

Large scale electronic structure calculations using the Lanczos method

Lin-Wang Wang, Alex Zunger

National Renewable Energy Laboratory, Golden, CO 80401, USA

Received 17 February 1994; accepted 24 February 1994

Abstract

The orthogonality requirement in either iterative diagonalization or conjugate gradient approaches to the single particle Schrödinger equation $\hat{H}\psi = E\psi$ leads to an overall N^3 scaling of the effort with the number N of atoms. We show that the Lanczos method circumvents this problem even when applied to all occupied states. Our implementation shows that the method is stable, exact, scales as N^2 for N around a few hundreds, and is thus optimally suited for such mid-size (100 – 1000 atoms) quantum systems. The analogy between the basic Lanczos equations and Anderson's localization in a disordered one-dimensional tight-binding chain is pointed out and used to gain some insights into improved convergence and stability of the method. For a 900-atom Si cluster tested here using pseudopotentials and a plane wave basis, the Lanczos method is about an order of magnitude faster than the state-of-the-art preconditioned conjugate gradient method using the same pseudopotentials and basis set.

1. Introduction

While recent developments in computational strategies [1] enable first principles electronic structure calculations for systems with up to 100 atoms, rapid experimental advances are constantly shifting interest to quantum systems with an ever increasing number of atoms. Examples include the ≥ 1000 -atom quantum dot and quantum wire structures [2], as well as superlattices and quantum wells [3]. The effort in state-of-the-art $\hat{H}\psi = E\psi$ electronic structure algorithms scales as N^3 , where N is the number of atoms in the system. Recently [4], we have demonstrated that one could find *exact* eigenfunctions of $\hat{H}\psi = E\psi$ in a desired "energy window" (e.g., E around the band gap of insulators and semiconductors) in a linear-in- N scaling. This is very useful for energy level

calculations of large quantum systems [5], but not for *total* energy calculations that require *all* occupied eigensolutions. Although there are several promising proposals for total-energy electronic structure method with a linear-in-size (N) scaling of the effort [6–10], these are still in their formative stages and the cross-over size of their cost with respect to the conventional (N^3 -scaling) methods is yet unknown. Here we present a method for finding all *exact* occupied eigenstates of a given Schrödinger equation based on the Lanczos algorithm [11]. The effort scales as N^2 for N around a few hundreds of atoms (basis set size around 30 000) and is an order of magnitude faster than the state-of-the-art preconditioned conjugate gradient method [1]. Although for large N (perhaps > 1000), the Lanczos method could be less effective than the currently proposed [6–10] *linear* scaling meth-

ods, our proposed method appears competitive for mid-range system sizes (100–1000 atoms). Using this method, a plane-wave calculation for a $\text{Si}_{617}\text{H}_{316}$ cluster (having 1392 occupied states and 30 000 basis functions) takes 2.5 Cray-YMP cpu hours for a given total potential $V(r)$, while our implementation of the preconditioned conjugate gradient is estimated to take 23 cpu hours for the same problem. Thus, the present method makes electronic structure calculations of all occupied states of a few hundred to one thousand atom quantum systems affordable.

We use the Kohn–Sham formalism [12] in which the total energy of a system with external potential $V_{\text{ext}}(r)$ (e.g. the ionic potential) is:

$$E_{\text{tot}} = 2 \sum_l^{\text{occupied}} \left[-\frac{\hbar^2}{2m} \right] \int \psi_l \nabla^2 \psi_l d^3r + \int V_{\text{ext}}(r) \rho(r) d^3r + \frac{e^2}{2} \int \frac{\rho(r)\rho(r')}{|r-r'|} d^3r d^3r' + E_{\text{xc}}[\rho(r)] + E_{\text{ion}}[\{\mathbf{R}_{\text{atom}}\}], \quad (1)$$

where E_{ion} is the ion–ion Coulomb repulsion energy of the atomic configuration $\{\mathbf{R}_{\text{atom}}\}$, $E_{\text{xc}}[\rho(r)]$ is the total exchange and correlation energy of the electronic charge density $\rho(r)$ given by the occupied single-particle wavefunctions $\psi_l(r)$:

$$\rho(r) = 2 \sum_l^{\text{occupied}} |\psi_l(r)|^2. \quad (2)$$

The wavefunctions can be obtained from the single particle self-consistent Schrödinger equation

$$\left[-\frac{\hbar^2}{2m} \nabla^2 + V(r) \right] \psi_l(r) = E_l \psi_l(r), \quad (3)$$

where $V(r)$ is the self-consistently determined total potential:

$$V(r) = V_{\text{ext}}(r) + \frac{e^2}{2} \int \frac{\rho(r')}{|r-r'|} d^3r' + \frac{\delta E_{\text{xc}}[\rho(r)]}{\delta \rho(r)}. \quad (4)$$

The preconditioned conjugate gradient (PCG) method is an effective method for solving Eqs. (3) and (4). Since however, the PCG method requires performing orthogonalization (a N^3 operation), it is limited to systems with fewer than about 100 atoms. Here, we concentrate on solving Eq. (3) for a *given* $V(r)$. Self-consistency can be subsequently obtained by repeating this step and generating $V(r)$ from Eq. (4) using the solutions $\{\psi_l(r)\}$ of Eq. (3).

We expand ψ in a plane wave basis:

$$\psi_l(r) = \sum_G a_l(G) e^{-iG \cdot r}, \quad (5)$$

where G is the reciprocal lattice vector of the system and $a_l(G)$ are the expansion coefficients to be determined. Then, for a *given* $V(r)$, Eq. (3) is a linear equation of the variables $a_l(G)$, where its dimension is the number of reciprocal lattice vector $\{G\}$; for the largest systems studied here, this is $\sim 30\,000$. In what follows we will assume that one can compute efficiently the matrix-by-vector product $\hat{H}\psi_l$ [e.g. using fast Fourier transform (FFT)], so *individual* matrix elements $H_{G_1, G_2} = \langle \exp(-iG_1 \cdot r) | \hat{H} | \exp(-iG_2 \cdot r) \rangle$ need not be computed. [Using the FFT, a single multiplication $\hat{H}\psi_l$ scales as $N \ln(N)$, so it is much faster than explicit matrix multiplication using the matrix elements H_{G_1, G_2} .]

It has been previously suggested [13,14] that the Lanczos algorithm could be an efficient way to solve for the eigenstates of large linear systems. However, although the Lanczos method has been recently used to calculate some (~ 30) of the lowest eigenstates of large systems [15], total energy Lanczos calculations (requiring *all* occupied eigenstates) have apparently not been widely used. Indeed, the alleged instability of the basic Lanczos scheme [11] and the success of iterative diagonalization [16] and the conjugate gradient method [17] in electronic structure calculations seem to have discouraged testing of the Lanczos method for large scale total energy calculations. To examine these paradigms, we will first introduce the Lanczos method and discuss the “Lanczos phenomena”. This analysis (de-

ogy between Lanczos convergence and the Anderson localization problem then suggest a particular implementation of the Lanczos approach, discussed in section 3. Section 4 describes applications to Si quantum dots and section 5 compares the method to the preconditioned conjugate gradient method, while section 6 provides a summary.

It is useful to illustrate some of the Lanczos phenomena on a concrete system. In what follows, we will use large Si quantum dots as examples. Here, the total potential $V(\mathbf{r})$ will be taken as a superposition $V(\mathbf{r}) = \sum_{\text{atom}} v_{\text{atom}}(|\mathbf{r} - \mathbf{R}_{\text{atom}}|)$ of spherically symmetric empirical pseudopotentials $v_{\text{atom}}(r)$ [5]. Our basic algorithm does not depend on this choice. The Si atoms have a bulk-like configuration and are arranged in a shape of a rectangular box; all surface dangling bonds are saturated by hydrogen atoms. We consider four quantum boxes: Si₄₇H₅₂, Si₁₄₇H₁₁₆, Si₃₂₉H₂₀₄ and Si₆₁₇H₃₁₆. We use a plane wave basis (Eq. (5)) with 5520, 11 012, 18 794 and 29 982 orbitals for the four clusters, respectively. We apply periodic boundary conditions on each unit cell that contains the quantum structure surrounded by vacuum. The action of \hat{H} on ψ is executed using FFT's with a real-space grid of up to $54 \times 54 \times 80$. Figures 1–5 illustrate various “Lanczos phenomena”, to be discussed below, on the Si₄₇H₅₂ quantum dot.

2. Lanczos method and Lanczos phenomena

The Lanczos method [11,13,14] starts with a random wavefunction u_1 , then iteratively generates “Lanczos vectors” $\{u_i\}$ from

$$\beta_{i+1}u_{i+1} = \hat{H}u_i - \alpha_i u_i - \beta_i u_{i-1}, \quad (6)$$

where $\hat{H} = -\hbar^2/2m\nabla^2 + V(r)$ is the Hamiltonian of the system (Eq. (3)), α_i is given by $\langle u_i|\hat{H}|u_i \rangle$ and β_{i+1} is determined by the normalization condition $\langle u_{i+1}|u_{i+1} \rangle = 1$. At the start of the iterations, we have $\beta_1 = 0$. Using induction, it is easy to show from Eq. (6) that, for the exact (e) Lanczos algorithm, u_{i+1}^e is orthogonal to all previous Lanczos vectors, i.e.

$$\langle u_i^e|u_j^e \rangle = \delta_{i,j}. \quad (7)$$

Thus, the exact u_i^e represents an orthonormal basis which tridiagonalizes (Eq. (6)) the original Hamiltonian \hat{H} . The element of this matrix H_{ij}^T are defined by $\langle u_i^e|\hat{H}|u_j^e \rangle$. However, if we only have u_i with a finite precision, H^T can be defined by Eq. (6) and written as:

$$H^T = \begin{pmatrix} \alpha_1 & \beta_2 & & & & \\ \beta_2 & \alpha_2 & \beta_3 & & & \\ & \beta_3 & \alpha_3 & \cdot & & \\ & & & \cdot & \cdot & \cdot \\ & & & & \cdot & \cdot & \beta_M \\ & & & & & & \beta_M & \alpha_M \end{pmatrix}. \quad (8)$$

All eigenvalues $\{E_i\}$ of this tridiagonal matrix can be obtained by standard techniques [14,18] with an effort proportional to $O(M)$, where M is the dimension of H^T . Then, the individual eigenfunctions $\{b_i^l\}$ of H^T can be calculated using “inverse iteration” [14] of H^T around the energy E_i . This is a very fast, $O(M)$ operation for each eigenfunction $\{b_i^l\}$. Finally, the eigenfunctions of \hat{H} can be constructed from the eigenfunction of H^T as

$$\psi_l(r) = \sum_i^M b_i^l u_i(r). \quad (9)$$

The above arguments are valid only for the exact Lanczos algorithm, i.e. when there are no round off errors in the numerical implementation. In practice, round off errors always exist, so some of the above equations, especially the orthogonality condition of Eq. (7), are not satisfied. In the following, we will discuss some properties of the finite precision Lanczos algorithm. These properties are referred to as “Lanczos phenomena”. We will identify here a number of problems with the Lanczos method. Cures will then be offered in our algorithm described in section 3.

2.1. Amplification of end-of-spectrum states

As the number of Lanczos iterations M increases, the eigenstates at the two ends of the

spectrum (“end-of-spectrum states”) are the first to converge. This is because there is a hidden driving mechanism which amplifies the amplitudes of the end-of-spectrum states. To see this effect, let’s first write u_i and u_{i+1} as sums of the eigenstates of \hat{H} , i.e. $u_i = \sum_l c_l^i \psi_l$ and $u_{i+1} = \sum_l c_l^{i+1} \psi_l$. Substituting these into the Lanczos iterative equation (6), we have:

$$c_l^{i+1} = \frac{1}{\beta_{i+1}} [(E_l - \bar{E}_i) - \beta_i] c_l^i, \quad (10)$$

where, $\bar{E} = \sum_l E_l c_l^i$ is the average eigenvalue of u_i . So if at the beginning of the iterative sequence, all c_l^i have the same order of magnitude, then $(E_l - \bar{E}_i)$ will have the largest value for the end-of-spectrum states l . Thus, the amplitude c_l^i of these states will increase exponentially as a function of iteration history i , until they are large enough to begin to change \bar{E}_i . After this, their amplitudes will decrease because the factor $(E_l - \bar{E}_i)$ becomes small for these l . In this discussion, we have ignored the effect of the β_i term in Eq. (10). But this term will not change the qualitative picture. This is confirmed in Fig. 1, which depicts the amplitude of the $\psi_l(r)$ component in $u_i(r)$ vs the iteration index i . We see that first $\langle \psi_l | u_i \rangle$ increases exponentially as a function of i and subsequently it decays. Because these end-of-spectrum ψ_l ’s appear in u_i with large amplitudes, we can expect that ψ_l can be accurately represented by some linear combinations of u_i . In other words, these ψ_l ’s have converged.

2.2. Loss of orthogonality among $\{u_i\}$

Because of the mechanism discussed in section 2.1 and because of numerical round off errors, the orthogonality condition (7) does not hold in actual finite-precision computations. The exact Lanczos algorithm constrains the process, so the previous ($j < i$) u_j will not have any components in the newly generated u_i . However, in finite precision calculations, there is always a small, but finite admixture of u_j into u_i . This admixture will be amplified exponentially by the driving mechanism described in section 2.1, and

will eventually lead to the violation of the orthogonality (Eq. (7)). This violation of the orthogonality is illustrated in Fig. 2, which depicts $\langle u_1 | u_i \rangle$ versus i . After $i > 180$, this overlap becomes large. This loss of orthogonality seem to have led to the belief [11] that the Lanczos algorithm breaks down unless one explicitly reorthogonalizes u_i to all the previous u_j ’s, (a rather expensive operation). However, as will be shown in section 2.3 below, the Lanczos algorithm without explicit reorthogonalization does work despite of the failure of Eq. (7).

2.3. Loss of orthogonality does not prevent finding accurate eigenstates

Let us just ignore for a moment the failure of Eq. (7) and go ahead and diagonalize the tridiagonal matrix H^T of Eq. (8), finding its eigenstates $\{b_l^i\}$. Then, $\{b_l^i\}$ satisfies:

$$\beta_i b_{i-1}^l + \alpha_i b_i^l + \beta_{i+1} b_{i+1}^l = E_l b_i^l \quad \text{for } 1 < i < M$$

and

$$\beta_M b_{M-1}^l + \alpha_M b_M^l = E_l b_M^l. \quad (11)$$

Now, using these coefficient $\{b_l^i\}$ to construct from Eq. (9) the wavefunction $\psi_l(r)$, and applying \hat{H} to this $\psi_l(r)$ using Eqs. (6) and (11), we find

$$\hat{H}\psi_l = E_l \psi_l + \beta_{M+1} b_M u_{M+1}. \quad (12)$$

We see that if b_M is very small, ψ_l is a numerically accurate eigenstate of \hat{H} . Equation (12) does not depend on the orthogonality condition of Eq. (7). If we had an orthogonal basis $\{u_i\}$, then a normalized $\{b_l^i\}$ would have guaranteed a normalized ψ_l . Thus, the lose of orthogonality means that ψ_l is no longer normalized. However, as long as the norm of ψ_l is much larger than b_M , this is not a real problem (since then the second term in Eq. (12) can be ignored). Thus, as will be demonstrated in our later numerical tests, accurate eigenstates of the original Hamiltonian \hat{H} can be obtained from Eq. (9). The reason for the smallness of b_M is discussed next.

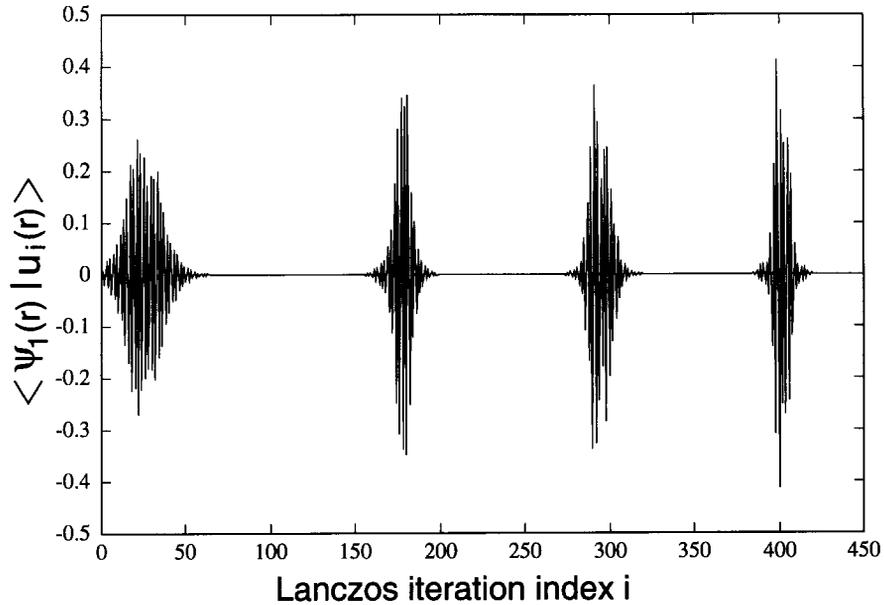


Fig. 1. The amplitude of the end-of-spectrum states $\psi_1(r)$ in the Lanczos wavefunction $u_i(r)$. The amplitude increases from $i = 1$ to $i = 25$, then it decreases from $i = 25$ to $i = 50$. At $i = 100$, $\psi_1(r)$ is converged. The next and other peaks are duplicates of $\psi_1(r)$ discussed in section 2.5. These are also indications of the failure of orthogonality among the different $u_i(r)$. The system studied is $\text{Si}_{47}\text{H}_{52}$.

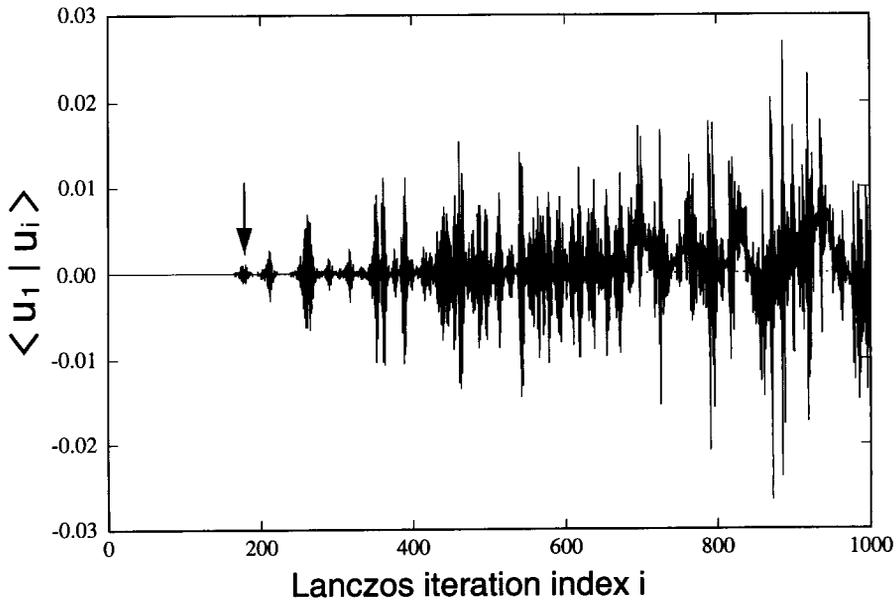


Fig. 2. The overlap coefficient $\langle u_1(r) | u_i(r) \rangle$. Note how $u_i(r)$ loses its orthogonality (i.e., nonzero overlap) with $u_1(r)$ after about $i = 200$ iterations. Note also that the first peak (indicated by an arrow at $i = 180$) is caused by the first duplicate of $\psi_1(r)$, shown in Fig. 1 ($i = 180$). The system studied is $\text{Si}_{47}\text{H}_{52}$.

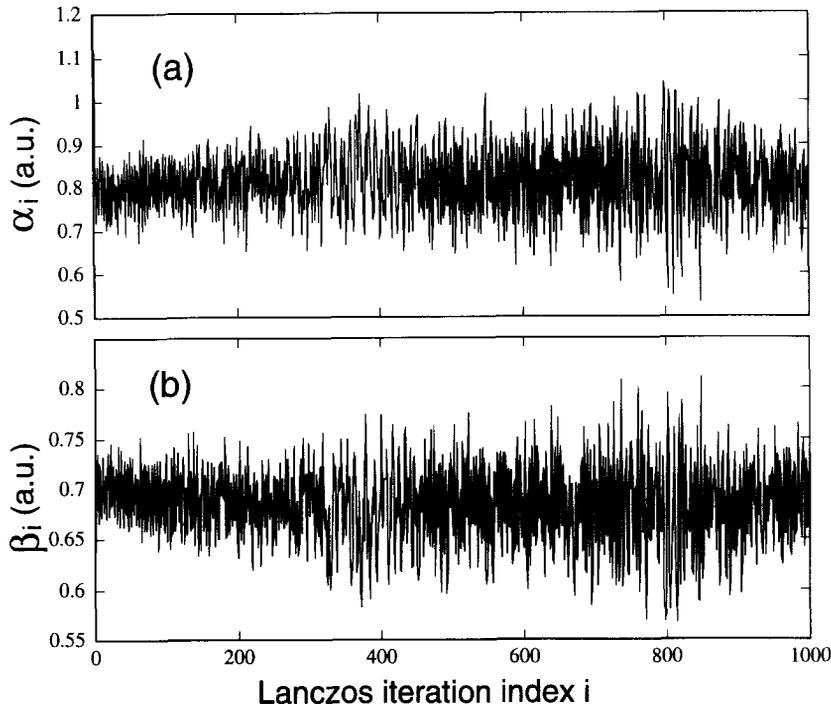


Fig. 3. The Lanczos coefficients α_i and β_i of Eq. (6) as functions of Lanczos iteration number i . Note the large random fluctuations in α_i and β_i . The system studied is $\text{Si}_{47}\text{H}_{52}$.

2.4. Anderson localization in b_i is the reason for the convergence of ψ_l

To understand why b_M can be so small, consider Eq. (11). Notice that this equation is isomorphous with the linear chain tight binding equation [19,20] for the eigenfunction $\{b_i^l\}$. In the linear chain problem one considers a single orbital per atom, where i is the atomic position and α_i and β_i are, respectively, the diagonal (“on-site”) and the off-diagonal nearest neighbor “hopping” Hamiltonian matrix element. Thus, the Lanczos iteration index $i=1,\dots,M$ is analogous to atomic positions in a chain and the Lanczos coefficients α and β are analogous to tight-binding matrix elements. Fig. 3 shows the value of α_i and β_i calculated from the Lanczos procedure of Eq. (6) for our test system $\text{Si}_{47}\text{H}_{52}$. We see that there are significant *random fluctuations* in the Lanczos value of α_i and β_i as a function of the iteration index i . Thus, Eq. (11) represents a *random* one-dimensional

tight-binding system. For such a system, it is well known that all eigenstates are localized with exponentially decaying tails (“Anderson localizations”) [19]. The decay length is determined by the magnitude of the randomness: the larger the randomness, the smaller the decay length, (and thus the faster the Lanczos convergence rate of the corresponding eigenstates). If site i on which localization occurs is far from site M , then b_M^l will be small and ψ_l will be well converged. The coefficients $\{b_i^l\}$ of typical converged eigenstates ψ_l ($l=1$ and $l=10$) are shown in Fig. 4. The localizations and the tails of $\{b_i^l\}$ are clearly seen. Because of the mechanism discussed in section 2.1, the localization site of the end-of-spectrum states always occurs at small position index i . Thus, these states will be converged even for small M . In practice, soon after M passes the localization site i of $\{b_i^l\}$, the wavefunction ψ_l and its eigenvalue E_l will converge. On the other hand, for a given M , there are eigenstates of Eq. (11) which have a

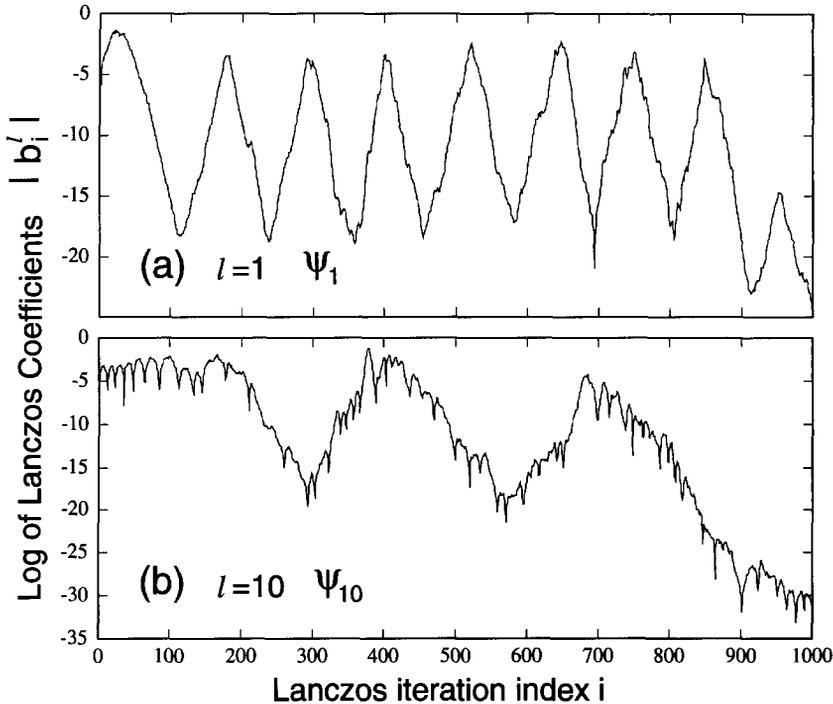


Fig. 4. The expansion coefficients $\{b_i^l\}$ of the basis $\{u_i(r)\}$ [Eq. (9)], for (a) $\psi_1(r)$, and (b) $\psi_{10}(r)$. The dimension of the tridiagonal matrix H^T used to calculate $\{b_i^l\}$ is $M = 1000$. The localization of $\{b_i^l\}$ is apparent from the exponential decay (straight line for $\ln(b_i^l)$) around the localization sites (the peak positions of $\ln(b_i^l)$). Note that $\{b_i^1\}$ has a much shorter decay length than $\{b_i^{10}\}$. The multiple peaks are a result of the coupling between the duplicates of $\psi_l(r)$. Because the last duplicate (centered around $i = 1000$) for $\{b_i^{10}\}$ has not been fully converged, there is no coupling to this state. Thus, this emerging (yet unconverged) duplicate does not spoil the convergence of the previously converged duplicates. The system studied is $\text{Si}_{47}\text{H}_{52}$.

localization site i close to M , so these eigenstates have large b_M^l value. These eigenstates are not well converged and their eigenvalues are scattered throughout the spectrum. However, because of their large b_M^l value, their eigenvalues are very sensitive to a change in M . Thus, their eigenvalues will change as we replace M by, e.g. $M - 1$. This is useful as a criterion to identify these spurious states. [The absolute amplitude of b_M^l is not used to judge the convergence, because the normalization of ψ_l constructed from Eq. (9) can be very different from unity.]

2.5. The occurrence of eigenvalue duplicates

Another consequence of the “driving mechanism” (section 2.1) and the loss of orthogonality

(section 2.2), is that the Lanczos procedure will produce many duplicates of a single eigenstate. In the language of the one-dimensional random system, one can say that there are many localized states with degenerate eigenvalues. The occurrence of duplicates is illustrated in Fig. 1 and Fig. 4. Note that b_i^l in Fig. 4 decays to machine precision $\epsilon (= 10^{-13})$ before it increases again to generate the next duplicate. The duplicates are generated at almost equal intervals d_l determined by:

$$d_l \cong -2\lambda_l \ln \epsilon, \quad (13)$$

where λ_l is the decay length of state $\{b_i^l\}$. From Anderson’s model [19], we know that λ_l depends on the magnitude of the randomness of α_i, β_i and the eigenvalue E_l . The smaller E_l , the

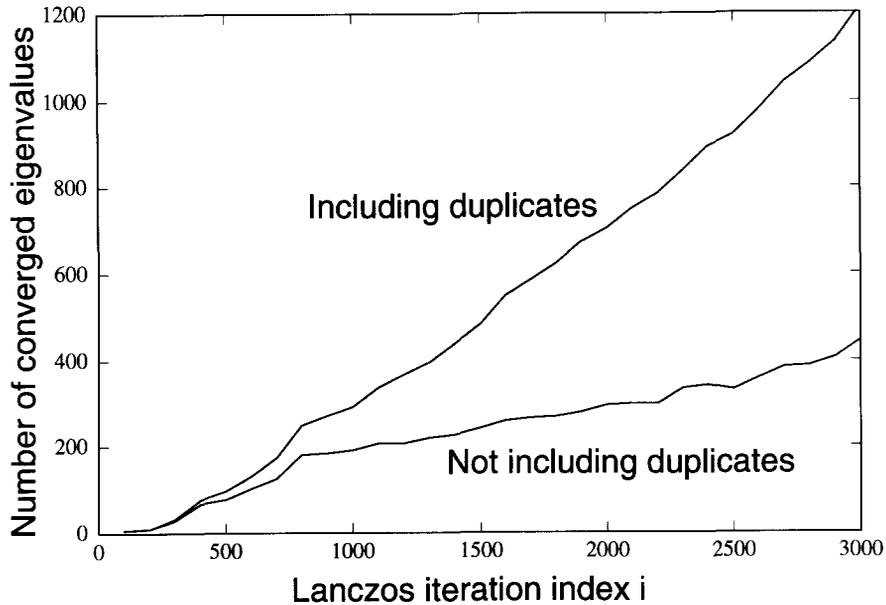


Fig. 5. The Number of converged eigenvalues as functions of the Lanczos iteration index i . Since this system has 120 occupied states, most of the occupied states should have already converged around $i \sim 1000$. The system studied is $\text{Si}_{47}\text{H}_{52}$.

shorter is λ_i , thus the more often it repeats itself. This is confirmed in Fig. 4 by the much shorter decay length of $\{b_i^1\}$ relative to that of $\{b_i^{10}\}$.

In actual calculation, the duplicate states can be easily recognized by noting their nearly identical eigenvalues. Thus, only a single $\{b_i^l\}$ and ψ_l will be calculated from these identical eigenvalues. However, if the original Hamiltonian \hat{H} supports truly degenerated states, this procedure will clearly miss one partner of a degenerate set. This problem will be solved in our algorithm (see below) by restarting the Lanczos iteration several times with different starting random wavefunctions u_1 .

A related concern is whether a yet unconverged duplicate of E_l which is just emerging, (i.e. the localization site of this new duplicate is very close to M), could spoil the convergence of the previously converged states of E_l . Considering the picture of a one-dimensional random system and actual numerical tests, we find that the already converged eigenvalue E_l will *not* be shifted by the newly emerging duplicate. Furthermore, the amplitude $\{b_i^l\}$ of the newly emerging state can be coupled to the previously

converged states only after the new state has been fully converged. This is illustrated in Fig. 4(b).

Another concern is that with all converged states generating their duplicates at a rate given by Eq. (13), could it be that after a while, the procedure will produce only duplicates of old states rather than new states. Fortunately, this does not happen, as illustrated in Fig. 5. We see that, while the duplicates can account for more than half of the total number of converged eigenstates, the number of unique converged states (without counting the duplicates) increases steadily almost as a linear function of the number M of Lanczos iterations. In the example shown in Fig. 5, there are two slopes: a larger slope for the states below the energy gap in the spectrum (i.e. the band gap of $\text{Si}_{47}\text{H}_{52}$), and a smaller slope for the higher energy states. Of course, after all the eigenstates of the original Hamiltonian have been generated, further iterations can only produce duplicates.

Finally, the ultimate test of accuracy and convergence of an eigenstate ψ_l is to apply on it the original Hamiltonian \hat{H} and examine the full-

fillment of the Schrödinger equation (3). This check will be used in our following algorithm.

3. The algorithm for electronic structure calculations

Based on our analysis of the Lanczos phenomena described in the previous section, we next present our Lanczos scheme for calculating all the occupied eigenstates $\{\psi_l(r)\}$ of Schrödinger equation (3). This algorithm will cure most problems identified in section 2. The algorithm is summarized in Fig. 6, and explained step by step in what follows.

Step (i): start with a randomly selected wavefunction $u_1(r)$ (constructed from choosing random coefficients of a given basis), and generate the Lanczos wavefunction $u_i(r)$ according to Eq. (6), without any reorthogonalization among the members of $\{u_i(r)\}$. If this is the first sweep of steps (i)–(vi) ($N_{\text{sweep}} = 1$), skip the following. If this is not the first sweep – so there are N_{tot}^c converged eigenstate $\{\psi_l(r)\}$ from previous sweeps – then orthogonalize each $u_i(r)$ to all $\{\psi_l(r)\}$ by subtracting the component of each $\psi_l(r)$ from $u_i(r)$, i.e. $u_i - \sum_l \langle \psi_l | u_i \rangle \psi_l$.

If disk space allows, store all generated $\{u_i\}$ on disk for later use. If not, discard the old u_i .

Step (ii): at some Lanczos iteration $i = N_t$ (e.g. every 50 or more Lanczos steps), diagonalize the tridiagonal Lanczos matrix H^T of Eq. (8) using standard programs (e.g. the LAPACK routines) and obtain the eigenvalues $\{E_l(N_t)\}$ without calculating the eigenvectors. Next, reduce the dimension of H^T by one and repeat the diagonalization of H^T , finding the set $\{E_l(N_t - 1)\}$. Compare $\{E_l(N_t)\}$ with $\{E_l(N_t - 1)\}$, keeping the eigenvalues that have changed by less than ϵ_t , where ϵ_t is a desired convergence criterion (e.g. 1×10^{-6} Hartree). For eigenvalues that are degenerated within ϵ_t , keep only a single copy. Denote these eigenvalues as $\{E_l^c(N_t)\}$. Then calculate the number $N_c(N_t)$ of $E_l^c(N_t)$'s which are below the Fermi energy, (the latter can be obtained for insulator from calculations on a small system). If $N_c(N_t)$ is larger than a desired percentage (e.g. 98%) of the total number N_{occ} of

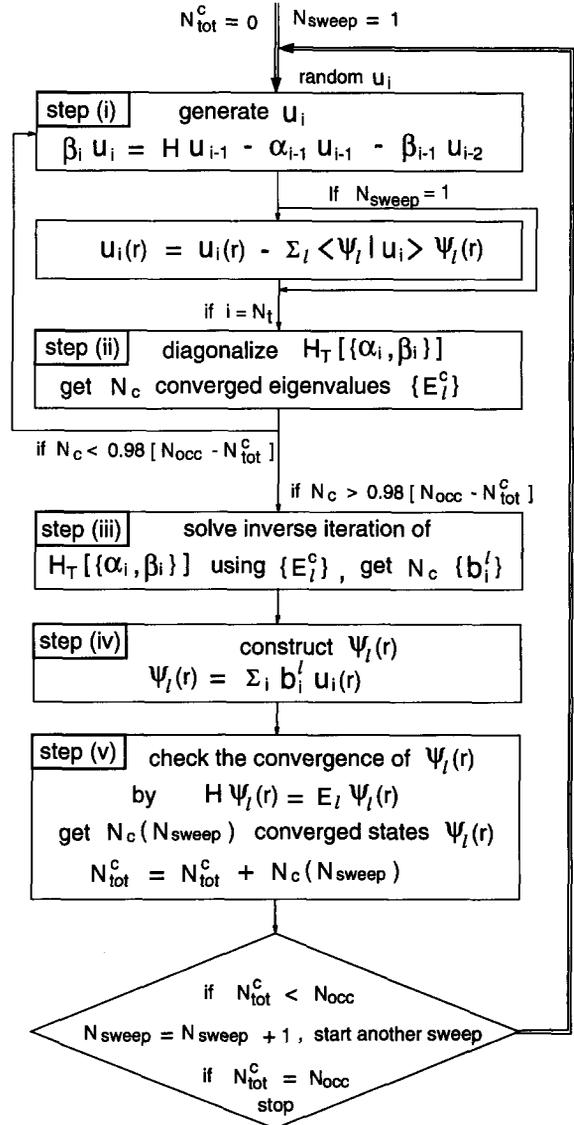


Fig. 6. Flow chart describing the Lanczos procedure for large scale electronic structure calculation. See section 3 for details.

occupied states, then go to step (iii). If not, continue step (i).

Step (iii): using $\{E_l^c(N_t)\}$ as the input energies, and applying the inverse iteration method to the tridiagonal matrix H^T (using standard programs, e.g. the LAPACK routine), calculate $N_c(N_t)$ eigenfunctions of H^T , $\{b_l^i\}$ in the basis of the Lanczos vector $\{u_i(r)\}$. As can be seen

from Fig. 4, most of the eigenstates have been converged far before N_l is reached, so for each E_l^c , it is worth doing a few (5 to 10) inverse iterations for different matrix dimensions M ($< N_l$), and examining whether the eigenvalue that resulted from the inverse iteration has changed from $E_l(N_l)$ (within a tolerance limit ϵ_l , which can be much smaller than ϵ_t and close to the machine precision). Choose the smallest M for which E_l does not change beyond ϵ_l , and denote this as M_l . Then, $\{b_i^l\}$ (which has M_l terms i) are the coefficients based on the M_l -dimensioned matrix H^T .

Step (iv): generate ψ_l from $\{b_i^l\}$ and $\{u_i\}$. If $\{u_i\}$ is stored on disk, retrieve $\{u_i\}$ from disk. If it is not, recalculate $\{u_i\}$ exactly as in step (i). While u_i are retrieved or recalculated one by one, construct the $N_c(N_l)$ $\psi_l(r)$'s from Eq. (9) and from the $\{b_i^l\}$ and M_l obtained in step (iii). Because of the nonorthogonality of $\{u_i\}$, the $\psi_l(r)$ from Eq. (9) is not normalized. So, at the end this step, normalize all $\psi_l(r)$'s.

Step (v): check the accuracy of the $N_c(N_l)$ ψ_l 's generated in step (iv) by evaluating the error $\delta_l \equiv \langle \psi_l | (H - E_l)^2 | \psi_l \rangle^{1/2}$. If $\delta_l < \epsilon_t$, accept ψ_l , else, discard it. If there are two eigenstates whose E_l 's are within the order of their δ_l 's, then orthogonalize them using the Gram–Schmidt scheme [18]. The error of the resulting wavefunction then needs to be checked again. Upon completion of step (v), we will have $N_c(N_{\text{sweep}})$ newly converged eigenstates. This brings the total number of converged eigenstates to $N_{\text{tot}}^c(N_{\text{sweep}}) = N_{\text{tot}}^c(N_{\text{sweep}} - 1) + N_c(N_{\text{sweep}})$.

Step (vi): if $N_{\text{tot}}^c(N_{\text{sweep}})$ is less than N_{occ} , go to step (i) and begin another sweep ($N_{\text{sweep}} = N_{\text{sweep}} + 1$). If $N_{\text{tot}}^c(N_{\text{sweep}})$ equals N_{occ} , stop.

We close this section by a number of comments on the algorithm:

(1) More recent sweeps ($N_{\text{sweep}} > 1$) operate similarly as the first sweep, except that the already converged states ψ_l have been removed from the spectrum of their effective Hamiltonian H^T . As a result, the Lanczos iteration of later sweeps have a faster convergence than that of the first sweep. Typically, one needs three or four sweeps to get all the occupied eigenstates. The number of Lanczos iterations and the number of

converged eigenstates generated at each sweep decay as a geometrical series.

(2) The Lanczos procedure described here is stable. It guarantees that each converged eigenstate ψ_l is accurate to within a given tolerance ϵ_l . Since we use *consecutive* sweeps (steps (i)–(vi)), the procedure also guarantees that all occupied eigenstates (irrespective to whether they are degenerate or not) are found. (An eigenstate missed from the first sweep can be found in subsequent sweeps.)

(3) All eigenstates are mutually orthogonal simply because they are found to be numerically accurate eigenstates of the *hermitian* Hamiltonian \hat{H} . Possible nonorthogonality among degenerate, or nearly degenerate states is avoid by explicit orthogonalizations in steps (i) and (v).

(4) The procedure works optimally for systems without explicit symmetry, thus without too many degenerated states. Otherwise, the first sweep cannot generate the majority of the eigenstates, and the procedure will be slow. One obvious solution for systems with explicit symmetry is to use group theory to break the original Hamiltonian \hat{H} into several blocks, and then apply the current procedure to diagonalize each block. Another solution is to apply, at the end of step (v), the symmetry operations on each newly converged eigenstate ψ_l , to see whether it can generate new (but energy degenerate) states. If it can, add them as converged eigenstates.

(5) For a typical nonsymmetric system, we find in our numerical tests that the first sweep can generate more than 90% of the total number of occupied eigenstates. Thus, the total number of Lanczos iterations in subsequent sweeps (which need explicit orthogonalization to the previously converged eigenstates ψ_l) is about 1/10 of the number of Lanczos iterations in the first sweep. Thus, practically speaking, the whole procedure can be considered as a Lanczos scheme without reorthogonalization.

4. Applications to Si quantum dots

We now apply the above Lanczos scheme to Si quantum dots described at the end of the In-

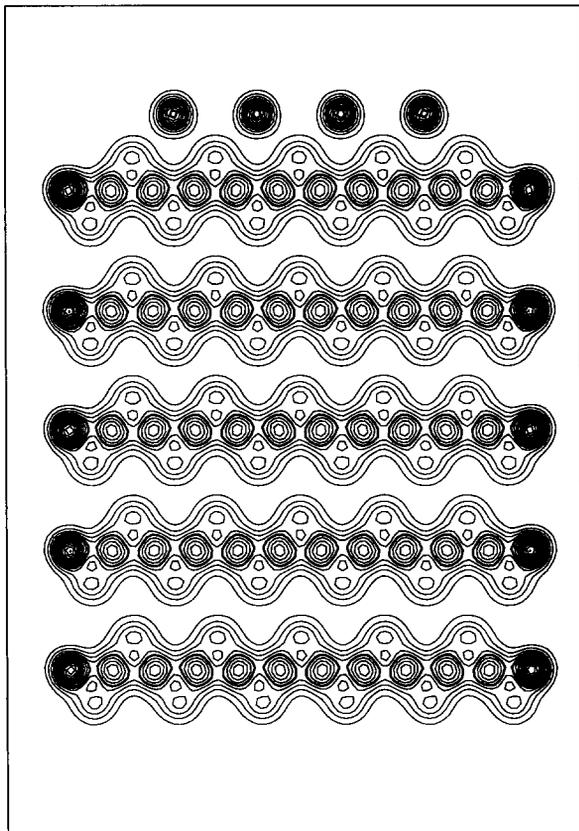


Fig. 7. Calculated total charge density contour plot for $\text{Si}_{617}\text{H}_{316}$. The contour plot is shown on a $[110]$ cross section at the center of the rectangular box. The Hydrogen atoms are on the periphery. The interior Si-Si bonds can be clearly seen.

roduction. The quantum dots studied here have approximate rectangular shapes, and all surface dangling bonds are passivated by hydrogen atoms. The symmetry of the pure Si skeleton is destroyed by the tilting of the surface H atoms. The pseudopotentials used for Si and H were given in ref. [5]. We use $\epsilon_t = 1 \times 10^{-6}$ Hartree as our convergence criterion. We do not store on disk the Lanczos wavefunctions u_i in step (i), but regenerated them in step (iv). The charge density contour plot of the $\text{Si}_{617}\text{H}_{316}$ system calculated by the current method is shown in Figs. 7 and 8 in two different planes. Hydrogen atoms are evident at the periphery of the cluster. In Fig. 7 [the (110) plane] we see clearly the Si-Si bonds. Figs. 7 and 8 also demonstrate

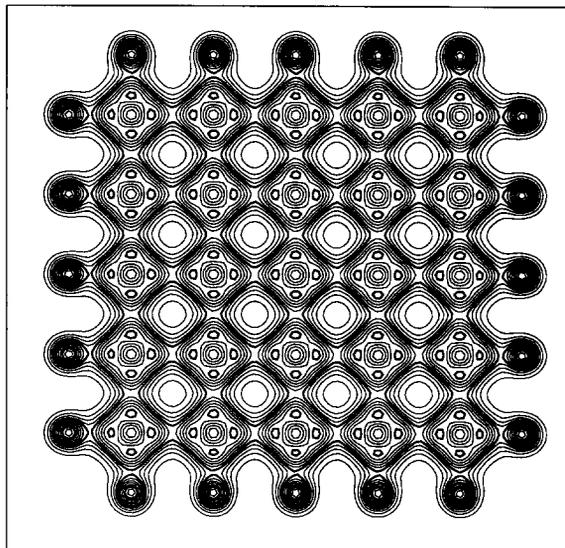


Fig. 8. Calculated total charge density contour plot for $\text{Si}_{617}\text{H}_{316}$. The contour plot is shown on a $[100]$ cross section at the middle of the rectangular box. The Hydrogen atoms are on the periphery.

that the charge density of the interior atoms is very close to that of bulk Si. At the surface, we find charge transfer that leads to the formation of strong Si-H bonds.

The time consumed by the different parts of the algorithm is given in Table 1 for the four systems considered here. For the $\text{Si}_{617}\text{H}_{316}$ system, this method takes about 2.5 Cray YMP cpu hours to calculate all 1392 occupied states. For this system, 41% of the time is spent on generating the Lanczos wavefunction u_i . This is a N^2 operation, where N is the size of the system. Thus, if $\{u_i(r)\}$ are stored on disk in step (i) instead of being regenerated in step (iv), 20% of the total time can be saved. About 28% of the time is spent on the constructing $\psi_l(r)$ using Eq. (9) and 20% of the time is spent on orthogonalizing $u_i(r)$ to previously converged $\psi_l(r)$ during the Lanczos iteration of $N_{\text{sweep}} > 1$. Both of these terms scale as N^3 . About 9% of the time is spent on using standard Lapack routines to diagonalize the tridiagonal matrix H^T and solve its inverse iteration. This corresponds roughly to a N^2 scaling. The remaining 2% of the time accounts for all other operations. The whole procedure scales between N^2 and N^3 , as shown in

Table 1

Time analysis (in Cray YMP cpu seconds) of the Lanczos method. “cpu for constructing ψ_l ” denotes the time consumed by constructing ψ_l using Eq. (9) in step (iv). “cpu for orth.” denotes the time consumed by making the u_i orthogonal to the converged ψ_l in the other-than-first sweeps. “cpu for diag. and inv. H^T ” denotes the time used to get the eigenvalues of H^T and to solve the inverse iteration to get the coefficients $\{b_i^l\}$.

	Si ₄₇ H ₅₂	Si ₁₄₇ H ₁₁₆	Si ₃₂₉ H ₂₀₄	Si ₆₁₇ H ₃₁₆
num. of occ. states	120	352	760	1392
FFT grid	32 × 32 × 48	40 × 40 × 60	48 × 48 × 72	54 × 54 × 80
memory required	6 Mb	33 Mb	120 Mb	360 Mb
total cpu time(sec)	71	471	2804	8986
cpu for $\hat{H}u_i$	61	330	1650	3686
cpu to construct ψ_l	3	50	456	2521
cpu for orth.	3	46	430	1792
cpu for diag. and inv. H^T	3	30	228	807
num. of sweeps	3	3	3	3
num. of Lanczos iter. in each sweep	1090	3359	9156	17657
	101	430	1285	1899
	101	151	166	256
num. of converged states in each sweep	113	319	661	1246
	6	28	86	125
	1	5	13	21
cpu of each sweep	59	373	2161	6730
	6	72	567	1975
	6	26	76	281

Fig. 9. For smaller size systems, the N^2 scaling of calculating $\hat{H}u_i$ dominates the computing time. At $N \simeq 1000$ atoms we find a cross-over of the N^2 to the N^3 scaling. Thus this method works best for systems with 100–1000 atoms.

One practical limitation is the run time memory. This is caused by the need to store in memory all the converged eigenstates $\{\psi_l(r)\}$. This aspect is the same as the conventional method (e.g. conjugate gradient method). The coefficient $\{b_i^l\}$ (30% of $\{\psi_l(r)\}$, in terms of memory space) can be stored on disk without slowing down the speed of the calculation. The possibility of storing $\{\psi_l(r)\}$ on the disk depends on (i) how rapidly can one retrieve them from disk, and (ii) how accurate would the orthogonalization of $u_i(r)$ to $\{\psi_l(r)\}$ be if this is only enforced every few Lanczos steps. This is left for future testing when memory space is not enough for truly large system calculations. In our example involving the 1392 occupied states of the Si₆₁₇H₃₁₆ cluster, the memory needed is

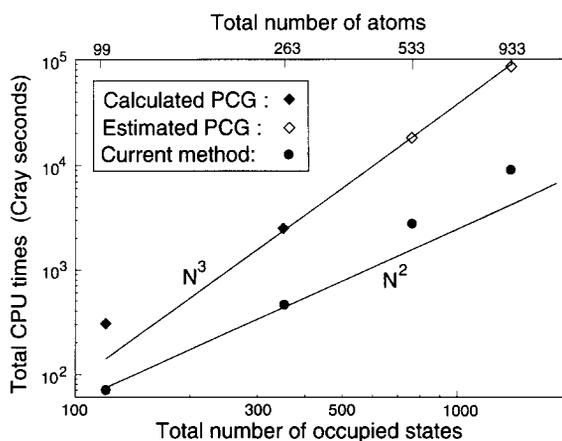


Fig. 9. The total computation times vs. system size for the preconditioned conjugate gradient (PCG) and the current Lanczos method. The computing times are in Cray YMP cpu seconds. The straight lines indicate the N^2 and N^3 scalings, where N is the size of the system. Note that for $N \sim 1000$ atoms, the current Lanczos method is an order of magnitude faster than the PCG.

360Mbyte.

To test whether this procedure is stable, we have repeated the calculation of the largest system $\text{Si}_{617}\text{H}_{316}$ using different starting random wavefunctions $u_1(r)$ in the Lanczos iteration (Eq. (6)). The two sets of the eigenvalues $\{E_l\}$ show a one-to-one correspondence, so no eigenstate is missed. The largest difference between the two sets of eigenvalues is $< 1 \times 10^{-10}$ Hartree. The mean square difference of the charge density $\rho(r)$ (on real space grid r) of these two runs is 0.0001%, indicating very accurate convergence. The sum of the eigenvalues (which is used for total energy calculations) differs by less than 1×10^{-10} Hartree. For most purposes, this accuracy is (more than) sufficient. Although the least converged eigenstates has a convergence error δ_l (defined in step (v)) close to ϵ_t , most eigenstates have much smaller convergence errors, i.e. close to machine precision. This and our numerical tests show that using a smaller ϵ_t does not affect much the computing time. The orthogonalization among different eigenstates ψ_l was checked explicitly by calculating $\langle \psi_l | \psi_{l'} \rangle$ for $l \neq l'$. The largest value found for this overlap element is 3×10^{-5} and occurs for the least converged eigenstates. For most states, however, this quantity is much smaller (i.e. close to the machine precision).

5. Comparisons with the preconditioned conjugate gradient method

We applied the preconditioned conjugate gradient (PCG) algorithm to the two smallest systems $\text{Si}_{47}\text{H}_{52}$ and $\text{Si}_{147}\text{H}_{116}$, using identical conditions and numerical techniques as in the Lanczos method for evaluation of the wavefunction $\psi_l(r)$ and for its multiplication by the Hamiltonian \hat{H} . Also, an identical convergence criterion $\epsilon_t = 1 \times 10^{-6}$, is used in both methods. The resulting computing times are shown in Table 2. We found that the eigenvalues calculated by the PCG method exhibit a one-to-one correspondence to their Lanczos counterparts, to within the convergence limit of 1×10^{-6} Hartree. The total charge density $\rho(r)$ calculated by PCG is

within 0.0001% of the Lanczos results. For systems larger than $\text{Si}_{147}\text{H}_{116}$, the computational times denoted by "E" in Table 2 have been estimated from the known scaling of the orthogonalization part (N^3 scaling) and the $H\psi(r)$ part (N^2 scaling). It can be seen from Table 2 that for systems larger than $\text{Si}_{47}\text{H}_{52}$ (the smallest considered), the computational time of PCG is dominated by the orthogonalization time. Thus, the cross-over size from N^2 to N^3 scaling for the PCG method is at $N \simeq 100$ atoms, compared to the 1000 atom cross-over size for the Lanczos method (Table 1 and Fig. 9). Also, Table 2 shows that for the smallest and the largest system, respectively, the Lanczos method is faster than the PCG method by a factor of 4 and 9. This is also illustrated in Fig. 9.

To understand why the Lanczos method is faster than the PCG method, recall that the conjugate gradient method is related to Lanczos method by the fact that both explore the same subspace. So, if only one or a few lowest eigenfunctions are needed, the conjugate gradient method should be more efficient because the preconditioning technique can be applied to the conjugate gradient method but not to the Lanczos method. However, if many eigenstates are needed, as in our examples, then the PCG method is less effective since it searches for each eigenstates *independently* (except for the orthogonalizations). This is analogous to starting a new Lanczos iteration series for each eigenstates. A set of Lanczos wavefunction $\{u_i(r)\}$ contain not only the informations of the lowest eigenstate $\psi_1(r)$, but also for other higher energy eigenstates. So, using $\{u_i(r)\}$ for just one eigenstate is not economical, compared with having a longer series of $\{u_i(r)\}$ and diagonalizing H^T of Eq. (8) to get all the eigenstates at once. In our examples, to get N occupied eigenstates, we need approximately $10N$ Lanczos iterations for the two smallest systems. For the same system, using the PCG method with preconditioning, getting each eigenstates requires about 30 conjugate gradient steps (line minimizations). Furthermore, note that each conjugate gradient step is more expensive than a Lanczos iteration step. More importantly, because of the independence

Table 2

Time analysis of the preconditioned conjugate gradient (PCG) method and the Lanczos method. “total num. of PCG iter.” is the total number of PCG line minimization steps. We use 1×10^{-6} Hartree as the convergence criterion. “cpu of $\hat{H}\psi_l$ ” is the time of applying \hat{H} to the wavefunction ψ_l , which is a N^2 operation. This step also includes the time of other N^2 scaling operations (e.g. operations like: $\psi_l(r) = \alpha\psi_l(r) + \beta p_l(r)$). “cpu of orth.” is the time of Gram–Schmidts explicit orthogonalization among $\{\psi_l\}$. This is a N^3 operation. Note that only the Si₄₇H₅₂ and Si₁₄₇H₁₁₆ quantum dots are actually calculated by the PCG method. The PCG computing times of other two systems indicated by “E” are estimated from their above noted scalings. All four quantum dots were actually calculated by the Lanczos method. All times are in Cray YMP cpu seconds.

	Si ₄₇ H ₅₂	Si ₁₄₇ H ₁₁₆	Si ₃₂₉ H ₂₀₄	Si ₆₁₇ H ₃₁₆
total num. of PCG iter.	31	32	32(E)	32(E)
total cpu time for PCG	305	2566	17800(E)	82100(E)
total cpu time for LANZOS	71	471	2804	8986
cpu of $\hat{H}\psi_l$ in PCG	206	1080	4900(E)	12900(E)
cpu of orth. in PCG	95	1470	12900(E)	69200(E)

of the searches of different eigenstates in the conjugate gradient method, the wavefunctions do not share any information, so their mutual orthogonality is not guaranteed, and has to be enforced explicitly. As Table 2 shows, this becomes the most time consuming part of the PCG algorithm. On the other hand, the Lanczos algorithm diagonalizes the tridiagonal matrix H^T of Eq. (8) and gets all eigenstates at the same time. Thus, the eigenstates “avoid each other” *automatically*. This leads to the orthogonalization among $\{\psi_l(r)\}$ without explicit enforcement. As a result, the algorithm is close to an N^2 scaling.

The above Lanczos algorithm is for all occupied state calculations. Not surprisingly, it is an order of magnitude slower than our novel method [4] used to calculate a *few* states around a given energy. If the band edge states (the top of valence band and the bottom of conduction band) are all that one wants, the above Lanczos method is not the best algorithm.

6. Conclusions

In order to utilize the Lanczos algorithm, several Lanczos phenomena have been explored in details. In particular, an unexpected relation was found between Lanczos convergence and one-dimensional Anderson localization. The understanding of the mechanism of exponential amplification of the end-of-spectrum states is dis-

cussed. Based on the insights into the Lanczos algorithm gained from this analysis, a Lanczos scheme is presented for large system total energy electronic structure calculations. Details of this algorithm are given. The programming of this method is no more complicated than that of the conjugate gradient method. This method is stable, all occupied eigenstates are guaranteed to be found and each eigenstate is guaranteed to be accurate to within a prescribed convergence limit. Thus, unlike other proposed schemes for large system calculations [6–10], the current approach is exact. For a few hundred atom system, this method scales as N^2 . When $N > 1000$, the scaling changes to N^3 . For the largest system tested here, this method is 9 times faster than the widely used preconditioned conjugate gradient method. Using this method, the charge density and 1392 eigenstates of a 933 atom quantum dot Si₆₁₇H₃₁₆ has been calculated within two and a half Cray YMP cpu hours. Thus, 100 – 1000 atom system total energy plane wave calculations become affordable by this method.

Acknowledgement

This work was supported by the office of Energy Research, Materials Science Division, U.S. Department of Energy, under Grant No. DE-AC02-83CH10093.

References

- [1] M.C. Payne, M.P. Teter, D.C. Allan, T.A. Arias and J.D. Joannopoulos, *Rev. Mod. Phys.* 64 (1992) 1045.
- [2] A.D. Yoffe, *Adv. Phys.* 42 (1993) 173.
- [3] G. Bastard, *Wave Mechanics Applied to Semiconductor Heterostructures*, (Les editions de physique, Les Ulis, 1988).
- [4] L.W. Wang and A. Zunger, *J. Chem. Phys.* 100 (1994) 2394.
- [5] L.W. Wang and A. Zunger, *J. Phys. Chem.*, in press, March 1994.
- [6] L.W. Wang, Ph.D thesis, (Cornell Univ.) 1991.
- [7] W. Yang, *Phys. Rev. Lett.* 66 (1991) 1438.
- [8] S. Baroni and P. Giannozzi, *Europhys. Lett.* 17 (1992) 547.
- [9] G. Galli and M. Parrinello, *Phys. Rev. Lett.* 69 (1992) 3547; F. Mauri, G. Galli and R. Car, *Phys. Rev. B* 47 (1993) 9973.
- [10] X.P. Li, R.W. Nunes and D. Vanderbilt, *Phys. Rev. B* 47 (1993) 10891.
- [11] C. Lanczos, *J. Res. Nat. Bur. Stand.* 45 (1950) 255.
- [12] W. Kohn and L. J. Sham, *Phys. Rev.* 140 (1965) A1133.
- [13] J.K. Cullum and R.A. Willoughby, *Lanczos Algorithms for Large Symmetric Eigenvalue Computations* (Birkhauser, Boston, 1985).
- [14] B.N. Parlett, *The Symmetric Eigenvalue Problem*, (Prentic-Hall, Englewood, Cliffs, N.Y., 1980).
- [15] F. Webster, P.J. Rossky and R.A. Friesner, *Comp. Phys. Comm.* 63 (1991) 494.
- [16] D.M. Wood and A. Zunger, *J. Phys. A* 18 (1985) 1343.
- [17] M.P. Teter, M.C. Payne and D.C. Allan, *Phys. Rev. B* 40 (1989) 12255.
- [18] G.H. Golub and C.F. Van Loan, *Matrix Computations*, (Johns Hopkins Univ. Press, Baltimore, 1989).
- [19] P.W. Anderson, *Phys. Rev.* 109 (1958) 1492.
- [20] I.M. Lifshitz, *Sov. Phys. USP.* 7 (1965) 549.