Improved accuracy and acceleration of variational order-N electronic-structure computations by projection techniques

Uwe Stephan

Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801

David A. Drabold

Department of Physics and Astronomy, Ohio University, Athens, Ohio 45701-2979

Richard M. Martin

Department of Physics, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801

We have performed minimizations of the order-N energy functional of Ordejón *et al.* [Phys. Rev. B **51**, 1456 (1995)] in models of fourfold-coordinated amorphous carbon containing 512 and 4096 atoms using the first-principle local-orbital Hamiltonian of Sankey and Niklewski. The total time for performing the minimization can be significantly reduced by directly projecting the initial functions to the occupied subspace of the Hamiltonian before applying a conjugate-gradients minimization. In addition, the energies achieved in the minimization can be significantly lowered by dynamically optimizing the localization regions of the projected functions using a population rather than distance criterion. Furthermore, the projection method also provides an efficient scheme for the linear-scaling computation of the Fermi energy in semiconductors, which is needed for the computation of the projected functions as well as for the evaluation of several other order-N functionals. [S0163-1829(98)01344-7]

I. INTRODUCTION

The development of computational techniques for molecular-dynamics (MD) simulations, which scale linearly with the number of atoms in the system has become a very fruitful area of current research. For instance, these methods expanded the scope of first-principles total-energy computations to systems containing up to a few thousand atoms. Several scientifically fascinating and technologically promising model systems have been considered so far, including giant single-shell fullerenes, multishell fullerenes, tubular systems, huge complicated molecules such as the DNA as well as large amorphous tetrahedrally-coordinated model structures.[1–4]

Most of the linear-scaling computations presented so far were based on the variational minimization of a given energy functional. This functional is constructed in such a way that it has a global or at least a local minimum at the correct one-particle ground-state energy of the given Hamiltonian. However, there exist two basically different approaches for the derivation of such an energy functional. The first scheme represents the band-structure energy as the trace $E_{bs} = \mathrm{Tr}[\hat{\rho}\hat{H}]$ over the density and Hamiltonian operator, in which case the minimization has to provide the matrix elements of the density operator in some basis of localized orbitals.[5–7] In the second approach, the trace is performed over the Hamiltonian operator alone, but this has to be done within a basis of Wannier-like states that span the occupied subspace of the Hamiltonian.[8–13] These functionals minimize the trace over the matrix product $\mathbf{S}^{-1}\mathbf{H}$, where $\mathbf{H}$ and $\mathbf{S}$ are the Hamiltonian and overlap matrices with respect to these Wannier-like states. Here it is important that the minimiza-

tion can actually be performed by replacing the inverse overlap matrix with some odd-order Taylor expansion such as $\mathbf{S}^{-1} \approx 2 - \mathbf{S}$. Furthermore, it is also possible to connect both schemes by minimizing the energy with respect to the density-matrix elements *and* the basis functions.[14]

The representation of the band-structure energy as the trace over the density and Hamiltonian operators also allows a nonvariational computation of E_{bs}. This can be accomplished by means of an explicit polynomial (Chebyshev) representation of the density operator, which is derived by an analogous expansion of the Fermi-Dirac distribution function at a certain appropriate temperature parameter. This so-called Fermi-operator expansion or projection method (this latter term was chosen according to the projection property of the density operator at zero temperature), was originally developed by Goedecker and co-workers[15] for orthonormal tight-binding-like Hamiltonians. Recently, we have published[16] a generalization of this method to the case of overlapping localized basis states, in which the density operator is represented by means of the ''contra-covariant'' or upper-lower indexed Hamiltonian matrix $\bar{\mathbf{H}} = \mathbf{S}^{-1}\mathbf{H}$. (Here, $\mathbf{H}$ and $\mathbf{S}$ are the ordinary Hermitian Hamiltonian and overlap matrices in the local basis chosen.) This matrix $\bar{\mathbf{H}}$ can be computed efficiently in a linear-scaling way by solving linear systems of equations $\mathbf{S}\bar{\mathbf{H}} = \mathbf{H}$ within localized regions of space.

Because the projection method is nonvariational, trace computations can be done without the need of providing any particular ''initial-guess'' functions or matrices. This is a particular advantage compared to those functional-minimization techniques that start from approximations of certain Wannier-like states in the system. Especially in disordered systems such initial functions are likely to be far

away from the final variational solutions. The minimization of the energy functional at the beginning of an MD simulation is therefore often the bottleneck in these computations. Another advantage of the projection method compared to all variational schemes for nonorthogonal basis states is that the former method scales linearly also with the number of atoms within the localization (LOC) regions.[17] This is due to the fact that no matrix elements between different projected functions have to be computed. This may be important when larger LOC regions are needed to increase the accuracy of order-N schemes. Furthermore, the projection method when using the original basis states is directly applicable to systems with arbitrary coordination numbers.

On the other hand, the nonvariational character of the projection method may restrict the pure application of this method within MD simulations. The reason for this is that the performance of the projections in a certain time step cannot draw advantage from the projections done in previous steps. All matrix-vector multiplications involved in the polynomial representation of $\hat{\rho}$ times a certain function have to be performed in each MD step anew, even though the final projected functions may change only slightly. It is therefore desirable to connect the projection method with a variational scheme.

In this paper, we will consider the projection scheme in conjunction with a Wannier-function-based minimization approach. This connection is possible because the projection method also allows to compute Wannier-like functions by applying the density operator to selected initial functions. These projected functions can then be used as new initial functions in the minimization of the energy functional. We will show that this direct projection of initial functions to the occupied subspace, followed by an approximate orthonormalization of the projected functions, results in more efficient functional-minimization schemes. The orthonormalization, of course, breaks the linear scaling of the projection method with respect to the size of the LOC regions, but, as we will see, this fact is not very relevant due to the small number of orthonormalization steps needed (≈ 5). Furthermore, the projection as well as the orthonormalization allows a straightforward definition of *dynamical* LOC regions for the Wannier-like states, which are based on a population rather than distance criterion. With these dynamical regions, the projection method results in a significant acceleration and improvement of the accuracy of functional-minimization schemes. This connection of a projection with a variational technique therefore allows to significantly reduce the ''initial MD step'' problem mentioned above. On the other hand, we will also see that the subsequent functional minimization improves the results obtained in the projection method by reducing errors incurred due to the confinement of the matrix $\bar{\mathbf{H}}$ or the numerical representation of the density operator.

However, this paper will not yet address the question of how to find optimal initial functions for the projection in general systems. Instead, we will consider disordered fourfold-coordinated systems where physically reasonable initial functions can easily be obtained. We then compare the efficiency and accuracy of projection and variational schemes with respect to these functions. As we will show, the most efficient minimization of the energy functional for a given size of the LOC regions can be obtained by using

larger LOC regions in the projection, which are reduced to optimal regions of the desired size after the orthonormalization process.

One potential problem of using the projection method within electronic-structure computations is that the polynomial construction of the density operator requires the knowledge of the extreme band edges and the Fermi energy E_F for the given problem. Furthermore, any information about the band gap, such as the density of defect states in the gap and the size of the HOMO-LUMO gap, is desirable to define an optimum temperature parameter in the Fermi distribution function.[16] Whereas the extreme band edges are easily and accurately obtainable in a linear-scaling way with a Lanczos procedure,[18] and the size and position of the inner gap and the defect states can be ascertained by more elaborate Lanczos, recursion, or maximum-entropy schemes,[20,21] it turns out that the accurate linear-scaling determination of the Fermi energy in semiconductors is a more difficult, ill-conditioned problem. This is due to the fact that E_F in semiconductors lies in a region of low density of states (DOS), which means that small errors in the DOS or integration routines may result in inaccuracies of the order of the gap width in the predicted values of E_F. Of course, if the projection method is followed by a minimization and as long as total energies are concerned, one may expect that projecting out the complete conduction band will be significant for the accelerated minimization. Nevertheless, if the projection method is used to compute local quantities such as local-charge densities and forces, an accurate value of E_F may be important. Furthermore, apart from the projection method, many linear-scaling implementations that are based on representing the occupied subspace by an overdimensional basis[5,13,14] require an estimate of the Fermi energy. We will therefore discuss a few schemes that are candidates for the linear-scaling determination of E_F. As we will show, the projection method itself provides the most powerful approach for the $O(N)$ computation of E_F in semiconductors.

II. ACCELERATION OF FUNCTIONAL MINIMIZATIONS

In the following, we present first-principles computations done with the Harris-functional local-orbital Hamiltonian derived by Sankey and Niklewski.[22] We first consider a model for amorphous tetrahedrally-coordinated carbon containing 512 atoms, which was originally generated by Djordjevic *et al.*[23] This model was subsequently relaxed by Drabold and Ordejón[24] using the Sankey-Niklewski Hamiltonian, which resulted in a completely fourfold-coordinated network with a HOMO-LUMO gap of 4.3 eV.

For the computations presented in Fig. 1, we did not use any localization restrictions for the Wannier-like functions, which allows us to compare the ''intrinsic'' properties of the minimization schemes used. The figure shows the relative deviation of the energy functional E_{bs}^{fct} (derived by Mauri *et al.*[10] and Ordejón *et al.*[11]) from the correct band-structure energy $E_{bs}^{(0)}$ within different stages of the minimization process. The dashed lines represent the pure conjugate-gradient (CG) minimization[25] of the energy functional whereas in the case of the solid lines a projection (P) and orthonormalization (O) is used prior to the CG treatment of the functional. In both cases, we started from the same initial functions that
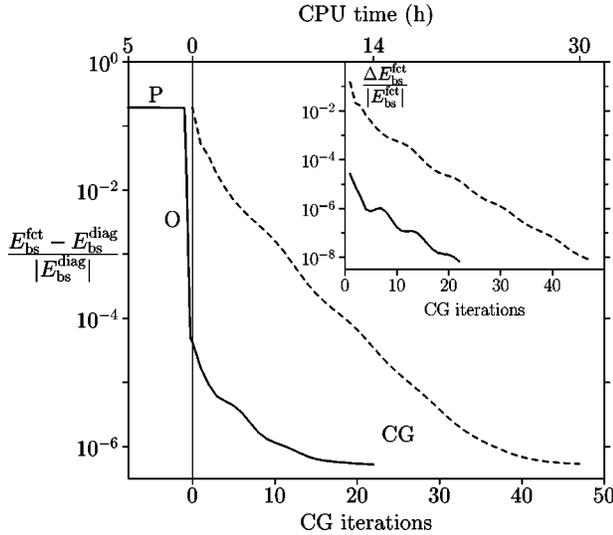
FIG. 1. Minimization of the energy functional (Ref. 11) in a 512-atom model of tetrahedrally coordinated amorphous carbon (Ref. 23) without localization restrictions. Main plot, relative deviation of the functional from the correct band-structure energy; inset, relative difference of two consecutive functional values. Dashed lines, pure conjugate-gradient (CG) minimization of the functional; solid lines, minimization consisting of projection (P), orthonormalization (O), and CG minimization. Also indicated are the CPU times (Ref. 26) needed for the different procedures.

were taken from the normalized but nonorthogonal bonding combinations of hybrid orbitals pointing into the bond directions. Within the projections for this system, we used the exact Fermi energy as obtained by direct diagonalization of the Hamiltonian. Furthermore, we used the value $\beta = 1/kT = 100$ (with respect to the Chebyshev interval $[-1,1]$) and 150 Chebyshev polynomials in the expansion of $\hat{\rho}$. These values have been chosen larger than necessary in order to get conclusions applicable also to systems with smaller gaps. According to the error analysis presented in Ref. 16, the relative error in E_{bs} for our system due to these parameters is of the order of 10^{-6}.

As can be seen in Fig. 1, the direct CG minimization of the energy functional leads to a saturated state very close to the exact band-structure energy after about 50 CG steps. Considering now the procedure that involves projection and orthonormalization, first note that the abscissa of the left-hand part of the figure was chosen such that its length corresponds to the CPU time needed relative to the time of the CG iterations.[26] An interesting fact is that the projection alone does not lead to a noticeable reduction in the functional value. This is caused by the property of the energy functional to overestimate the energy of nonorthonormal states. However, a subsequent orthonormalization of the projected states leads to an abrupt drop in the functional value. (The orthonormalization was performed with four first-order Löwdin iterations,[16] which reduced the maximum remaining overlap value to 4×10^{-5}.) The CG minimization using these orthonormalized states achieves the saturated energy state after about 20 iterations. This also minimizes the total CPU time needed. The connection of the CG technique with a projection and orthonormalization scheme therefore results in a more efficient minimization of the energy functional.
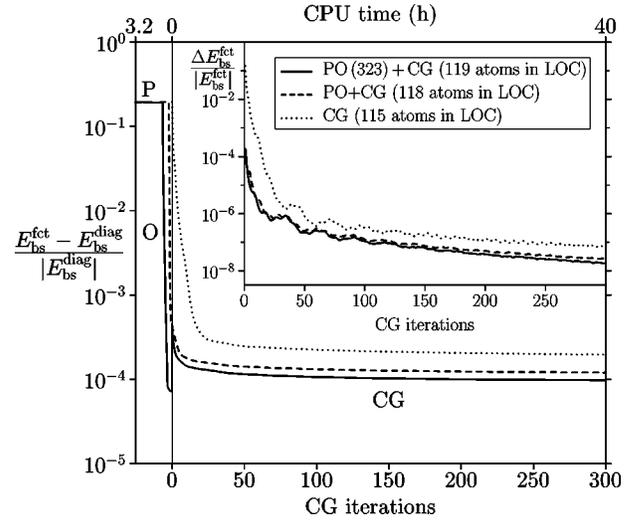
FIG. 2. Minimization of the energy functional (Ref. 11) in a 512-atom model of tetrahedrally coordinated amorphous carbon (Ref. 23) with localization restrictions. Main plot and inset as in Fig. 1. Dotted lines, pure conjugate-gradient (CG) minimization of the functional; dashed and solid lines, minimization consisting of projection (P), orthonormalization (O), and CG minimization, starting from different sizes of the localization (LOC) regions. The P and CG parts of the minimizations have been done with fixed LOC regions while during the orthonormalization these regions were dynamically optimized without significantly changing their size. For orientation, the CPU times (Ref. 26) needed are also indicated.

Also note in the inset of Fig. 1 that the minimum-energy state in both computations was reached at a relative change in the energy functional of 10^{-8}.

Figure 2 depicts the corresponding results in the presence of LOC regions for the Wannier-like functions. These localization restrictions are necessary to obtain a linear-scaling procedure. Comparing this figure with the previous one, the reader will immediately notice that the presence of such LOC regions results in a much less accurate approach to the correct band-structure energy. In the pure CG computation represented by the dotted line we used fixed LOC regions, which included all the atoms within four bond steps from a central bond. In the given system, this leads to regions containing, on average, 115 atoms. The projection computation indicated by the dashed lines in Fig. 2 started from the same LOC regions. However, in the approximate order-N orthonormalization of the projected functions (using five Löwdin iterations), we allowed a reshaping of these LOC regions to include atoms that acquired significant weight (measured as Mulliken's net population) in the orthonormalization process.[16,27] The number of atoms in the LOC regions, on the other hand, was permitted to increase only slightly, resulting in 118 atoms on average. Note again, as in Fig. 1, that performing a few orthonormalization steps on the projected functions leads to a very steep reduction of the functional value. The important new point, however, is that the dynamically redefined LOC regions are now much better capable of describing the Wannier-like functions in the system than the original regions. This can be seen from the fact that the CG minimization of the energy functional using these new LOC regions results in a significantly lower energy. This is especially important considering the very slow minimization of

the energy functional after the first ≈ 50 CG iterations. The use of (approximately) orthonormalized projected functions within the CG scheme now not only increases the efficiency of the minimization but also improves the accuracy of the energy values and the quality of the Wannier-like functions obtained.

Dynamical LOC regions can be introduced in a straightforward manner whenever the localized functions are iteratively obtained by repeatedly multiplying them with the Hamiltonian or overlap matrices.[28] Such multiplications, which are performed in the projection as well as orthonormalization procedures, repeatedly create nonzero weights of these functions outside their current LOC regions. This can be used to define new LOC regions according to these weights instead of using any fixed distance criterion. Furthermore, such dynamical regions are naturally capable of distinguishing between different "kinds" of Wannier-like functions such as those originating from σ- or π-like states or, as in the functional of Kim *et al.*,[13] from *s*- or *p*-like valence orbitals. In our scheme, using dynamical LOC regions in the orthonormalization has the advantage that the additional numerical effort related to sorting the weights and updating the localization and overlap index matrices is small compared to the total computational time because of the small number of orthonormalization steps.

To test the effect of using more flexible functions also within the projections, we used a different approach in this work. The linear scaling of the numerical effort of the projection scheme with respect to the size of the LOC regions allows one to use larger, but fixed regions within the projections and to reduce them to the desired size after the orthonormalization process. The results of this procedure are indicated by the solid line in Fig. 2. The reduction of the LOC regions leads to an increase of the functional value, of course, but the energies obtained in a subsequent minimization are again lower than the values achieved in the previous schemes. Note that the time effort for these projections and orthonormalizations within larger LOC regions is still small compared to accomplishing a significant number of CG iterations.

We emphasize here that this procedure of reducing larger LOC regions assumes that smaller LOC regions for the Wannier-like functions are needed within any subsequent treatment of these functions. For instance, this would be the case if larger LOC regions lead to prohibitive CPU times when a certain minimum relative change of the functional value is required (as in force computations). If total energies are the primary concern, one would directly minimize the energy functional after the orthonormalization or, when the functional evaluation becomes too expensive for large LOC regions, use a pure projection computation. Furthermore, our computations also showed that the initial CG iterations always result in a noticeable reduction of the functional value even when projected functions are used. This is a consequence of the fact that the functional minimization reduces errors remaining in the projection scheme due to the confinement of the matrix $\bar{\bar{H}}$ or the numerical representation of the density operator. Whenever Wannier-like functions are to be computed and the LOC regions used are not too large, the performance of at least a few CG iterations after the projection and orthonormalization is therefore recommended.

In the presence of LOC regions, the relative deviation of two consecutive functional values alone is no longer a good criterion for the accuracy achieved in the CG minimization. This is a direct consequence of the variational restrictions imposed due to the LOC regions and does not necessarily imply that a particular computation has been trapped in a local energy minimum. This latter possibility has recently been discussed by some authors,[11,13] but the use of an accurate projection scheme is one way to reduce the chance for such a trapping.[15] Our result as discussed above is that the increase in the variational freedom by selecting optimum atoms for LOC regions of a given size is of primary importance for achieving more accurate CG minimizations. Considering the band-structure energy, such a scheme allows one to perform relatively few CG iterations and thus to obtain a more efficient *and* simultaneously more accurate representation of the occupied subspace. For nonvariational quantities such as forces it is essential that the computations, which include projection and dynamical LOC regions also lead to smaller relative deviations of the functional values as shown in the inset of Fig. 2. These computations therefore converge within fewer CG steps. This reduction in the relative deviations is particularly important taking the slow decay of these values in Fig. 2 into account. We therefore expect a significant reduction in the number of CG iterations also for force computations. This, however, needs to be confirmed by further investigations.

Here, we have also investigated our projection scheme within the functional minimization for the huge "amorphous diamond" model containing 4096 atoms, which was generated by the same group.[23] For this model, the exact band-structure energy and the position of the Fermi energy cannot be computed by a diagonalization routine. In Table I, we present computations in which we applied the Fermi energy returned by an accurate projection procedure as described in Sec. III A. This computation took 15 h (Ref. 31) and yielded the Fermi energy $E_F = -2.8$ eV (with a remaining systematic inaccuracy of about 0.05 eV). Considering the DOS obtained by Röder *et al.*[30] for the same system and with the same Hamiltonian,[22] this value for E_F lies in the middle of the largest gap between defect states found between the sigma bands in this system. According to the results of this dynamical projection procedure, we could use the same values for β and the number of Chebyshev polynomials as applied to the 512-atom model.

Our results are summarized in Table I. Note that the same conclusions that have been discussed for the smaller model can also be drawn from this table. Starting with a projection calculation that uses about 200 atoms in the LOC regions (five bond steps, middle part of Table I), and reducing these regions after the orthonormalization, the functional value after three CG iterations has almost reached the value obtained after 200 iterations in the pure CG scheme (upper part of Table I). Employing more CG steps, of course, gives a better approximation of the band-structure energy. The relative change of the functional value achieved after 200 CG iterations in the pure minimization (1.4×10^{-7}) has been reached after half as many steps in the method, which includes projection. The latter method, which ends up with essentially the same number of atoms in the LOC regions, again turns out to be faster as well as more accurate. If total energies are con-

TABLE I. Different procedures for the minimization of the energy functional (Ref. 11) in a 4096-atom model of tetrahedrally coordinated amorphous carbon (Ref. 23) starting from the same bond-centered initial functions and using localization (LOC) restrictions. Upper part, pure CG minimization; middle part, minimization using reduced projected functions; lower part, minimization using unreduced projected functions. The average number of atoms in the LOC regions, the energy value of the functional, the relative change of the functional value, and the elapsed CPU time (Ref. 31) are given after generation of the initial functions (INIT), projection (P), approximate orthonormalization (ON), reduction of the LOC regions (RED), and the specified number of conjugate-gradient iterations. The projections were done using the Fermi energy $E_F = -2.8$ eV.

Iter.	Atoms in LOC	Energy (eV)	Relative change	CPU (h)
INIT		-43.750		0.0
1	114	-51.7880	1.5×10^{-1}	0.3
2	114	-52.9488	2.2×10^{-2}	0.7
3	114	-53.8063	1.6×10^{-2}	1.1
50	114	-53.8511	1.6×10^{-6}	18
100	114	-54.8526	3.1×10^{-7}	36
200	114	-54.8539	1.4×10^{-7}	72
INIT		-43.750		0.0
P	201	-43.931		4.0
ON	205	-54.8571		5.8
RED	116	-54.8367		5.8
1	116	-54.8485	2.2×10^{-4}	6.1
2	116	-54.8520	6.3×10^{-5}	6.4
3	116	-54.8535	2.7×10^{-5}	6.7
50	116	-54.8584	2.5×10^{-7}	23
100	116	-54.8590	1.4×10^{-7}	40
INIT		-43.750		0.0
P	201	-43.931		4.0
ON	205	-54.8571		5.8
1	205	-54.8583	2.4×10^{-5}	7.3
2	205	-54.8594	2.1×10^{-5}	8.9
3	205	-54.8604	1.8×10^{-5}	10
10	205	-54.8623	1.3×10^{-6}	22

cerned only, one may omit the reduction of the LOC regions and perform a few minimization steps of the orthonormalized projected functions. This is shown in the lower part of Table I where we directly applied the orthonormalized functions as computed above. When comparing these results, also note that the energy reduction in the first CG iterations after the orthonormalization becomes smaller for larger LOC regions, i.e., when the energies obtained are already closer to the correct ground-state energy. Using this scheme with 10 CG iterations, we obtained our minimum-energy value for the given system after 22 h CPU time. Let us finally emphasize here that, to our knowledge, this is the first computation of E_{bs} in such a large system with use of the Sankey-Niklewski Hamiltonian.[22]

In Sec. III, we will discuss methods for the linear-scaling computation of the Fermi energy. Here, however, let us investigate how sensitive the observed decrease in the functional value is dependent on the chosen value for E_F in the
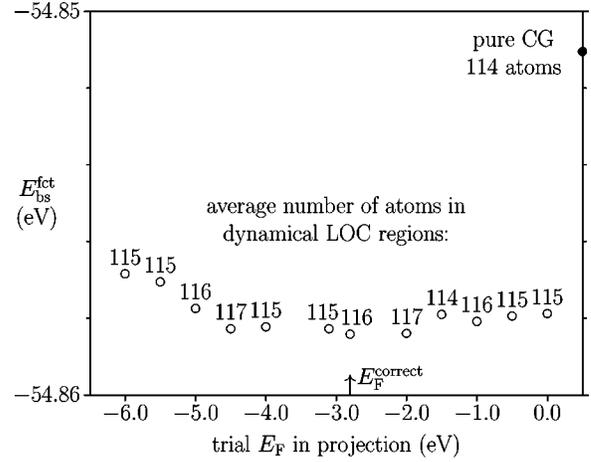
FIG. 3. Dependence of the band-structure energy value in a-C4096 (Ref. 23) obtained after 50 CG minimization iterations of the energy functional (Ref. 11) on the parameter E_F chosen within a preceding projection computation of Wannier-like initial functions. For each computation, we have also indicated the average number of atoms obtained within the dynamical definition of localization regions applied during the orthonormalization of the projected functions. The solid dot describes the corresponding energy value obtained without performing a projection computation and using fixed LOC regions. The arrow marks the position of the correct Fermi energy.

projection scheme. In Fig. 3, we have done several computations with different trial Fermi energies in the projection for the same 4096-atom model and applied the scheme presented in the middle part of Table I. Plotted are the values of the energy functional after 50 subsequent CG iterations. These values approximately describe the differences in the energy values when approaching the saturated states as in Fig. 2. We have also indicated the average number of atoms in the LOC regions obtained from the dynamical orthonormalization. These numbers were adjusted to be not smaller than but close to the number of 114 atoms (four bond steps) used in the pure CG minimization.[27]

Let us first mention that the relatively small differences in the number of atoms in the LOC regions, as a careful look at Fig. 3 reveals, still have a noticeable influence on the functional values achieved. This emphasizes our conclusion of using as large LOC regions as reasonably possible throughout the scheme. The main result of Fig. 3, however, is that the differences in most of the "projection+CG" computations shown are about one order of magnitude smaller than the difference to the pure CG minimization (solid dot in Fig. 3). Having a coarse estimate of the Fermi energy within the band gap therefore proves to be sufficient to achieve a significant improvement of the band-structure energy using a preceding projection computation. (Note that the edges of the valence and conduction bands as computed in Ref. 30 are situated at about -5 and -1 eV, while defects states exist at about -4 and -2 eV.) On the other hand, the minimum energies in Fig. 3 have been obtained for trial Fermi energies close to the correct one (at -2.8 eV). This is an expected result. Placing the Fermi energy above its correct value increases the contributions from unoccupied states remaining in the projection. Decreasing E_F below the correct value, however, increases the linear dependence of the obtained

projected functions. This leads to greater overlap values between these functions and again increases the value of the energy functional. It is clear that this effect for total energies is relatively weak when the position of E_F is wrong by a few defect states, but it will become more serious when the E_F parameter chosen approaches or is even situated within the bands. Also, as already emphasized in the Introduction, one may expect that the knowledge of the exact Fermi energy is important when local quantities such as forces are concerned. In this case, the wrong occupation of a localized state may appreciably influence quantities within the region of this state. Due to these reasons, we want to discuss some methods that can be used for the linear-scaling computation of E_F.

III. ORDER-N COMPUTATION OF THE FERMI ENERGY

As we have emphasized above, it is somewhat difficult to obtain an accurate estimation of the Fermi energy in semiconductors by a linear-scaling method. This is related to the fact that E_F itself is a purely global quantity. Since an approximation of E_F is necessary to perform a projection calculation, we want to summarize here our results for the order-N computation of E_F. As to our experience, the highest efficiency for computing E_F can be obtained by applying the projection method itself. We also want to discuss briefly our results for using a quadrature recursion method with adjusted first moments. Furthermore, we will add a few comments about factorization methods, which, though being quite efficient in certain cases, are, in general, not $O(N)$.

A. Projection method

An order-N estimate of the Fermi energy can be obtained by the projection method itself. Recall that the number of electrons in a system is given by[16]

$$N_{el} = \mathrm{Tr}[\hat{\rho}] = \sum_{\alpha} \rho^{\alpha}{}_{\alpha}, \qquad (1)$$

where $\rho^{\alpha}{}_{\beta}$ are the elements of the contra-covariant or upper-lower indexed density matrix $\bar{\boldsymbol{\rho}} = 2\,F(\bar{\mathbf{H}})$ in the original local basis, $\bar{\mathbf{H}} = \mathbf{S}^{-1}\mathbf{H}$, and $F(E)$ being the Fermi-distribution function. This allows an iterative process in which a trial value for E_F is successively improved. For the efficient performance of this process note that the Chebyshev vectors do not depend on the Fermi energy. This enables one to compute columns of several density matrices simultaneously.[29] Furthermore, as pointed out by Goedecker,[18] since evaluation of Eq. (1) only requires knowledge of the diagonal elements of the Chebyshev vectors, the orbital sums of these elements can be stored and then used repeatedly for all necessary Chebyshev expansions. This, in principle, solves the problem of computing E_F. However, without knowledge of the Fermi energy and, consequently, the HOMO-LUMO gap it is not clear at the outset which value for the temperature parameter and how many Chebyshev iterations should be used. Due to the sensitivity of the estimate of E_F on the accuracy of the computational scheme, using too smooth a distribution function or too few Chebyshev iterations may result in significant errors in E_F. On the other hand, perform-
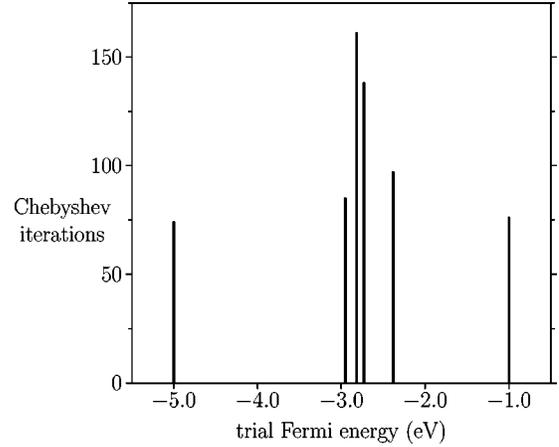
FIG. 4. Number of Chebyshev iterations needed to obtain a stable difference (as described in Sec. III A) between the number of electrons estimated at the indicated trial Fermi energies [Eq. (1)] and the correct N_{el} for the 4096-atom model of tetrahedrally coordinated amorphous carbon (Ref. 23).

ing too many iterations will reduce the efficiency of the scheme. We therefore propose a technique that dynamically increases the number of Chebyshev iterations as necessary. This can be done by using the Chebyshev approximation of $F(E)$ as an expansion, which can be broken off as soon as a sufficiently accurate value for $N_{el}(E_T)$ with E_F set to the trial energy E_T has been found. For this method, it suffices to have a stable sign of the deviation of $N_{el}(E_T)$ from the correct number of electrons. The termination of the Chebyshev expansion is possible due to the (more or less) fast decay of the Chebyshev coefficients and the property of the Chebyshev polynomials to deviate least from zero. For each Chebyshev iteration, the matrix-vector multiplications $\bar{\mathbf{H}}\varphi_{\alpha}$ are done simultaneously for all local orbitals $|\varphi_{\alpha}\rangle$, the sum of the diagonal elements of the Chebyshev vectors is stored, and the current value of $N_{el}(E_T)$ is computed. When approaching the correct Fermi energy, the number of Chebyshev iterations needed to achieve a relatively stable difference $N_{el}(E_T) - N_{el}$ increases.

The price that has to be paid for this method is to store two Chebyshev vectors for all valence orbitals ($4N$ in carbon) simultaneously. This results in somewhat increased memory requirements compared to the CG minimization of the energy functional itself.[31]

We applied this scheme to the 4096-atom model used in Sec. II, starting from the initial energy interval $[-5.0, -1.0]$ eV for the trial Fermi energy. After four weighted bisection iterations, we obtained the estimate $E_F \approx -2.80$ eV with a final systematic energy uncertainty of about ± 0.05 eV. As mentioned in Sec. II, this value of E_F lies in the middle of the largest gap between defect states in this system as computed by Röder *et al.* In Fig. 4, we depicted the energy value chosen and the number of Chebyshev iterations needed in each of these iterations. Note that the relatively large HOMO-LUMO gap in this system (situated between about -4 and -2 eV) has been identified after two bisection iterations already. The iterations can therefore be terminated earlier if an accurate DOS structure is taken into account.

The computation just mentioned used an extended range for the matrix $\bar{\bar{\mathbf{H}}}$ and yielded the final result after 15 h of CPU time.[31] Using the same cutoff[16] for $\mathbf{H}$ and $\bar{\bar{\mathbf{H}}}$, the iterations converged on $E_F \approx -2.75 \pm 0.1$ eV after about 7 h. This is practically equivalent to the result presented above.

B. Recursion method

All methods that compute the total electronic DOS in a system with a given number of electrons are, in principle, capable of providing an approximation for E_F. However, the explicit numerical integration of a DOS as obtained, e.g., by maximum-entropy schemes is not a recommended procedure because this integration inevitably introduces additional inaccuracies. Both the calculation of the DOS, which naturally is a more structured quantity than its integral as well as the integration routine require very high numerical accuracy to locate E_F within the low-DOS gap region. In many cases, small errors in these procedures give rise to wrong locations of E_F somewhere in the band tails. Another problem with order-N spectral methods comes from the fact that all these methods require the repeated use of random vectors to obtain the total DOS or related quantities.[32–37] The sensitivity in the computation of E_F implies that the Fermi energies derived with these vectors usually spread freely throughout the gap region even if the DOS itself is already well converged.

To minimize the problems when integrating a DOS, we used the widely known recursion method developed by Haydock and co-workers[20] associated with a Gaussian quadrature procedure,[38,39] which allows to obtain the integrated DOS directly from the recursion coefficients. Whereas this quadrature method is still capable of providing sufficiently well-resolved localized states within the gap as well as van Hove singularities in the bands, the bands returned in amorphous systems are naturally rather smooth.[40] An alternative scheme is the kernel polynomial method as developed by Silver and co-workers[30,35,36] in which the integrated DOS can be obtained directly by evaluating the coefficients of a smoothed Chebyshev expansion of the DOS.

To improve the convergence of quantities, which are derived by integrating over the DOS, with the number of random initial vectors, some authors have proposed to use the exactly [in $O(N)$] calculatable first moments of the DOS to detect and exclude random vectors $|\xi\rangle$ with significantly wrong moments $\mu_\xi^{(k)} = \langle\xi|\hat{H}^k|\xi\rangle$. Drabold and Sankey[37] used a CG scheme to adjust the moments $\mu_\xi^{(0)}$ through $\mu_\xi^{(2)}$ of each random vector to the corresponding exact moments $\mu^{(k)}$. Varga and Krempaský[33] earlier simply omitted all those random vectors that had first moments outside a certain tolerance interval around $\mu^{(1)}$. We compared both schemes with respect to calculations of E_F and found indeed a similar significant improvement when already using this first moment only (and the normalization). However, in the presence of an overlap matrix for the local-basis states, the CG approach is much more efficient.[41] In Fig. 5, we therefore present our results for this method. In order to be able to compute a distribution of estimates for E_F, we performed a large number of calculations using a simpler non-self-consistent local-basis LDA Hamiltonian.[42] The computations were done in the original 512-atom model.[23]
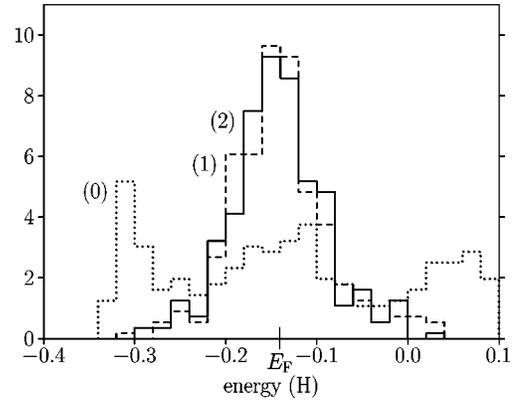
FIG. 5. Distribution of estimates for the Fermi energy in a 512-atom model of tetrahedrally coordinated amorphous carbon (Ref. 23) obtained by the recursion method with random initial vectors, using the Hamiltonian of Ref. 42. The numbers in parentheses denote the highest moment used in the conjugate-gradient adjustment of the initial vectors. Each single estimate was derived by summing partial continued fractions over 32 **k** points. The correct Fermi energy is at -0.1406 H as indicated.

Each distribution in Fig. 5 was obtained by computing the Fermi energy within 280 recursion cycles. In each cycle E_F was derived from the total continued fraction (CF) obtained by summing[43] over 32 **k** points and starting from different random initial vectors. These initial vectors were constructed by randomly taking 1's and -1's as the orbital entries.[20,32,33] (Another frequently applied choice consists in taking the orbital coefficients from a Gaussian distribution with unit variance.[34,35]) In each recursion run, we computed 40 pairs of CF coefficients, which turned out to be sufficient compared with the variations due to the random vectors.

The dotted line in Fig. 5 gives the distribution of the E_F estimates when the random vectors are normalized only (0). This has to be compared with the dashed and solid lines that represent the corresponding distributions when adjusting the first (1) and the first and second moments (2), respectively. The improvement in the convergence from (0) to (1) or (2) is clearly visible. Note, however, that the inclusion of the second moment does not significantly change the distribution compared to the first-moment case. This is in agreement with our finding that the second moments of the original random vectors vary much less (about 10%) compared to the variation in the first moments. The reason for this result should be that the second moment $\langle\xi|\hat{H}^2|\xi\rangle$, similar to the normalization product $\langle\xi|\xi\rangle$, appears as a "square" of a vector $\hat{H}|\xi\rangle$, which reduces the influence of the random signs in $|\xi\rangle$. Furthermore, unlike the first moment, the computation of the second moment requires the non-Hermitian Hamiltonian matrix $\bar{\bar{\mathbf{H}}}$ (see Ref. 41). The accuracy in fixing the second moment can therefore not be higher than the accuracy achieved in computing $\bar{\bar{\mathbf{H}}}$.

The final estimate of E_F should be derived from the total CF summed over all recursion cycles where this summation can be done very efficiently using the quadrature approach.[38] However, the convergence is often better when simply performing an average over the Fermi energies obtained from the single CF's. But note, this average is influenced to some extent by different magnitudes of the DOS below and above

E_F. For the 4096-atom model as mentioned above we obtained the averaged Fermi energy $E_F \approx -3.1$ eV after 120 recursion cycles using four **k** points. This resulted in 480 random initial vectors. Performing three averages over 40 cycles, each yielded quite dispersed estimates $(-3.4, -2.5, -3.5)$ eV, which, however, are all situated within the same gap.[30] On the other hand, the dispersion for the correctly summed Fermi energies was even higher: $(-4.2, -1.5, -4.8)$ eV. These computations took 46 h (Ref. 31) and were therefore significantly slower than the projection computations reported in Sec. III A.

We conclude from these results that the number of random vectors needed to get converged Fermi energies can be quite large if the density of states in the gap region is very small. Of course, one may comment again that the computation can be significantly abridged if the position of the defect states within the gap is known from an accurate Lanczos or maximum-entropy calculation. In this case, the computation can be terminated as soon as the HOMO-LUMO gap has been identified.

Furthermore, it is important to note that the relative error in the DOS (in a finite energy interval) obtained by using random initial vectors decreases with system size as $1/\sqrt{N}$ (see the Appendix). This effect occurs due to the statistical averaging of fluctuations in the state coefficients for computing the DOS. As a consequence, the integration over continuous bands in the DOS requires fewer random vectors for large system sizes. On the other hand, the fluctuations in the state coefficients themselves and, hence, the statistical errors in the weights of isolated states cannot be reduced by increasing the number of atoms in the system. In other words, there is no statistical averaging with increasing system size for localized states. This affects the correct evaluation of integrals over the DOS in the gap region. The exact determination of the Fermi energy, i.e., the identification of the highest occupied and lowest unoccupied states, therefore, cannot be made more efficient in terms of random vectors for larger systems.

C. Comments on factorization methods

Let us finally mention a group of methods that allow an efficient computation of the Fermi energy but, in general, have a higher-order scaling than $O(N)$. First, notice that the number of positive, negative, and zero eigenvalues of the generalized eigenvalue problem $\mathbf{HC} = \varepsilon \mathbf{SC}$ with the indefinite Hermitian matrix $\mathbf{H}$ and the positive definite matrix $\mathbf{S}$ equal the corresponding numbers of eigenvalues of the matrix $\mathbf{H}$. This can be seen by converting the general problem to the ordinary one $\mathbf{S}^{-1/2}\mathbf{HS}^{-1/2}(\mathbf{S}^{1/2}\mathbf{C}) = \varepsilon(\mathbf{S}^{1/2}\mathbf{C})$ and applying Sylvester's law of inertia.[44] The Fermi energy can therefore be derived by means of an iterative process computing the number of negative eigenvalues of the shifted equation $(\mathbf{H} - \mu\mathbf{S})\mathbf{C} = (\varepsilon - \mu)\mathbf{SC}$. This number can be obtained without explicitly computing the eigenvalues by performing a factorization of the matrix $\mathbf{J} \equiv \mathbf{H} - \mu\mathbf{S} = \mathbf{LTL}^T$, where $\mathbf{L}$ is a unit lower triangular matrix and $\mathbf{T}$ is a tridiagonal matrix, which in the most efficient factorization methods is even block diagonal. Applying Sylvester's law again, the number of negative eigenvalues of $\mathbf{J}$ equals the number of negative eigenvalues of $\mathbf{T}$, which is easily accessible. Unfor-

tunately, the factorization that clearly scales as $O(N^3)$ for full matrices does not scale linearly even for sparse matrices. The reason is that a stable method with μ not close to one of the extreme band edges requires to apply pivoting techniques, which in turn results in more or less serious fill-in of $\mathbf{J}$. The reader interested in these undoubtedly fascinating methods should consult Refs. 45 and 46.

IV. CONCLUSIONS

In this paper, we have shown how the projection method can be used to improve order-N electronic-structure computations, which are based on the minimization of an energy functional with respect to a set of Wannier-like states. Starting from a set of appropriate initial functions, the direct projection of these functions to the occupied subspace of the Hamiltonian and the dynamical definition of localization regions within an approximate orthonormalization of the projected functions appreciably reduces the value of the energy functional. Because the CPU time necessary for the projection and orthonormalization is small compared to performing a significant number of conjugate-gradient steps, the method may also serve to accelerate the functional minimization.

An optimum reduction of the functional value for a given number of atoms in the localization regions can be achieved by using larger LOC regions in the projection and reducing these regions to the desired size after the orthonormalization. This procedure maximizes the variational freedom of the Wannier-like functions for a given size of the LOC regions.

One problem when using the projection method compared to straightforward functional-minimization techniques may be that estimates of the band edges and, in particular, of the Fermi energy are necessary to perform the projection. We therefore have discussed two linear-scaling techniques with regard to the accuracy achieved and the CPU time needed when computing approximations of the Fermi energy in semiconductors. Especially the projection method itself provides a reliable tool for getting sufficiently accurate Fermi energies within a minimum of time.

Our results were obtained by performing first-principle linear-scaling computations of the energy functional[10,11] within two models of fourfold-coordinated amorphous carbon containing 512 and 4096 atoms.

APPENDIX

In this appendix, we will give a simple derivation for the dependence of the relative error in a total normalized DOS

on the number of random vectors. The projected DOS

$$n_\Phi(E) = \sum_i |\langle \Phi | \psi_i \rangle|^2 \, \delta(E - E_i),$$

with eigenstates ψ_i and eigenenergies E_i approximates the total DOS when averaging over different independent random initial vectors

$$|\Phi_j\rangle = \frac{1}{\sqrt{N}} \sum_\mu \zeta_{j\mu} |\varphi_\mu\rangle,$$

where $|\varphi_\mu\rangle$ are orthonormalized basis states and $\zeta_{j\mu}$ are random variables, which, in the simplest case, obey $\zeta_{j\mu} = \pm 1$ with equal probability.[20,32,33] The expectation value $\mathcal{E}(c_i)$ of the state coefficients $c_i = |\langle \Phi | \psi_i \rangle|^2$ then equals $1/N$ while their variance when averaging over J random vectors is given by[40]

$$\sigma_i^2 = \mathcal{E}\left(\left| \frac{1}{J} \sum_j^J |\langle \Phi_j | \psi_i \rangle|^2 - \frac{1}{N} \right|^2 \right)$$

$$= \mathcal{E}\left(\left| \frac{1}{JN} \sum_j^J \sum_{\mu \neq \nu} \langle \psi_i | \zeta_{j\mu} \varphi_\mu \rangle \langle \zeta_{j\nu} \varphi_\nu | \psi_i \rangle \right|^2 \right)$$

$$\leq \frac{2}{JN^2} \sum_{\mu \neq \nu} |\langle \psi_i | \varphi_\mu \rangle|^2 |\langle \varphi_\nu | \psi_i \rangle|^2 \leq \frac{2}{JN^2}.$$

The relative error in the state coefficients

$$\frac{\sqrt{\sigma_i^2}}{c_i} \leq \sqrt{\frac{2}{J}}$$

therefore decreases with J but does not depend on the number of atoms N. This is important for the weight of isolated states in the DOS.

To obtain the relative error in a continuous DOS function, we consider a small but finite energy interval dE at energy E, which contains about $n(E)NdE$ eigenstates where $n(E)$ is the true DOS. The sum of the state coefficients within dE can be considered to obey a normal distribution with expectation value and variance approximately given as the sums of the mean values and variances of the single coefficients. Hence, the expectation value of the random quantity $n_\Phi(E)dE$ follows as $\mathcal{E}[n_\Phi(E)dE] = n(E)NdE/N = n(E)dE$ as it has to be, but the variance in the same quantity is given by

$$\sigma^2[n_\Phi(E)dE] \leq n(E)NdE \frac{2}{JN^2} = \frac{2n(E)dE}{JN}.$$

Therefore, we obtain the relative error in a finite DOS interval:

$$\frac{\sqrt{\sigma^2[n_\Phi(E)]}}{n(E)} \leq \frac{1}{\sqrt{JN}} \sqrt{\frac{2}{n(E)dE}}.$$

These results agree with an error discussion given by Skilling[34] for the maximum-entropy scheme and by Silver and Röder[35] for the kernel polynomial method. The relative fluctuations in the state coefficients, which do not depend on N average out in a finite DOS interval and decrease with increasing system size. By virtue of this statistical effect, the total computational effort of spectral order-N methods can even be sublinear for large system sizes.[35]

---

[1] J. P. Lu and W. Yang, Phys. Rev. B **49**, 11 421 (1994).

[2] P. Ordejón, D. A. Drabold, R. M. Martin, and S. Itoh, Phys. Rev. Lett. **75**, 1324 (1995).

[3] S. Itoh, P. Ordejón, D. A. Drabold, and R. M. Martin, Phys. Rev. B **53**, 2132 (1996); S. Itoh, p. Ordejón, D. A. Drabold, R. M. Martin, and S. Ihara, Sci. Rep. Res. Inst. Tohoku Univ. A **41**, 163 (1996).

[4] D. Sánchez-Portal, P. Ordejón, E. Artacho, and J. M. Soler, Int. J. Quantum Chem. **65**, 453 (1997).

[5] X.-P. Li, R. W. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10 891 (1993); R. W. Nunes and D. Vanderbilt, *ibid.* **50**, 17 611 (1994).

[6] A. E. Carlsson, Phys. Rev. B **51**, 13 935 (1995).

[7] W. Kohn, Phys. Rev. Lett. **76**, 3168 (1996).

[8] L. W. Wang and M. P. Teter, Phys. Rev. B **46**, 12 798 (1992).

[9] G. Galli and M. Parrinello, Phys. Rev. Lett. **69**, 3547 (1992).

[10] F. Mauri, G. Galli, and R. Car, Phys. Rev. B **47**, 9973 (1993); F. Mauri and G. Galli, *ibid.* **50**, 4316 (1994).

[11] P. Ordejón, D. A. Drabold, M. P. Grumbach, and R. M. Martin, Phys. Rev. B **48**, 14 646 (1993); P. Ordejón, D. A. Drabold, R. M. Martin, and M. P. Grumbach, *ibid.* **51**, 1456 (1995).

[12] E. B. Stechel, A. R. Williams, and P. J. Feibelman, Phys. Rev. B **49**, 10 088 (1994); W. Hierse and E. B. Stechel, *ibid.* **50**, 17 811 (1994).

[13] J. Kim, F. Mauri, and G. Galli, Phys. Rev. B **52**, 1640 (1995).

[14] E. Hernández and M. J. Gillan, Phys. Rev. B **51**, 10 157 (1995); E. Hernández, M. F. Gillan, and C. M. Goringe, *ibid.* **53**, 7147 (1996); **55**, 13 485 (1997).

[15] S. Goedecker and L. Colombo, Phys. Rev. Lett. **73**, 122 (1994); S. Goedecker and M. Teter, Phys. Rev. B **51**, 9455 (1995). The projection idea was also explored by O. F. Sankey, D. A. Drabold, and A. Gibson, *ibid.* **50**, 1376 (1994).

[16] U. Stephan and D. A. Drabold, Phys. Rev. B **57**, 6391 (1998).

[17] However, when the error related to the finite extent of the LOC regions used within the projection becomes smaller than the error introduced by the confinement radius for the matrix $\bar{\mathbf{H}}$, this latter radius has to be increased as well resulting in a deviation from the linear dependence mentioned, cf. Ref. 16.

[18] As pointed out by Goedecker (Ref. 19), the extreme band edges can also be obtained by the projection method using appropriate auxiliary functions of $\hat{H}$. However, we prefer to use Lanczos's method for this purpose, which allows an efficient direct computation of these band edges with high accuracy.

[19] S. Goedecker, Rev. Mod. Phys. (to be published).

[20] V. Heine, D. W. Bullett, R. Haydock, and M. J. Kelly, in *Solid State Physics*, edited by F. Seitz, C. Turnbull, and H. Ehrenreich (Academic, New York, 1980), Vol. 35.

[21] L. R. Mead and N. Papanicolaou, J. Math. Phys. **25**, 2404 (1984);

R. H. Brown and A. E. Carlsson, Phys. Rev. B **32**, 6125 (1985).

[22] O. F. Sankey and D. J. Niklewski, Phys. Rev. B **40**, 3979 (1989).

[23] B. R. Djordjevic, M. F. Thorpe, and F. Wooten, Phys. Rev. B **52**, 5685 (1995). This reference presents and discusses the big 4096-atom model. Results for the smaller 512-atom model are unpublished.

[24] D. A. Drabold and P. Ordejón (unpublished).

[25] The CG minimization of the energy functional was done by using an exact line minimization, which in turn was implemented by means of the analytic derivative of the functional with respect to the line parameter. This procedure leads to a speed-up of about 20% compared to computing functional values at sample points on the line.

[26] The computations for the 512-atom model were done on a workstation DEC Alpha 2000 4/275 and required about 26 Mb for the projection/minimization computations without and less than 13 Mb with localization restrictions.

[27] Similar to our description in Ref. 16, we defined the LOC regions for the Wannier-like states by first defining a critical tolerance interval for the atomic net populations of these states, then computing the populations for a particular state at all atoms included in the orthonormalization, and finally putting the boundary of the pertaining LOC region within the largest population gap found within the tolerance interval. The number of atoms in the LOC regions can be adjusted by changing the position of this population interval.

[28] Dynamical LOC regions for the computation of Wannier-like states have been introduced independently by Stephan and Drabold (Ref. 16) and by Baer and Head-Gordon (Ref. 29).

[29] R. Baer and M. Head-Grodon, J. Chem. Phys. **107**, 10 003 (1997).

[30] H. Röder, R. N. Silver, D. A. Drabold, and J. J. Dong, Phys. Rev. B **55**, 15 382 (1997).

[31] The computations for the 4096-atom model were done on a 500 MHz Microway Alpha workstation and required a maximum of about 180 Mb for performing the orthonormalizations in Table I, and about 120 Mb for the computation of the Fermi energy with the projection method.

[32] M. Krajči, J. Phys. F **17**, 2217 (1987).

[33] Š. Varga and J. Krempaský, J. Phys.: Condens. Matter **1**, 7851 (1989); Š. Varga, *ibid.* **2**, 8303 (1990).

[34] J. Skilling, *Maximum Entropy and Bayesian Methods* (Kluwer, Dordrecht, 1989), p. 455.

[35] R. N. Silver and H. Röder, Int. J. Mod. Phys. C **5**, 735 (1994).

[36] R. N. Silver, H. Röder, A. F. Voter, and J. D. Kress, J. Comput. Phys. **124**, 115 (1996); R. N. Silver and H. Röder, Phys. Rev. E **56**, 4822 (1997).

[37] D. A. Drabold and O. F. Sankey, Phys. Rev. Lett. **70**, 3631 (1993).

[38] C. M. M. Nex, J. Phys. A **11**, 653 (1978); Comput. Phys. Commun. **34**, 101 (1984).

[39] S. Glanville, A. T. Paxton, and M. W. Finnis, J. Phys. F **18**, 693 (1988).

[40] U. Stephan, Ph.D. thesis, Technische UniversitätChemnitz-Zwickau, 1995, available under http://archiv.tu-chemnitz.de/pub/1996/0007/.

[41] We compute the random-vector moments $\langle \xi | \hat{H}^k | \xi \rangle$ in the following way. First note that the coefficients C_ξ of the random vectors $| \xi \rangle$ in the recursion method have to be defined with respect to Löwdin orthonormalized basis states. We then perform a Löwdin transformation of these coefficients to obtain a new representation of the random vectors with respect to the original overlapping basis states, $C'_\xi = \mathbf{S}^{-1/2} C_\xi$. This transformation is done by means of the accelerated Taylor expansion of $\mathbf{S}^{-1/2}$ described in Ref. 16 [Eq. (A3)]. Finally, the moments are given by $C'_\xi \, \mathbf{S}\bar{\mathbf{H}}^k C'_\xi = C'_\xi \, \mathbf{H}\bar{\mathbf{H}}^{k-1} C'_\xi$ with the precomputed non-Hermitian Hamiltonian matrix $\bar{\mathbf{H}}$.

[42] P. Blaudeck, T. Frauenheim, D. Porezag, G. Seifert, and E. Fromm, J. Phys.: Condens. Matter **4**, 6389 (1992); U. Stephan, T. Frauenheim, p. Blaudeck, and G. Jungnickel, Phys. Rev. B **50**, 1489 (1994).

[43] A. Crodelli, G. Grosso, and G. Pastori Parravicini, Phys. Rev. B **38**, 2154 (1988).

[44] Proofs of this and of similar statements for more general cases are given in R. G. Grimes, J. G. Lewis, and H. D. Simon, SIAM J. Matrix Anal. Appl. **15**, 228 (1994).

[45] G. H. Golub and C. F. van Loan, *Matrix Computations*, 2nd ed. (Johns Hopkins University Press, Baltimore, 1989), Sec. 4.4.

[46] I. S. Duff, J. K. Reid, N. Munksgaard, and H. B. Nielsen, J. Inst. Math. Appl. **23**, 235 (1979); I. S. Duff and J. K. Reid, ACM Trans. Math. Softw. **9**, 302 (1983); M. T. Jones and M. L. Patrick, SIAM J. Matrix Anal. Appl. **14**, 553 (1993); **15**, 273 (1994).