density approximation (LDA), in which this term is written as

$$E_{\mathrm{XC}}(\rho) = \int_{\mathbb{R}^3}\rho(\boldsymbol{x})\varepsilon_{\mathrm{XC}}[\rho(\boldsymbol{x})]d\boldsymbol{x}, \tag{3}$$

where $\varepsilon_{\mathrm{XC}}$ is a function of the density alone. The generalized gradient approximation (GGA),[21,22] in which $\varepsilon_{\mathrm{XC}}$ depends on the gradient of ρ as well, can be used instead and this will not change the discussions below in any significant way.

We consider ρ to be the density of valence electrons, and their electrostatic interactions with the nuclei and the core electrons are described by a so-called pseudopotential. The pseudopotential is completely defined by an ionic function, m, and an ionic pseudopotential operator, $\hat{V}_{\mathrm{PS}}$. The ionic function is

$$m(\boldsymbol{x}) = \sum_{j=1}^{N_a} m^a(\boldsymbol{x}-\boldsymbol{R}_j),$$

where m^a is a localized function (usually with exponential decay) and $\{\boldsymbol{R}_j\}_{j=1}^{N_a}$ are the locations of the atoms. In this article, we use the norm conserving Troullier-Martins pseudopotential[23] in the Kleinman-Bylander form.[24] We therefore define the ionic pseudopotential operator as

$$\hat{V}_{\mathrm{PS}}\psi(\boldsymbol{x}) = \sum_{j=1}^{N_a}\left(V^j_{\mathrm{Local}}(\boldsymbol{x}-\boldsymbol{R}_j)\psi(\boldsymbol{x}) + \sum_{l=0}^{l_{\max}}\sum_{m=-l}^{l}\int_{\mathbb{R}^3}\beta^j_{lm}(\boldsymbol{y}-\boldsymbol{R}_j)\psi(\boldsymbol{y})d\boldsymbol{y}\,\beta^j_{lm}(\boldsymbol{x}-\boldsymbol{R}_j)\right), \tag{4}$$

where V^j_{Local} is the local potential for the jth atom (by which we mean that it acts by pointwise multiplication) and $\{\beta^j_{lm}\}$ projects ψ onto the subspace generated by the corresponding pseudowave function of the jth atom. The indices $\{l,m\}$ indicate the angular-momentum components of the pseudowave function. By introducing the ionic function m, we can make sure that V^j_{Local} and $\{\beta^j_{lm}\}$ are all finite supported. This is important in order to achieve linear scaling.

In Eq. (1), the pseudopotential energy E_{PS} is defined as

$$E_{\mathrm{PS}}[\{\psi_i\}] = 2\sum_{i=1}^{N}\int_{\mathbb{R}^3}\psi_i(\boldsymbol{x})\hat{V}_{\mathrm{PS}}\psi_i(\boldsymbol{x})d\boldsymbol{x}. \tag{5}$$

When the Kohn-Sham functional (1) is minimized under the orthogonality constraint (2), the Euler-Lagrange equations lead to the following nonlinear eigenvalue problem:

$$\mathbf{H}\psi_i = \varepsilon_i \psi_i; \quad i = 1, \ldots, N. \tag{6}$$

The Hamiltonian in Eq. (6) is defined as

$$\mathbf{H} = -\frac{1}{2}\Delta + V_{\text{eff}}[\rho], \tag{7}$$

where

$$V_{\text{eff}}[\rho](\boldsymbol{x}) = \int_{\mathbb{R}^3} \frac{(\rho - m)(\boldsymbol{y})}{|\boldsymbol{x} - \boldsymbol{y}|} d\boldsymbol{y} + V_{\text{PS}}(\boldsymbol{x}) + \frac{\delta E_{\text{XC}}[\rho]}{\delta \rho}.$$

The starting point of our linear-scaling approach is the reformulation of the Kohn-Sham problem in terms of nonorthogonal wave functions, as introduced in Ref. 12. Given N linearly independent wave functions, $\{\psi_j\}_{j=1}^N$, define the overlap or Gram matrix as

$$\mathbf{S}_{jk} = \int_{\mathbb{R}^3} \psi_j \psi_k, \quad j,k = 1, \ldots, N. \tag{8}$$

Then,

$$E_{\text{KS}}[\{\psi_j\}] = 2\sum_{j,k} (\mathbf{S}^{-1})_{jk} \int_{\mathbb{R}^3} \psi_j \left(-\frac{1}{2}\Delta \psi_k\right) d\boldsymbol{x} + E_{\text{XC}}[\rho] + \frac{1}{2}\int_{\mathbb{R}^3}\int_{\mathbb{R}^3} \frac{(\rho - m)(\boldsymbol{x})(\rho - m)(\boldsymbol{y})}{|\boldsymbol{x} - \boldsymbol{y}|} d\boldsymbol{x} d\boldsymbol{y} + E_{\text{PS}}[\{\psi_i\}], \tag{9}$$

where the electron density is now defined by

$$\rho(\boldsymbol{x}) = 2\sum_{jk} \psi_j(\boldsymbol{x})(\mathbf{S}^{-1})_{jk}\psi_k(\boldsymbol{x}), \tag{10}$$

and the pseudopotential energy is now

$$E_{\text{PS}}[\{\psi_i\}] = 2\sum_{j,k} (\mathbf{S}^{-1})_{jk} \int_{\mathbb{R}^3} \psi_j(\boldsymbol{x}) \hat{V}_{\text{PS}} \psi_k(\boldsymbol{x}) d\boldsymbol{x}. \tag{11}$$

It is easy to show that the Kohn-Sham functional (9) is invariant under nonsingular linear transformations of the N-dimensional subspace spanned by the wave functions

$$\text{span}\{\psi_i\} = \text{span}\{\phi_i\} \Rightarrow \text{E}_{\text{KS}}[\{\psi_i\}] = \text{E}_{\text{KS}}[\{\phi_i\}]. \tag{12}$$

As a consequence, the Kohn-Sham functional can be thought of as a functional acting on subspaces and the goal is to find the *occupied subspace*. This is the subspace generated by the eigenfunctions corresponding to the smallest eigenvalues of the self-consistent Hamiltonian. The advantage of this viewpoint is that the specific representation of the subspace is not relevant and therefore we can choose the representation that is most convenient for our purposes. Linear scaling can be achieved by choosing a representation in terms of optimally localized nonorthogonal wave functions, as described in Ref. 19. Nonorthogonal wave functions have better localization properties than orthogonal orbitals,[25,26] which constitutes an additional advantage of this approach.

B. Numerical algorithms

A numerical algorithm for the Kohn-Sham problem should have the following basic components:

(1) A discretization procedure that translates the continuous problem into a discrete (finite-dimensional) problem. Typical discretization procedures include the plane-wave methods, methods based on a linear combination of atomic orbitals, finite difference, and finite element methods.[27]

(2) An algorithm for handling the linearized finite-dimensional problem. In particular, assuming that the electron density is given, we need to find the optimal eigensubspace for the Kohn-Sham operator with a known potential. This can be done by diagonalization or some other procedure.[18,27]

(3) A nonlinear iteration to achieve self-consistency using some mixing procedure for the density.[27]

In the present paper, we will focus on the second component, i.e., finding the minimizing subspace of the Kohn-Sham operator. For the numerical results presented below, we use a simple second-order finite difference discretization and simple linear mixing

$$\rho^{\text{new}} = \alpha\rho^{\text{old}} + (1 - \alpha)\rho^*, \tag{13}$$

where ρ^* is the outcome of step 2 above. These are chosen for their simplicity. There is no difficulty in using higher order finite difference, finite element discretizations, or other mixing schemes. However, since we will focus on localized wave functions, our work does not apply to the case when plane waves are used. As reported in Ref. 28, the number of iterations required for self-consistency may depend on the size of the system. Given that we are using the same types of mixing as in Ref. 28, the same phenomenon is expected.

II. LOCALIZED SUBSPACE ITERATION FOR KOHN-SHAM DENSITY FUNCTIONAL THEORY

A. Subspace iteration

Consider the problem of computing the largest eigenvalue (in magnitude) of a symmetric matrix $\mathbf{H}$. For simplicity of presentation, let us assume that the eigenvalues of the matrix $\mathbf{H}$ are $0 \le \lambda_1 \le \lambda_2 \le \ldots \le \lambda_{N-1} < \lambda_N$. An effective approach is the power method:[29]

(1) Given an initial vector $\mathbf{v}^0$.

(2) For $k \ge 0$, define

$$\text{(a)} \quad \mathbf{v}^{k+1} = \frac{\mathbf{H}\mathbf{v}^k}{\|\mathbf{H}\mathbf{v}^k\|},$$

$$\text{(b)} \quad \mu^{k+1} = (\mathbf{v}^{k+1})^T \cdot \mathbf{H}\mathbf{v}^{k+1}. \tag{14}$$

(3) Repeat until $|\mu^{k+1} - \mu^k| \le$ tolerance.

One can show that the power method converges, i.e., $\mu^k \to \lambda_N$, with convergence rate

$$\tau = \frac{\lambda_{N-1}}{\lambda_N} < 1. \tag{15}$$

If other eigenvalues are required, one can use the shifted power method or the inverse shifted power method, which are essentially the power method applied to $\mathbf{H} - \lambda\mathbf{I}$ and $(\mathbf{H} - \lambda\mathbf{I})^{-1}$, respectively.

The subspace iteration is a straightforward generalization of the power method, applied to a subspace. A typical procedure is the following:

(1) Given an initial space V_0 of dimension $M<N$, for each $k\geq 1$,

(a) Calculate $W_k=\mathbf{H}V_k$.

(b) Orthogonalize the basis (QR decomposition, for example): $W_k=Q_kR_k$.

(c) Let $V_k=Q_k$.

(2) Repeat until convergence.

The orthogonalization step is necessary in order to ensure the linear independence of the vectors in the new space.

It can be proved that the previous algorithm converges and that the rate of convergence is[29]

$$\tau=\frac{\lambda_M}{\lambda_{M+1}}<1. \tag{16}$$

The cost of the orthogonalization step scales like $O(N^3)$. A key idea in the algorithm that we present here is to replace this step by a localization procedure. Given a candidate subspace, we find an optimally localized basis for this subspace. The details of this procedure are explained below.

B. Localization

Localization has been recognized by several authors as one of the key ideas in achieving linear scaling.[6,10,30] One of the first implementations of locality in Kohn-Sham DFT consisted on computing the so-called *maximally localized Wannier functions*,[30] in periodic systems, and for the orthogonal formulation of Kohn-Sham DFT (1).

This concept was generalized to the nonorthogonal case in Ref. 19. Given a linear space $V=\text{span}\{\psi_j\}_{j=1}^N$ of dimension N, and a given smooth weight function $w\geq 0$, the maximally localized nonorthogonal wave function $\tilde{\psi}$ is defined as

$$\tilde{\psi}=\arg\min_{\phi\in V,\|\phi\|=1}\int_{\mathbb{R}^3} w(\boldsymbol{x})|\phi(\boldsymbol{x})|^2 d\boldsymbol{x}. \tag{17}$$

Nonorthogonal wave functions have better localization properties than orthogonal Wannier functions. Consider the weight function $w(\boldsymbol{x})=|\boldsymbol{x}-\boldsymbol{x}_c|^{2p}$, where p is a positive integer (the maximally localized Wannier function essentially corresponds to the choice $p=1$). It is proved in Ref. 19 that for metallic systems, the resulting $\tilde{\psi}$ from the minimization decays asymptotically as $|\boldsymbol{x}-\boldsymbol{x}_c|^{-p-1}$. For systems with a small spectral gap, the nonorthogonal wave functions also decay faster in the algebraic regime (before exponential decay is dominant) and the nonorthogonal wave function will have the same exponential decay property as the orthogonal ones for insulators.[26]

Given a set of wave functions, $\{\psi_j\}_{j=1}^N$, centered at the locations $\{\boldsymbol{b}_j\}_{j=1}^N$, respectively, we obtain a localized representation of $V=\text{span}\{\psi_j\}_{j=1}^N$ by minimizing, for each $j\in\{1,\ldots,N\}$,

$$F[\phi]=\frac{\displaystyle\int_{\mathbb{R}^3}|\boldsymbol{x}-\boldsymbol{b}_j|^{2p}|\phi(\boldsymbol{x})|^2 d\boldsymbol{x}}{\displaystyle\int_{\mathbb{R}^3}|\phi(\boldsymbol{x})|^2 d\boldsymbol{x}} \tag{18}$$

among functions ϕ of the form

$$\phi(\boldsymbol{x})=\sum_{k=1}^{r}\alpha_k\psi_k(\boldsymbol{x})\in V. \tag{19}$$

The minimization (18) leads to $\mathbf{W}a=\lambda\mathbf{S}a$, where

$$W_{kl}=\int_{\mathbb{R}^3}|\boldsymbol{x}-\boldsymbol{b}_j|^{2p}\psi_k(\boldsymbol{x})\psi_l(\boldsymbol{x})d\boldsymbol{x},\quad k,l=1,\cdots,r,$$

and λ is the smallest generalized eigenvalue.

In the implementation of the algorithm, the wave functions $\{\psi_j\}_{j=1}^N$ are truncated beyond a certain cutoff radius, which is determined at the beginning of the simulation. Therefore, the wave functions used are always finite supported. As a result, only a fixed number of wave functions appear in Eq. (19) and the localized basis can be obtained with $O(N)$ cost. The actual value of r depends on the geometry of the system, i.e., the location of the atoms, but can otherwise be bounded independently of N. It is important to notice that the new localized wave functions span the same space as the one we started with. Since the new localized wave functions span the same subspace as the original ones, their Kohn-Sham energy is the same.

There are several additional important technical components of our localized subspace-iteration algorithm that need to be explained:

(1) The filtering procedure, which is an improvement over the power method,

(2) the computation of the Fermi energy, necessary to define an effective filter, and

(3) the computation of $\mathbf{S}^{-1}$.

We now explain these components in detail.

C. Filtering

The power method, and its variants, can be understood as a simple filtering strategy by which a certain portion of the spectrum of $\mathbf{H}$ is amplified more than others. In view of its convergence rate (16), it is clear that one can construct a more efficient filter by applying a subspace iteration to $P(\mathbf{H})$, where P is a polynomial of degree n. Assuming that the polynomial P splits the spectrum of $\mathbf{H}$, in the sense that

$$P(\lambda_i)\leq P(\lambda_M),\quad i=1,\ldots,M, \tag{20}$$

$$P(\lambda_j)\geq P(\lambda_{M+1}),\quad j=M+2,\ldots,M, \tag{21}$$

the rate of convergence of the polynomial filtered subspace iteration is[29]

$$\kappa=\left|\frac{P(\lambda_M)}{P(\lambda_{M+1})}\right|. \tag{22}$$

By a judicious choice of P, we can get $\kappa\ll\tau$ [see Eq. (16)], thereby accelerating the convergence of the subspace iteration considerably. It can be shown that the optimal choice for P is the Chebyshev polynomial of degree n, T_n.

The Chebyshev polynomial of degree n is defined as

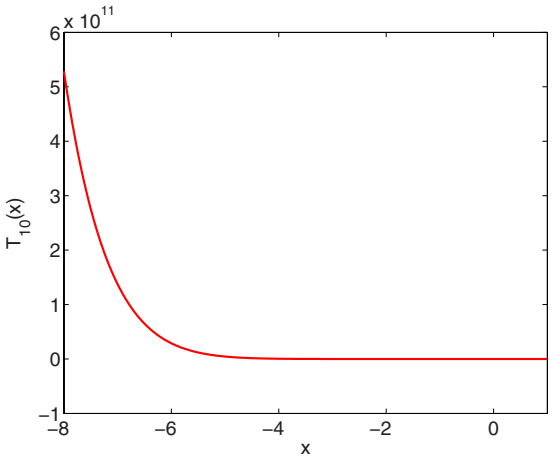

FIG. 1. (Color online) Chebyshev polynomial of degree 10, depicting its fast growth outside the interval $[-1,1]$.

$$T_n(x) = \begin{cases} \cos(n \cos^{-1} x) & \text{if} |x| \leq 1, \\ (-1)^n \cosh(n \cosh^{-1}|x|) & \text{if} |x| \geq 1. \end{cases} \tag{23}$$

The Chebyshev polynomials have the property that $|T_n(x)| \leq 1$ for $x \in [-1,1]$ and $|T_n(x)| \gg 1$ for $|x| > 1$. In addition, they satisfy the three-term recursion

$$T_0(x) = 1; T_1(x) = x; T_{k+1} = 2xT_k(x) - T_{k-1}(x), k \geq 1, \tag{24}$$

which can be used to evaluate $T_n(\mathbf{H})$ efficiently without explicitly computing the operator. The Chebyshev polynomial T_{10} is plotted in Fig. 1, to illustrate the fast growth outside $[-1,1]$.

To use the Chebyshev polynomial as an effective filter, the spectrum of the Hamiltonian is divided into $\lambda_{\min} < \varepsilon_F < \lambda_{\max}$. The map $x \to 2\frac{x-\varepsilon_F}{\lambda_{\max}-\varepsilon_F} - 1$ transforms the interval $[\varepsilon_F, \lambda_{\max}]$ into $[-1,1]$, and the interval $[\lambda_{\min}, \varepsilon_F]$ is mapped into a subinterval of $(-\infty, -1]$. Therefore the rescaled Chebyshev polynomial

$$T_n\left(2\frac{x-\varepsilon_F}{\lambda_{\max}-\varepsilon_F} - 1\right)$$

amplifies the lower end of the spectrum, while leaving the high end unchanged in magnitude, effectively filtering out the high end of the spectrum.

In the context of electronic structure analysis, subspace iteration with Chebyshev filtering has been used by Zhou *et al.*[20] In our algorithm, linear scaling is achieved by replacing the diagonalization and orthogonalization step in the subspace-iteration algorithm introduced in Ref. 20 by the localization procedure described in Sec. II B.

D. Computation of the Fermi energy

The Fermi energy ε_F is necessary in our algorithm in order to define an effective Chebyshev filter, as described in Sec. II C. Given that we have no knowledge of the spectrum, this is not as straightforward as in codes that involve diagonalization.[20] We estimate the Fermi energy by the maximum eigenvalue of the Ritz matrix. This eigenvalue can be approximated with linear complexity in the following way. Consider a basis of the subspace V, represented by $\Psi = [\psi_1, \ldots, \psi_N]$. We know that $\Psi \mathbf{S}^{-1/2}$ is an orthonormal basis, where the Gram matrix $\mathbf{S}$ is defined in Eq. (8). The Ritz matrix $\mathbf{R} \in \mathbb{R}^{N \times N}$ is defined as

$$\mathbf{R} = (\Psi \mathbf{S}^{-1/2})^T \mathbf{H} \Psi \mathbf{S}^{-1/2} = \mathbf{S}^{-1/2} \Psi^T \mathbf{H} \Psi \mathbf{S}^{-1/2}. \tag{25}$$

We are interested in the eigenvalues of $\mathbf{R}$, which are the same as the eigenvalues of $\mathbf{S}^{1/2}\mathbf{R}\mathbf{S}^{-1/2} = \Psi^T \mathbf{H} \Psi \mathbf{S}^{-1}$. The Fermi energy is then approximated as the maximum eigenvalue of $\mathbf{R}$, which can be approximated with a few steps of the power method. Note that the power method will converge to the maximum eigenvalue in magnitude. Therefore, to approximate the Fermi energy, we compute both the maximum and minimum eigenvalues in magnitude of the Ritz matrix, respectively, λ_M and λ_m. Once the maximum eigenvalue is approximated, the shifted power method can be used to obtain the minimum eigenvalue (in magnitude). Note that in the application of the power method, we do not need $\mathbf{S}^{-1}$ explicitly, but only $\mathbf{S}^{-1}\mathbf{v}$ for some vector $\mathbf{v}$. This can be obtained by solving the system of equations $\mathbf{S}\mathbf{y} = \mathbf{v}$. Assuming that the condition number of $\mathbf{S}$ is bounded independently of N and given that $\mathbf{S}$ is sparse, the system of equations can be solved with the preconditioned conjugate gradient, using the main tridiagonal part of $\mathbf{S}$ as a preconditioner.[31] Since the basis Ψ is localized, the Ritz matrix is sparse and the required matrix-vector multiplications can be performed with linear scaling.

E. Computation of $\mathbf{S}^{-1}$

The inverse of the Gram matrix $\mathbf{S}$, defined in Eq. (8), plays a fundamental role in the nonorthogonal formulation of Kohn-Sham DFT, as it appears in the energy functional itself (9), the definition of the electron density (10), and the pseudopotential (11). A direct computation of $\mathbf{S}^{-1}$ scales as $O(N^3)$. However, under the assumption that both $\mathbf{S}$ and $\mathbf{S}^{-1}$ decay exponentially fast away from the diagonal, $\mathbf{S}^{-1}$ can be computed with $O(N)$ complexity. We have implemented the scaled third-order Newton-Schulz iteration described in Ref. 32 for the inverse matrix square root. Given the Gram matrix $\mathbf{S}$, the iteration solves equation

$$\mathbf{ZSZ} = \mathbf{I}. \tag{26}$$

To compute $\mathbf{Z} = \mathbf{S}^{-1/2}$ in the following way:

(1) Let $\lambda = \frac{1}{\|\mathbf{S}\|_1}$, $\mathbf{Z}_0 = \mathbf{I}$, $\mathbf{Y}_0 = \mathbf{S}$, and $k = 0$;

(2) Set $\mathbf{X}_k = \lambda \mathbf{Y}_k \mathbf{Z}_k$, $\mathbf{T}_k = \frac{1}{8}(15\mathbf{I} - 10\mathbf{X}_k + 3\mathbf{X}_k^2)$, $\mathbf{Z}_{k+1} = \mathbf{Z}_k \mathbf{T}_k$, and $\mathbf{Y}_{k+1} = \mathbf{T}_k \mathbf{Y}_k$.

(3) Repeat the previous step until $\|\mathbf{Z}_{k+1} - \mathbf{Z}_k\|_1 \leq$ tolerance.

(4) $\mathbf{S}^{-1/2} = \lambda^{1/2} \mathbf{Z}_k$ and $\mathbf{S}^{-1} = \lambda \mathbf{Z}_k^2$.

All operations are carried out using sparse matrix computations. The truncation of the successive matrices is done using the algorithm described in Ref. 33. The main idea in this approach is to control the l_1-norm error at each step by discarding the blocks whose contribution to the next iterate is below a given tolerance.

F. LSI Algorithm for Kohn-Sham DFT

The algorithm for Kohn-Sham DFT, based on a Chebyshev-filtered subspace iteration with optimally localized nonorthogonal wave functions is:[18]

Algorithm 1. Localized subspace iteration for Kohn-Sham DFT.

(1) Given (localized) wave functions Ψ_0.
(2) repeat ([*self-consistency loop (SCF)*])
(3) Compute the electron density, ρ.
(4) Compute the effective potential, $V_{\text{eff}}[\rho]$.
(5) repeat [(*localized subspace iteration*)]
(6) Filtering step: $\Phi = T_n(H)\Psi$.
(7) Localization Step: Localize ψ_r for $r=1,\cdots,N$.
(8) Truncate beyond the cutoff radius.
(9) Until convergence of the linear iteration.
(10) Update the electron density (mixing).
(11) Until Convergence of the self-consistent iteration.

III. PERFORMANCE EVALUATION USING A ONE-DIMENSIONAL MODEL

The original Kohn-Sham problem is a rather complicated problem in three dimensions. In order to better understand the essential features of the proposed algorithm, we will first study its behavior on a much simplified one-dimensional model. This one-dimensional model shares the most important features of a linearized Kohn-Sham model. It can be used as a test bed for systematically analyzing the accuracy and convergence of DFT algorithms.

A. One-dimensional model

We consider an infinite array of atoms on a line with unit spacing: $X_i=i$, for $i \in \mathbb{Z}$. Each atom has one valence electron and we ignore spin degeneracy. The electrons are noninteracting, so that the electronic structure of the system is determined by solving linear eigenvalue problems (instead of nonlinear eigenvalue problems as in the full Kohn-Sham case)

$$H\psi_i = \epsilon_i \psi_i, \tag{27}$$

where the Hamiltonian is given by

$$H = -\frac{1}{2}\frac{d^2}{dx^2} + V(x). \tag{28}$$

The effective potential V is a sum of Gaussian wells located at the atom sites

$$V(x) = -\sum_{i\in\mathbb{Z}} \frac{a}{\sqrt{2\pi\sigma^2}} \exp[-(x-X_i)^2/2\sigma^2]. \tag{29}$$

This model has two parameters: a, which characterizes the depth of the wells, and σ, which characterizes its width.

The band structure of this model is shown in Fig. 2 for several sets of parameter values. The first band is fully occupied and the second band is empty. The band gap E_{gap} is the difference between the highest eigenvalue in the first band and the lowest eigenvalue in the second band [both obtained at the edge of the first Brillouin zone $k=\pi$ (Ref. 34)]. We are also interested in the ratio between the band gap and the width of the first band E_{width}. The relation between the band gap and the parameters a and σ of the potential is shown in Fig. 3. It is observed that the gap is proportional to $\sqrt{a}/\sigma$. By changing parameters, we may change the model from a well-gapped insulator to a metal-like system. We are going to test the performance of LSI algorithms in different situations.

B. Convergence issue of the LSI and related algorithms

Recall that each LSI iteration contains three steps: filtering, localization, and truncation. Starting from a given subspace, the filtering step makes the subspace closer to the occupied subspace by filtering out the higher spectrum. The localization step finds a better representation while keeping the subspace unchanged. It is clear that without truncation, the LSI iteration will converge and will give the correct occupied subspace of the Hamiltonian, like other subspace-iteration algorithms.

However, truncation makes things much more complicated since the subspace deviates from the correct occupied subspace after truncation. As a result, the iteration process might not converge, as illustrated in Fig. 4. This is a common feature for linear-scaling algorithms involving truncation. For instance, for the orbital minimization algorithm,[14] the truncation step in general increases the energy while the minimization step in general decreases the energy. Hence, even if a variational algorithm is used, it is not guaranteed that it will converge to the minimizer (not to mention other issues, such as the issue of local minimizers in the orbital minimization algorithm[14]). The general behavior of the trajectory of the energy in the LSI iteration is shown in Fig. 4.

In Figs. 4(a) and 4(b), the difference between the energy at each step and the accurate result obtained by direct diagonalization is shown in a semilog plot. In Fig. 4(b), we also plot the energy before the truncation step. It is obvious that the truncation will increase the energy since the subspace deviates from the occupied subspace after truncation. It can be seen that the error first decays exponentially and then starts to fluctuate around a value that is small but different from the round-off error. In Figs. 4(c) and 4(d), the fluctuations are shown for two different cases. We will call this behavior *remanent fluctuation* and we will refer to the error between the numerical solution and the true minimizer (without truncation) *remanent error*. The remanent error is caused by the truncation of the wave functions.

We will quantify the size of the remanent fluctuation by its variance between iterations N_0 and N_0+N_s, defined as

$$(\Delta E)^2 = \frac{1}{N_s-1}\sum_{i=N_0+1}^{N_0+N_s} (E_i - \bar{E})^2, \tag{30}$$

where $\bar{E}$ is the mean value of $\{E_i\}$, $i=N_0+1,\cdots,N_0+N_s$. The remanent error in the energy is then the difference between $\bar{E}$ and the true energy of the system. Obviously, ΔE depends on N_0 and N_s. In the following, we choose $N_s=1000$ and N_0 sufficiently large such that the iteration has already started to fluctuate before step N_0. We calculate the variance ΔE for

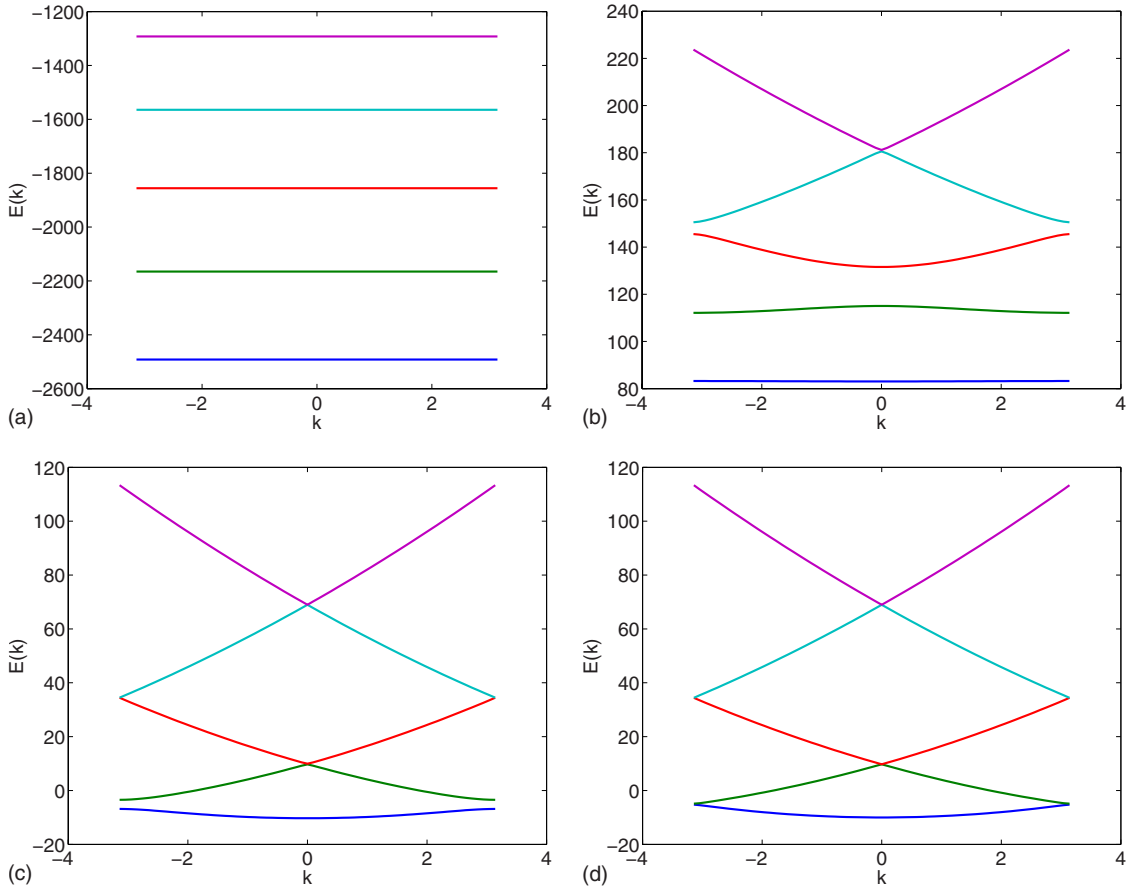

FIG. 2. (Color online) Band structure for different values of σ and a. Choosing a and σ appropriately we can control the size of the energy-band gap. In figures (a) $a=1000$; $\sigma=0.15$, (b) $a=100$; $\sigma=0.3$, (c) $a=10$; $\sigma=0.3$, we show insulator systems with different size gaps. The gap in the system shown in figure (d) ($a=10$; $\sigma=0.45$) is almost zero, closer to what occurs in a metal.

different parameters (band gap, cutoff radius R_c, order of Chebyshev filter N_{Cheb}) to try to understand the effects of these parameters on the dynamical behavior of LSI iteration. The results are collected in Table I. It is observed that the variance decreases when either the cutoff radius, the band gap, or the order of the Chebyshev filter is increased. The effect of the cutoff radius is the most significant in the three cases considered. Indeed, the variance decreases exponentially as the cutoff radius is increased.

We may understand the behavior of the LSI and other related algorithms that include a truncation step by considering a local error estimate. Let us denote by V_0 the true occupied subspace of the Hamiltonian and by V_n the subspace in the nth step. The error is the difference $d(V_n, V_0)$, where d is some choice of subspace distance (for example, the two norms of the difference in the projection matrix). We use F, L, and T to represent abstractly the filtering, localization, and truncation steps, respectively. Define the truncation error δ as

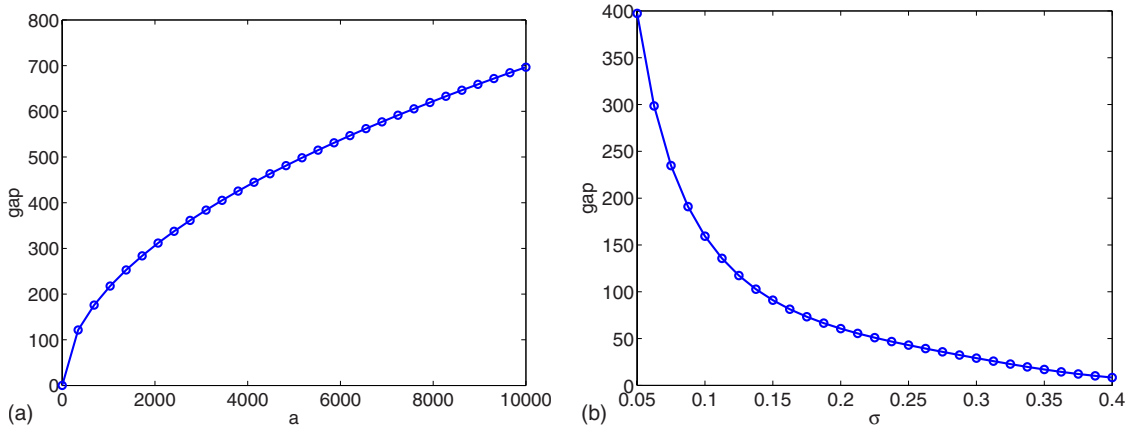

FIG. 3. (Color online) Band gap as function of a and σ. (a) Fixed σ, a varies from 0 to 10 000; the gap scales like $\sqrt{a}$. (b) Fixed a, σ varies from 0.05 to 0.4; the gap scales like $1/\sigma$.

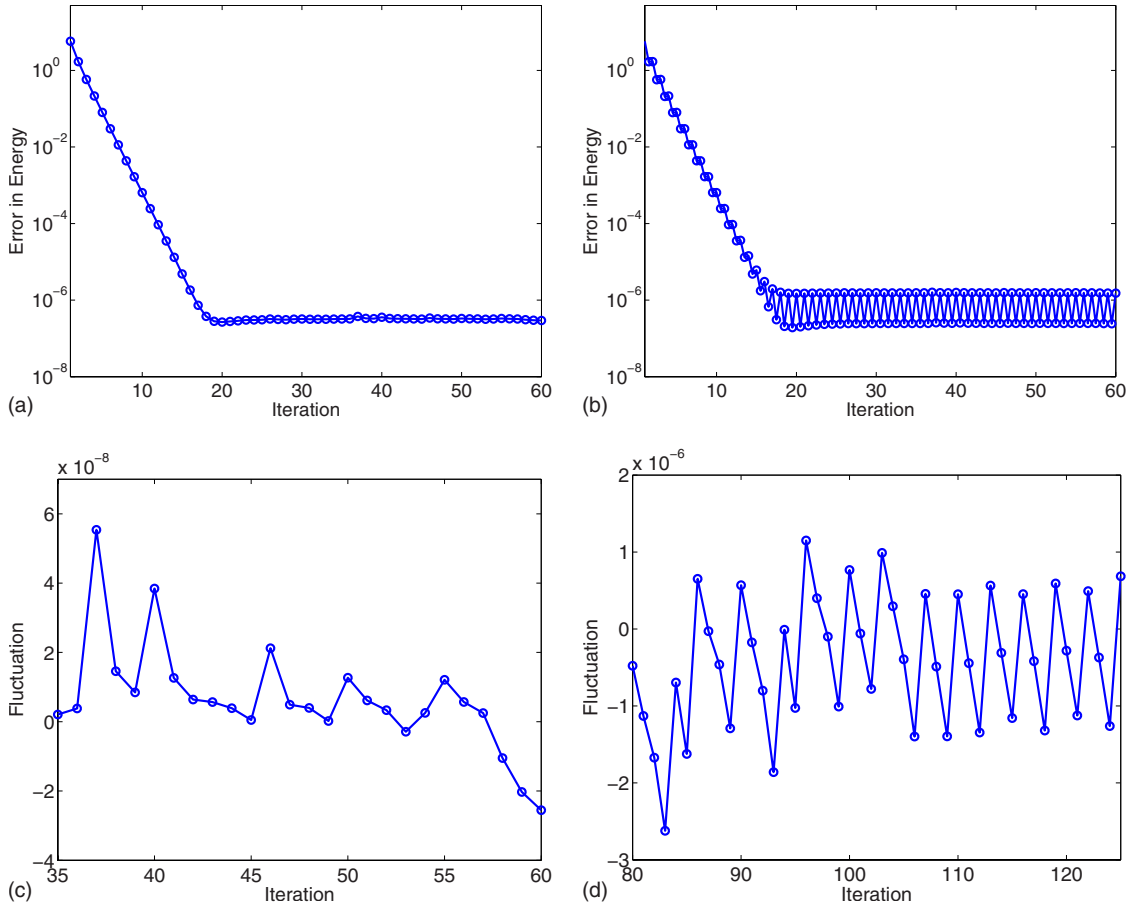

FIG. 4. (Color online) Trajectory of the energy in the LSI iteration. (a) Logarithmic plot of the error in the energy for each step in the LSI iteration ($a=50, \sigma=0.3, R_c=2.0$). (b) Logarithmic plot of the error in the energy in the LSI iteration ($a=50, \sigma=0.3, R_c=2.0$); for each step, the errors before and after the truncation of the wave functions are both shown. (c) Remanent fluctuation in the energy ($a=50, \sigma=0.3, R_c=2.0$). (d) Remanent fluctuation in the energy ($a=10, \sigma=0.4, R_c=2.0$).

TABLE I. The relationship between the size of remanent fluctuation (quantified by its variance) and the parameters in the LSI iteration.

a	σ	R_c	N_{Cheb}	ΔE
10	0.2	1.0	20	1.2246×10^{-5}
10	0.2	2.0	20	2.1078×10^{-7}
10	0.2	3.0	20	5.0868×10^{-11}
10	0.3	1.0	20	2.0360×10^{-5}
10	0.3	2.0	20	4.6898×10^{-7}
10	0.3	3.0	20	2.0813×10^{-11}
10	0.4	1.0	20	3.3369×10^{-5}
10	0.4	2.0	20	1.1842×10^{-6}
10	0.4	3.0	20	2.6947×10^{-10}
25	0.3	1.0	20	1.2325×10^{-5}
25	0.3	1.0	12	1.2639×10^{-5}
25	0.3	1.0	6	6.3202×10^{-5}
25	0.3	1.0	4	3.3415×10^{-4}

$$\delta = d(TL(V_0), V_0). \tag{31}$$

So δ quantifies the error caused by truncating the localized representation of the true occupied subspace. δ is determined by the localization property of the Hamiltonian, and hence, δ is decreased as the band gap is enlarged. Since $d[TL(\cdot),\cdot]$ is a continuous function of the subspace, we have

$$d[TL(V), V] \leq C\delta \tag{32}$$

if subspace V is sufficiently close to V_0, say if $d(V, V_0) < M$.

For the filtering step, the error is reduced as in traditional subspace-iteration methods

$$d[F(V_n), V_0] \leq \lambda d(V_n, V_0) = \lambda e_n, \tag{33}$$

where $\lambda < 1$ is the amplification factor determined by the band gap and order of the Chebyshev polynomial used. Now assume that we start from a subspace V_n that is inside a small neighborhood of V_0. Then, by Eq. (32), we have

$$\begin{aligned} e_{n+1} &= d[TLF(V_n), V_0] \leq d[TLF(V_n), F(V_n)] + d[F(V_n), V_0] \\ &\leq \lambda e_n + C\delta. \end{aligned} \tag{34}$$

Here, the estimate (34) is valid when the current subspace is

TABLE II. Comparison between the results from the LSI and direct diagonalization.

	Energy	Energy (LSI)	Error	Error in the density
Case 1	−2492.72	−2492.72	0.00%	3.5315×10^{-8}
Case 2	−8.365 27	−8.35749	0.09%	1.1488×10^{-2}
Case 3	−8.882 99	−8.88282	0.02%	3.9734×10^{-5}

close enough to the true occupied subspace. It is a local result. Note that, without truncation, the true occupied subspace is an attractor for the subspace iteration. With truncation, it is in general not even a fixed point. The estimate (34), however, guarantees that the numerical results stay within a neighborhood of the true occupied space if they come close enough to it. The size of the neighborhood, which is clearly related to the variance of the remanent fluctuation, can be estimated from Eq. (34) and it is given by

$$e_\infty \leq C\delta/(1-\lambda). \tag{35}$$

Note that δ decreases exponentially as the cutoff radius increases and λ becomes smaller as the band gap or the order of the filter increases. At least qualitatively, this estimate agrees well with the numerical results. We will leave the detailed discussion of the dynamical behavior of LSI and other related algorithms to future publications.[26]

It is clear that, in practice, the LSI iteration should be terminated when the remanent fluctuation is reached. When the specified error tolerance is smaller than the remanent error, the iteration is terminated before the remanent fluctuation is reached. Later, when we talk about the accuracy and convergence of LSI, it should be understood in this sense.

C. Performance of the LSI for the one-dimensional model

We will focus on three sets of parameter values:

(1) $a=1000$ and $\sigma=0.15$. The system is an insulator with a fairly large band gap ($E_{\text{gap}}=323.974$ and $E_{\text{gap}}/E_{\text{width}}=3193.766$). The potential is quite narrowly peaked in this case.

(2) $a=10$ and $\sigma=0.45$. The system is close to being a metal in this case, with only a tiny gap between the first and second bands ($E_{\text{gap}}=0.731\,622$ and $E_{\text{gap}}/E_{\text{width}}=0.159\,585$). The external potential is almost flat in this case.

(3) $a=10$ and $\sigma=0.3$. This is an intermediate situation. There is a clear band gap but it is not so large ($E_{\text{gap}}=3.438\,81$ and $E_{\text{gap}}/E_{\text{width}}=0.998\,883$). The system can be regarded as a semiconductor.

Since the system is periodic, for the exact solution we will use k-point sampling with 64 k points in the first Brillouin zone $[-\pi,\pi]$. We use a second-order finite difference discretization with grid size 1/64.

First let us study the accuracy of the LSI. We choose a cutoff radius $R_c=4.0$ (recall that the lattice constant is 1). The results are shown in Table II. In the table, for each case, the energy calculated from direct diagonalization and from the LSI algorithm, the relative error in the energy, and the error in the electronic density in the L^1 norm are listed. For the insulator (case 1) and the semiconductor (case 3), the results from the LSI algorithm are very accurate. For the metal-like system (case 2), since the gap is quite small, the performance of linear-scaling algorithms is expected to deteriorate, but the LSI algorithm still gives results with fairly good accuracy.

As noted above, the cutoff radius in the results presented in Table II was fixed at 4.0. We now study the effect that the cutoff radius has on the accuracy of the LSI. Naturally we expect that the larger the cutoff radius used, the more accurate the result will be. As an illustration, we take the parameter values $a=100$ and $\sigma=0.3$ with a cutoff radius that varies from 0.1 to 4.0 lattice constants. The accuracy of the energy is shown in a logarithmic scale in Fig. 5(a). It is observed that as the cutoff radius increases, the error decreases exponentially before reaching the scale of 10^{-8}, which is the stopping criterion used in the iteration for remanent fluctuation. This agrees well with the expected exponential decaying property of the Wannier function for insulators. It is also of interest to see how the computed wave functions change as the cutoff radius is changed. This can be seen in Fig. 5(b), where we show the wave functions resulting from the LSI with cutoff radii $R_c=0.2$, 0.6, and 1.0. With a small cutoff radius $R_c=0.2$, the result is quite far from the correct result.

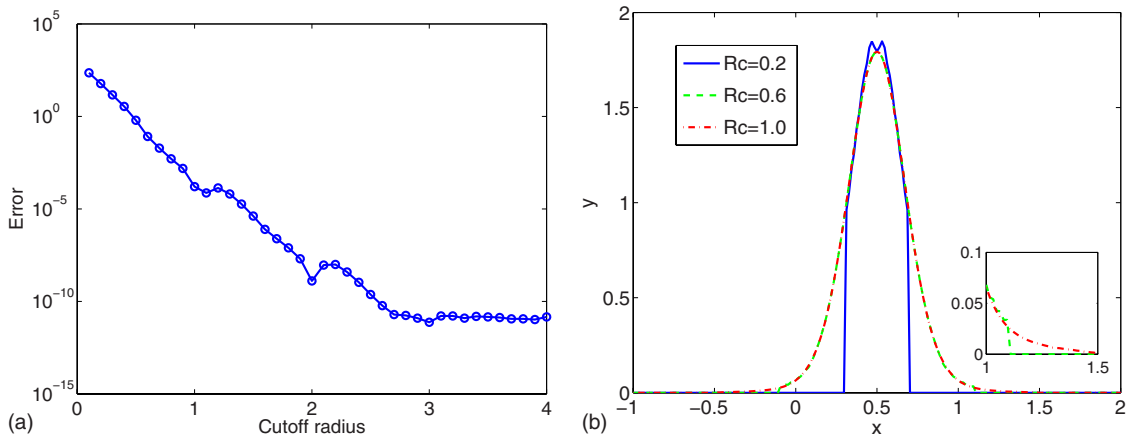

FIG. 5. (Color online) The effect of the cutoff radius on the accuracy for the insulator case. (a) Logarithmic plot of the error in the energy for different cutoff radii. (b) The resulting wave functions for cutoff radii 0.2, 0.6, and 1.0.

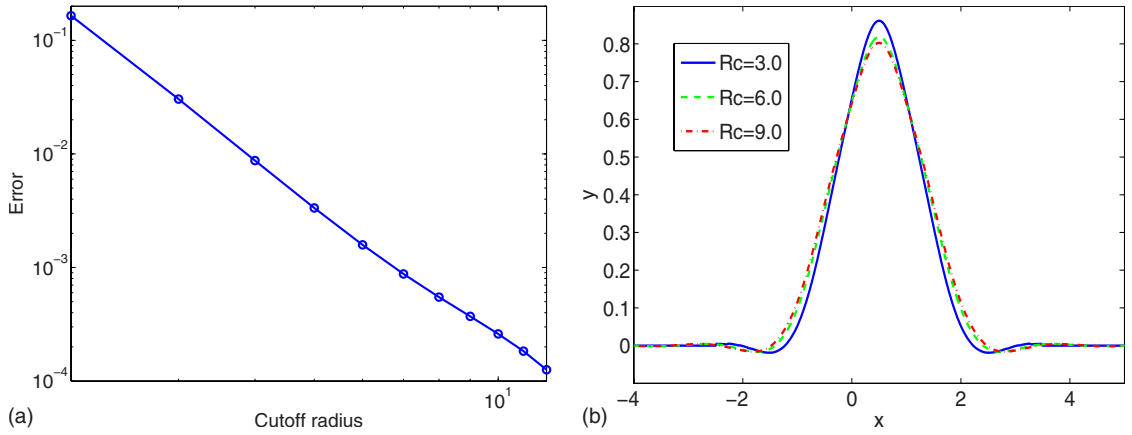

FIG. 6. (Color online) The effect of the cutoff radius on the accuracy for the metallic case. (a) Logarithmic plot of the error in the energy for different cutoff radii. (b) The resulting wave functions for cutoff radii 3.0, 6.0, and 9.0.

For $R_c=0.6$ and 1.0, however, the wave functions are only notably different from each other at the edge of the cutoff radius. The wave function obtained using an even larger cutoff radius is almost identical to the one obtained with $R_c=1.0$ and is therefore not included in the figure.

More interesting is the case when the system is metal-like, since this will help us understand how the LSI fails. Therefore we will focus next on the metal-like case when $a=10$ and $\sigma=0.45$. In Fig. 6(a), the error in energy is shown in a log-log plot as the cutoff radius is varied from 2.0 to 13.0. It is observed that the error decays algebraically as the cutoff radius increases.[35] The resulting wave functions corresponding to cutoff radii 3, 6, and 9 are shown in Fig. 6. The wave function expands and becomes closer to the correct solution as the cutoff radius grows. Note also that a larger cutoff radius is required to obtain accurate results for metal-like systems than for insulator systems ($R_c=1.0$ is enough for the previous insulator case).

IV. NUMERICAL EXAMPLES: VALIDATION AND LINEAR SCALING

We have implemented the LSI algorithm for Kohn-Sham DFT in three dimensions. To achieve linear scaling, it is important to use a sparse representation for the wave functions, as well as the pseudopotential. Both the wave functions and the pseudopotential components are compactly supported, and therefore once the grid size is fixed, each array contains only a small fraction of grid values.

Consider the domain $\Omega=[0,L]\times[0,D]\times[0,H]$, containing the support of all the wave functions. This domain contains also the support of all the pseudopotential components (local and nonlocal), as well as the support of the atomic function. We discretize Ω using a uniform mesh with grid sizes $\Delta x=L/n_x$, $\Delta y=D/n_y$, and $\Delta z=H/n_z$. We define the values of the wave functions and the density at the center of the cells $\psi_{i,j,k}\approx\psi(x_i,y_j,z_k)$, where

$$x_i=\left(i-\frac{1}{2}\right)\Delta x,\quad i=0,\ldots,n_x+1,$$

$$y_j=\left(j-\frac{1}{2}\right)\Delta y,\quad j=0,\ldots,n_y+1,$$

$$z_k=\left(k-\frac{1}{2}\right)\Delta z,\quad k=0,\ldots,n_z+1. \tag{36}$$

The points with indices 0 or n_x+1, n_y+1, and n_z+1 are ghost cells outside the computational domain and are used only to impose the boundary conditions. Since the wave functions are zero on the boundary, we define the ghost values by reflection, which on the $x=0$ boundary would be $\psi_{0jk}=-\psi_{1jk}$. A similar expression is used on the other boundaries.

The Laplacian is discretized using second-order centered differences with the standard seven-point stencil. If we define

$$\delta_{xx}u_{ijk}=\frac{u_{i+1,j,k}-2u_{i,j,k}+u_{i-1,j,k}}{(\Delta x)^2},$$

$$\delta_{yy}u_{ijk}=\frac{u_{i,j+1,k}-2u_{i,j,k}+u_{i,j-1,k}}{(\Delta y)^2},$$

$$\delta_{zz}u_{ijk}=\frac{u_{i,j,k+1}-2u_{i,j,k}+u_{i,j,k-1}}{(\Delta z)^2},$$

the Laplacian is

$$\Delta u_{ijk}=\delta_{xx}u_{ijk}+\delta_{yy}u_{ijk}+\delta_{zz}u_{ijk}. \tag{37}$$

The Coulomb term is approximated as a discrete convolution and this convolution is computed using the fast Fourier transform (FFT) with zero padding as described in Refs. 18 and 36.

To validate our code, we have compared the result obtained with our algorithm to the results obtained using the package PARSEC.[37,38] We considered two alkane chains $CH_3(CH_2)_nCH_3$, with $n=0$ (ethane) and with $n=10$. The geometric structure of the alkanes was determined using the geometry optimization option in the GAUSSIAN package[39] using the Gaussian-type orbitals STO-3G. The molecule so obtained for $n=10$ is plotted in Fig. 7(a) to illustrate the geometric structure. Although the molecule is a linear chain,

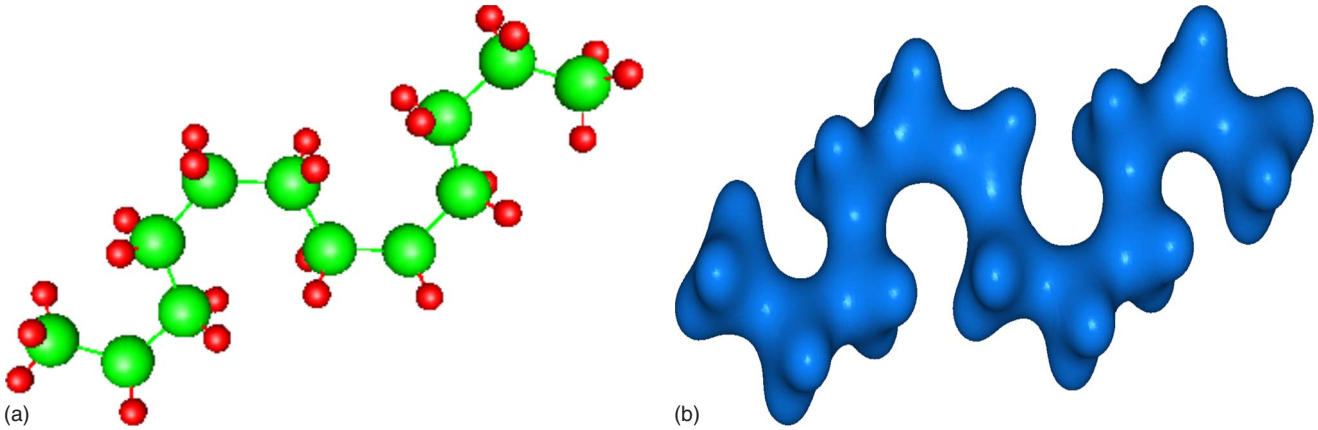

FIG. 7. (Color online) Electron density of an alkane chain obtained with the LSI algorithm. (a) Atomic configuration, obtained using the package GAUSSIAN (Ref. 39). (b) Electron density.

there is a nontrivial three-dimensional structure due to the hydrogen atoms. The cutoff radius for the wave functions was chosen to be twice the length of the smallest bond length.

We used the Troullier-Martins pseudopotentials in the Kleinman-Bylander form.[23,24] For hydrogen, the pseudopotential has one local component for the $1s$ orbital and no nonlocal components; for carbon, however, we need to consider the $2s$ and $2p$ nonlocal components of the pseudopotential. We choose the $2p$ pseudopotential to be the local component in order to avoid using spherical harmonics. The nonlocal pseudopotential is therefore the $2s$ component [so $l_{\max}=0$ in Eq. (4)]. We generate the pseudopotentials using the code developed by Paolo Gianozzi and his collaborators.[40] For the exchange and correlation term in the energy, we consider the homogeneous electron-gas approximation of Ceperley and Alder,[41] as parameterized by Perdew and Zunger.[42]

The pseudopotential operator acts on the wave function ψ in the following way:

$$\hat{V}_{\mathrm{PS}}\psi(\boldsymbol{x}) = \sum_{j=1}^{N_a} [V_{\mathrm{local}}(\boldsymbol{x}-\boldsymbol{R}_j)\psi(\boldsymbol{x}) + \alpha_j \beta_{00}^j(\boldsymbol{x}-\boldsymbol{R}_j)], \quad (38)$$

where we have defined

$$\alpha_j = \int_{\mathbb{R}^3} \beta_{00}^j(\boldsymbol{y}-\boldsymbol{R}_j)\psi(\boldsymbol{y})d\boldsymbol{y}, \quad j=1,\ldots,N_a. \quad (39)$$

Both $\beta_{00}^j(\cdot-\boldsymbol{R}_j)$ and ψ are compactly supported, so the integrals only need to be carried out if the supports of $\beta_{00}^j(\cdot-\boldsymbol{R}_j)$ and ψ intersect. The integral is approximated using the midpoint rule[43]

$$\alpha_j = \Delta V \sum_{klm} \beta_{00}^j(\boldsymbol{y}_{klm}-\boldsymbol{R}_j)\psi_{klm}, \quad (40)$$

where $\Delta V=\Delta x\Delta y\Delta z$ is the volume of each computational cell. The wave functions are centered at the center of the chemical bonds and initially they are chosen to be Gaussians. We also considered the case of random initial wave functions and the code converged as well, although the number of iterations was higher. The self-consistent iteration was terminated when the difference between two consecutive electron densities, measured in the L_2 norm, was less than 10^{-5}. In principle, the filtering and localization steps must be performed until the linear iteration has converged, as described in *Algorithm 1*. In practice, however, we only perform these steps once before updating the electron density and the effective potential.

For validation purposes, we have compared the energies obtained with the LSI code and the corresponding energies obtained with PARSEC. The result of the comparison is shown in Table III. As can be seen, the difference is of the order of 2%–3%. This discrepancy is due to the differences in the formulation and implementation of the two codes, such as domain size, discretization, evaluation of the Coulomb potential, etc., and is of no consequence for the purposes of this article. The PARSEC code was used only as a reference to validate our own. The electron density obtained for $N=38$ atoms is shown in Fig. 7.

To illustrate the linear-scaling behavior of our code, we have collected the timings for the filtering, localization, inversion of the Gram matrix, and computation of the Fermi

TABLE III. Comparison between our code and PARSEC. The ionic energy E_{ion} includes the Coulomb and pseudopotential energies. All energies are measured in Hartrees. The results of the LSI are in good agreement with the results from PARSEC.

Comparison with PARSEC			
	C_2H_6 ; $N=8$		
	E_{ion}	E_{XC}	E_{KS}
E_{LSI}	−5.626 058 47	−22.157 325 46	−14.382 176 32
E_{PARSEC}	−5.538 702 40	−22.351 239 17	−14.893 958 80
E_{diff}	1.58%	0.87%	3.43%
	$CH_3(CH_2)_{10}CH_2$; $N=38$		
	E_{ion}	E_{XC}	E_{KS}
E_{LSI}	−30.206 844 90	−124.296 984 11	−81.752 116 08
E_{PARSEC}	−29.850 574 5	−125.827 971 34	−83.544 015 155
E_{diff}	1.19%	1.21%	2.14%

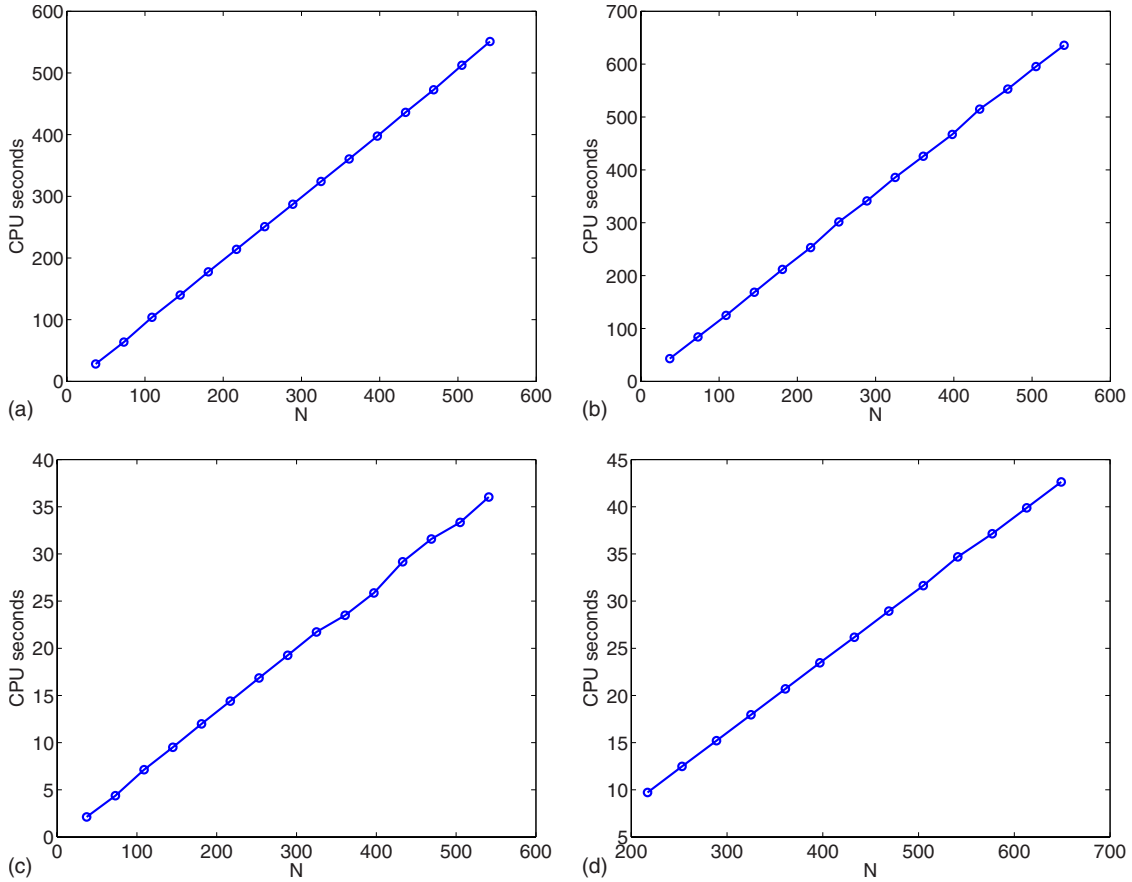

FIG. 8. (Color online) Timings obtained with the LSI code. Linear scaling is observed. (a) Localization, (b) Chebyshev filter; $n=10$, (c) Fermi energy, (d) inverse of the Gram matrix.

energy for alkane chains with sizes ranging between 38 and 542 atoms. The results are shown in Fig. 8.

We have also computed the error in the approximation of $\mathbf{S}^{-1}$ in the Frobenius norm. The results are shown in Table IV. The tolerance was set to $\varepsilon=10^{-5}$. The error appears to be independent of the size of the molecule. A reduction in the error can be achieved by simply reducing the tolerance.

V. CONCLUSION

We have described a linear-scaling algorithm for Kohn-Sham DFT, the LSI. The algorithm is based on the nonorthogonal formulation of the Kohn-Sham DFT and consists of a Chebyshev-accelerated subspace iteration with optimally localized wave functions. The use of the nonorthogonal formulation is important, at least for two reasons: it allows us to bypass the orthogonalization step and it allows us to use optimally localized wave functions. We have also seen that the added difficulty associated with the nontrivial Gram matrix of the nonorthogonal wave functions can be overcome rather easily.

We have presented and analyzed a controllable one-dimensional model with which we can reproduce the main characteristics of insulating and metallic systems. This model has been used to illustrate the convergence properties of the LSI and we plan to use it as a benchmark for the analysis of other methodologies.

We have also implemented the LSI in a realistic three-dimensional setting and tested its accuracy, as well as its linear-scaling properties. All together, these demonstrate that

TABLE IV. Error in the computation of the inverse of the Gram matrix as a function of molecular size. The errors appear to be independent of the size of the system in this case.

N	$\|\mathbf{S}^{-1}-\widetilde{\mathbf{S}^{-1}}\|_F$
145	$2.850\,907\times10^{-5}$
181	$2.885\,210\times10^{-5}$
217	$2.849\,086\times10^{-5}$
253	$2.851\,860\times10^{-5}$
289	$2.850\,546\times10^{-5}$
325	$3.030\,061\times10^{-5}$
361	$3.122\,993\times10^{-5}$
397	$3.152\,983\times10^{-5}$
433	$3.054\,400\times10^{-5}$
469	$3.226\,956\times10^{-5}$
505	$3.301\,841\times10^{-5}$
541	$3.279\,986\times10^{-5}$
577	$3.348\,167\times10^{-5}$
649	$3.298\,259\times10^{-5}$

the LSI is a very promising tool for the electronic structure analysis of insulators.

ACKNOWLEDGMENTS

The authors would like to thank Weiguo Gao and Emil Prodan for interesting discussions. The work of C.J.G.-C. is supported by and NSF CAREER (Grant No. DMS-0645766). The work of W.E, J.L., and Y.X. is supported in part by ONR under Grant No. N00014-01-1-0674, DOE under Grant No. DOE DE-FG02-03ER25587, and NSF under Grants No. DMS-0407866 and No. DMR-0611562.

*cgarcia@math.ucsb.edu
†jianfeng@math.princeton.edu
‡weinan@princeton.edu
§weinan@princeton.edu

[1] W. Kohn and L. Sham, Phys. Rev. **140**, A1133 (1965).
[2] R. Parr and W. Yang, *Density-Functional Theory of Atoms and Molecules, International Series of Monographs on Chemistry* (Oxford University Press, New York, 1989).
[3] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964).
[4] W. Kohn, Rev. Mod. Phys. **71**, 1253 (1999).
[5] S. Goedecker, Rev. Mod. Phys. **71**, 1085 (1999).
[6] W. Kohn, Phys. Rev. Lett. **76**, 3168 (1996).
[7] W. Yang, J. Mol. Struct.: THEOCHEM **255**, 461 (1992).
[8] M. Barrault, E. Cancès, W. W. Hager, and C. L. Bris, J. Comput. Phys. **222**, 86 (2007).
[9] X.-P. Li, R. W. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10891 (1993).
[10] P. D. Haynes, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne, Phys. Status Solidi B **243**, 2489 (2006).
[11] R. McWeeny, Rev. Mod. Phys. **32**, 335 (1960).
[12] G. Galli and M. Parrinello, Phys. Rev. Lett. **69**, 3547 (1992).
[13] F. Mauri, G. Galli, and R. Car, Phys. Rev. B **47**, 9973 (1993).
[14] W. Gao and W. E, Discrete Contin. Dyn. Syst. **23**, 249 (2009).
[15] J. Kim, F. Mauri, and G. Galli, Phys. Rev. B **52**, 1640 (1995).
[16] W. Yang, Phys. Rev. B **56**, 9294 (1997).
[17] S. K. Burger and W. Yang, J. Phys.: Condens. Matter **20**, 294209 (2008).
[18] C. García-Cervera, J. Lu, and W. E, Commun. Math. Sci. **5**, 999 (2007).
[19] W. E, T. Li, and J. Lu (unpublished).
[20] Y. Zhou, Y. Saad, M. Tiago, and J. Chelikowsky, J. Comput. Phys. **219**, 172 (2006).
[21] D. C. Langreth and J. P. Perdew, Phys. Rev. B **21**, 5469 (1980).
[22] D. C. Langreth and M. J. Mehl, Phys. Rev. B **28**, 1809 (1983).
[23] N. Troullier and J. L. Martins, Phys. Rev. B **43**, 1993 (1991).
[24] L. Kleinman and D. M. Bylander, Phys. Rev. Lett. **48**, 1425 (1982).
[25] W. Kohn, Phys. Rev. **115**, 809 (1959).
[26] J. Lu, Ph.D. thesis, Princeton University, 2009.
[27] R. Martin, *Electronic Structure: Basic Theory and Practical Methods* (Cambridge University Press, Cambridge, 2005).
[28] J. F. Annet, Comput. Mater. Sci. **4**, 23 (1995).
[29] B. Parlett, *The Symmetric Eigenvalue Problem*, Classics in Applied Mathematics Vol. 20 (Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998).
[30] N. Marzari and D. Vanderbilt, Phys. Rev. B **56**, 12847 (1997).
[31] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins Studies in the Mathematical Sciences (Johns Hopkins University Press, Baltimore, MD, 1996).
[32] B. Jansík, S. Høst, P. Jørgensen, J. Olsen, and T. Helgaker, J. Chem. Phys. **126**, 124104 (2007).
[33] E. Rubensson and P. Salek, J. Comput. Chem. **26**, 1628 (2005).
[34] M. Reed and B. Simon, *Methods of Modern Mathematical Physics IV. Analysis of Operators* (Academic, New York, 1978).
[35] Strictly speaking, the system still has a tiny gap, but the exponent of exponential decay is so small that the algebraic behavior is dominant.
[36] C. García-Cervera, Comm. Comp. Phys. **2**, 334 (2007).
[37] J. R. Chelikowsky, N. Troullier, and Y. Saad, Phys. Rev. Lett. **72**, 1240 (1994).
[38] J. Chelikowsky, J. Phys. D: Appl. Phys. **33**, R33 (2000).
[39] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, J. A. Montgomery, Jr., T. Vreven, K. N. Kudin, J. C. Burant *et al.* (unpublished).
[40] http://www.fisica.uniud.it/giannozz/
[41] D. M. Ceperley and B. J. Alder, Phys. Rev. Lett. **45**, 566 (1980).
[42] J. P. Perdew and A. Zunger, Phys. Rev. B **23**, 5048 (1981).
[43] E. Isaacson and H. B. Keller, *Analysis of Numerical Methods* (Wiley, New York, 1966).