PHYSICAL REVIEW B VOLUME 50, NUMBER 24 15 DECEMBER 1994-II

Order-N methods in self-consistent density-functional calculations

W. Hierse* and E. B. Stechel
Sandia National Laboratories, MS 0345, Albuquerque, New Mexico 87185-0345
(Received 4 March 1994; revised manuscript received 9 August 1994)

We discuss the potential impact of N-scaling algorithms on self-consistent density-functional calculations. N-scaling algorithms can increase numerical efficiency in two qualitatively different ways: First, by eliminating the $O(N^3)$ scaling of numerical diagonalizations or orthogonalizations, and second, through the transferability of localized electronic-structure information between chemically related, but globally different systems. We argue that the second aspect is potentially of great practical importance to self-consistent density-functional calculations. We describe how the transferability of electronic-structure information can be exploited and give numerical examples.

I. INTRODUCTION

In a variety of disciplines, ranging from solid state physics to materials science to biochemistry, there is great need for a quantum-mechanically correct understanding of systems consisting of many (i.e., considerably more than 100) inequivalent atoms. Density functional theory[1] (DFT) in the local-density approximation[2] (LDA) is the most promising candidate to provide the theoretical framework for accurate and reliable electronic-structure calculations on very large systems. In conventional LDA-DFT algorithms involving full or iterative diagonalizations, electronic structure is obtained in terms of one-particle states that generally extend throughout the system under consideration. The numerical effort to compute such extended states scales as N^3. (N is the number of occupied electronic states.) This unfavorable scaling behavior (the "N^3 bottleneck") is one of the numerical obstacles that must ultimately be overcome before DFT calculations will be routinely applied to large systems.

During the past few years, a number of algorithms[3-12] has been proposed that obtain the equivalent information of a diagonalization step, but display linear scaling of the numerical effort with N. Until now, these N-scaling methods have predominantly been applied to non-self-consistent (NSC) or fixed-potential Hamiltonians [tight-binding models (TB), local basis density functional (LBDF),[14] etc.] and proven to be useful in this context, enabling calculations on systems of unprecedented size.[9,15]

One would like to realize analogous gains in efficiency for fully self-consistent DFT calculations. For instance, it can be shown that for localized basis sets and in the limit of large N, all computational steps of a DFT algorithm *except for the numerical diagonalization* can scale as $O(N)$ or $O(N \ln N)$. Hence, replacement of the $O(N^3)$ diagonalization with an N-scaling algorithm makes the entire DFT algorithm scale roughly as $O(N)$ instead of $O(N^3)$. Unfortunately, this does not necessarily mean that in practice a large gain in efficiency must result. The problem is that the above consideration refers to *asymptotic* scaling only, i.e., to the limit of $N \rightarrow \infty$. For typical *finite* values of N treatable today, DFT calculations are often dominated by the nearly N-scaling nondiagonalization parts and *not* by the diagonalization. Thus, a mere replacement of the numerical diagonalization with more efficient procedures is not likely to result in a large overall speedup.

On the other hand, the overall CPU time of any self-consistent DFT calculation is roughly proportional to the number of self-consistent iterations. We show that suitably formulated N-scaling algorithms can reduce the number of self-consistent iterations if the transferability of localized electronic-structure information is exploited. We discuss similarities and differences of this "chemical" transferability with uses of transferability already in DFT algorithms. Using N-scaling algorithms specifically adapted for the purpose of "chemical" transferability, we obtain a significant gain in numerical efficiency even for relatively small systems. The computational speedup is mostly due to a dramatic reduction of the self-consistency cycle, rather than linear scaling *per se*. Thus, N-scaling algorithms can allow for efficiency gains totally unrelated to scaling.

II. CRITERIA FOR THE USE OF N-SCALING ALGORITHMS WITH DFT

It is evident that the introduction of linear scaling into DFT algorithms will make sense only if a significant gain in numerical efficiency results for interesting problems run on available computers. However, the maximal efficiency gain obtainable from linear scaling is much smaller in self-consistent DFT calculations than in NSC calculations. This is because in self-consistent DFT, a large fraction of the computational work is associated with computational steps other than diagonalization, such as construction of the charge density, solving Poisson's equation for the Hartree potential, and integrations to determine Hamiltonian matrix elements. This nondiagonalization work can scale as $O(N)$ or $O(N \ln N)$. For localized basis sets and large N, the number of nonzero matrix elements is $O(N)$ because very distant basis functions have zero overlap. The solution of Poisson's equation to determine the Hartree potential

from the charge density is $O(N \ln N)$ if fast Fourier transforms are used. Nevertheless, the prefactor of this almost-linearly scaling nondiagonalization workload is normally so large that the $O(N^3)$ scaling of the diagonalization begins to dominate only at relatively large N. For instance, in the linear combination of atomic orbitals (LCAO) algorithm by Feibelman[16] in the parallelized version by Sears and Schultz[17] that we used for the present work, the effort associated with diagonalization becomes significant only for systems of several hundred atoms. Several hundred atoms are close to the limit of what can conveniently be treated on the largest computers available today. Thus, just replacing the diagonalization by an N-scaling algorithm is not likely to attack a significant (and often even dominant) part of the total computational work of a typical DFT algorithm. The situation is radically different in NSC applications whose nondiagonalization workload is much smaller. Here, systems corresponding to thousands of atoms can conveniently be run on computers generally available and replacing the diagonalization by an N-scaling algorithm can bring about revolutionary efficiency gains, as have been documented in the literature.[9,15]

Another issue that should be addressed in a discussion of N-scaling algorithms in DFT is accuracy. This is very complex, since the accuracy requirements in DFT algorithms are as varied as the problems to which they are being applied. In this paper, we address accuracy only within self-consistent minimial basis LCAO-DFT and nonorthogonal TB. The main focus, however, is on the question of efficiency. This leads us naturally to the concept of chemical transferability and nonorthogonal localized occupied orbitals, as described in the following sections. In Sec. V, we argue that without basing the solution on nonorthogonal localized occupied orbitals, accuracy and strong localization cannot be simultaneously achieved. Strong localization is necessary to have practical transferability. The numerical results presented in Sec. V suggest that reasonable to high accuracy is obtainable with nonorthogonal orbitals.

III. TRANSFERABILITY AND N-SCALING ALGORITHMS

Transferability of electronic-structure information enters the methodology of N-scaling algorithms in a very natural way.[8] To demonstrate this, we recall that N-scaling methods are based on the fact that electronic structure is *locally determined.*[18] This means that the structure of the electron system on a given atom is mostly determined by the properties of this atom and its immediate neighbors. In general, the influence of more distant atoms vanishes more or less rapidly with increasing distance. Thus, the electronic structure at a given place in a large system is mainly determined by the properties of a local environment whose radius is typically several interatomic distances. This insight is central to important chemical concepts such as bonding and valence.

Given the fact that electronic structure is locally determined in a physical sense, it should also be *locally computable.* To compute the electronic structure in a small part of a large system, it should be sufficient to use information only from a certain environment of this subsystem, rather than information taken from the entire system, if the system is large. Splitting up a large system into small overlapping subsystems, and treating each subsystem in this way, one directly arrives at an N-scaling electronic-structure algorithm. The various N-scaling algorithms proposed so far give prescriptions as to how electronic structure can be parametrized locally, and how the local computation of electronic structure is to be carried out.

Now it is easy to see why the localized electronic-structure information obtained from N-scaling methods should also be transferable. Assume that in two large systems A and B (which may be very different globally), two subsystems (like certain functional groups of a large molecule) and their respective local environments are similar. Then the localized electronic-structure information computed for these two subsystems should display a corresponding degree of similarity. Thus, having carried out an N-scaling calculation on system A, it should be possible to *transfer* the electronic-structure information of the matching subsystem to the "first guess" of localized electronic-structure information with which a calculation on system B is initialized. "Recycling" electronic-structure information in this way, a reduction of the number of self-consistent iterations should result from the improved quality of the start values.

It should be noted that the idea of transferring information from previous calculations to the start values of other calculations is not new at all, and is routinely being used (in one form or another) in many existing DFT schemes. For instance, Car-Parrinello-like algorithms[19] owe much of their efficiency to the complete reutilization of the electronic-structure information as preconverged start values for each new molecular dynamics time step. In the Car-Parrinello situation, the transfer occurs between incrementally different geometries of the *same* system. In the present context, however, we aim at an equally efficient transfer of electronic-structure information between *globally different* (but locally similar) systems. Another example for existing implementations of the transferability concept is the transfer of charge density and/or potential parameters from calculations on isolated atoms (or other simple systems) to more complicated situations. By contrast, we go beyond that by also transferring the information that specifies the occupied subspace, in the form of occupied localized orbitals. (In calculations using numerical diagonalizations, this information is contained in the self-consistent extended eigenstates which are not transferable due to their sensitivity to boundary conditions.) Furthermore, the information that we transfer includes a large part of the interactions between atoms such as bonding and induced polarizations. Thus, the number of self-consistent iterations to converge can be significantly reduced. In the example calculations presented in Sec. V, the cycle of self-consistent iterations is *eliminated* by suitable transfer of results from calculations on smaller systems. This holds the promise of a benefit of N-scaling algorithms that is specific to self-consistent calculations and has not really

been exploited yet. More importantly, it also addresses the workload problem that we discussed in Sec. II. Even in situations where the diagonalization workload of a DFT calculation is small in comparison with the nondiagonalization workload, efficiency *can* be gained through the introduction of N-scaling methods if transferability enables a reduction in self-consistent iterations. This aspect is of relevance to the potential usefulness of N-scaling algorithms in the context of self-consistent DFT calculations.

IV. PRACTICAL TRANSFERABILITY AND NONORTHOGONAL ORBITALS

According to the arguments of the preceding section, any N-scaling algorithm based on the locality of electronic structure should yield electronic-structure information that is, in principle, transferable. In practical terms, however, not all N-scaling algorithms are equally well adapted for that purpose. We mentioned in the preceding section that in order to apply transferability, two large systems must contain "matching" subsystems with similar arrangements of atoms. If transferability worked only for fairly large matching subsystems, its applicability would be limited to only those cases where large numbers of atoms have matching or incrementally different positions. Here, however, we want to find a formulation that allows the transfer of electronic-structure information between relatively small subsystems, such as functional groups of large molecules. As we will see, expressing the electronic structure in terms of nonorthogonal localized orbitals provides such a formulation.

To show this, we recall that the basic approximations of most N-scaling algorithms amount to a truncation of the real-space density matrix beyond a certain cutoff radius R:

$$\rho(\mathbf{r},\mathbf{r}')=0 \quad \text{for } |\mathbf{r}-\mathbf{r}'|>R \ . \tag{1}$$

This is the approximation that is essential for the $O(N)$ scaling of the numerical algorithm. The exact self-consistent solution of the density matrix (without the above approximation) would be

$$\rho(\mathbf{r},\mathbf{r}')=\sum_{i}^{N}\psi_i^*(\mathbf{r})\psi_i(\mathbf{r}') \ , \tag{2}$$

with $\psi_i(\mathbf{r})$ the ith self-consistent eigenstate and the sum running over the N lowest eigenvalues of the Kohn-Sham Schrödinger equation. (For the purposes of this general discussion, we completely suppress electron spin.) Since (2) decays to 0 as $|\mathbf{r}-\mathbf{r}'|$ increases, approximation (1) is physically motivated. [In insulating systems, this decay is exponential, as can be seen from the properties of Wannier functions in insulating systems.[20–22] In (ordered) metallic systems, the decay must be algebraic. This can be proven from the fact that the $\mathbf{k}$-occupation numbers are discontinuous.] Typically, R must be several (i.e., more than 3) interatomic distances to permit a degree of accuracy that is acceptable for DFT calculations (see the results in Sec. V).

The exact ground-state density matrix is idempotent, i.e.,

$$\int d^3r''\rho(\mathbf{r},\mathbf{r}'')\rho(\mathbf{r}'',\mathbf{r}')=\rho(\mathbf{r},\mathbf{r}') \ . \tag{3}$$

Since we want to start an N-scaling DFT calculation with a first guess that is "close" to the exact solution, this first guess should at least be approximately idempotent. If we want to assemble this first guess from subblocks of converged density matrices obtained in previous calculations *and* at the same time fulfill the condition of approximate idempotency, an analysis of Eq. (3) in connection with (1) shows that the transferable blocks must correspond to subsystems with a radius larger than R. Only then can we expect that the idempotency of the previous results will approximately carry over into the first guess for the density matrix. (This conclusion is quite independent of the particular, algorithm-dependent parametrization in which the density matrix is given.) Because R is, in general, at least three interatomic distances, the corresponding subsystems would typically contain more than a hundred atoms in three-dimensional systems. Evidently, this would virtually eliminate any prospect for the practical applicability of transferability. It follows that in any practical transferability scheme the length scale of the transferable subunits must be separated from the length scale of the localization cutoff R that determines the accuracy of the N-scaling algorithm.

In order to bring about this separation, we formulate the electronic structure in terms of nonorthogonal localized orbitals. Specifically, we parametrize the density matrix in the form

$$\rho(\mathbf{r},\mathbf{r}')=\sum_{ij}^{N}\phi_i^*(\mathbf{r})\mathcal{D}_{ij}\phi_j(\mathbf{r}') \ . \tag{4}$$

Here, $\phi_i(\mathbf{r})$ are nonorthogonal localized orbitals. The coefficients $\mathcal{D}_{ij}$ form a matrix that compensates for the nonorthogonality of the $\phi_i(\mathbf{r})$ and restores the (approximate) idempotency of $\rho(\mathbf{r},\mathbf{r}')$. From this it follows that for exact idempotency, we must have

$$\sum_{j}^{N}\mathcal{S}_{ij}\mathcal{D}_{jk}=\delta_{ik} \ , \tag{5}$$

$$\mathcal{S}_{ij}=\int d^3r\phi_i^*(\mathbf{r})\phi_j(\mathbf{r}) \ . \tag{6}$$

In this sense, the matrix elements $\mathcal{D}_{ij}$ are not independent parameters of the density matrix, but dependent quantities. Once the $\phi_i(\mathbf{r})$'s are given, the $\mathcal{D}_{ij}$'s can, in principle, be obtained through the inversion of the overlap matrix (6). [However, an actual $O(N^3)$ numerical inversion of the overlap matrix is never carried out within the N-scaling algorithm. Instead, we proceed by $O(N)$ iterative updates[12] on $\mathcal{D}_{ij}$, as described in the Appendix.] The point of the ansatz (4) for $\rho(\mathbf{r},\mathbf{r}')$ is that because of their nonorthogonality, the localized orbitals $\phi_i(\mathbf{r})$ can be made relatively short ranged.[23] In the next section, we present results showing that we can obtain very accurate density matrices with localized orbitals extending over relatively few atoms. The comparatively long range of $\rho(\mathbf{r},\mathbf{r}')$ that is necessary for sufficient accuracy, however, is maintained by a long-range cutoff that is imposed on the matrix $\{\mathcal{D}_{ij}\}$. (For details, see the Appendix.) Thus, instead of only one cutoff radius, we now have two

different cutoff radii: A relatively long-range cutoff on the $\mathcal{D}_{ij}$ coefficients that essentially corresponds to the $\rho(\mathbf{r},\mathbf{r}')$ cutoff R introduced above, and a short-range cutoff for the truncation of the localized orbitals $\phi_i(\mathbf{r})$. Since the $\phi_i(\mathbf{r})$ are the actual independent quantities parametrizing $\rho(\mathbf{r},\mathbf{r}')$, they must be the transferable subunits that carry electronic-structure information from one system into another. Because they are relatively short ranged, only small subsections of different systems have to match in order to enable the practical transfer of electronic structure.

In the Appendix, we discuss the technical procedures by which electronic structure can be obtained in terms of nonorthogonal localized orbitals. Here, we only remark that the numerical handling of nonorthogonal orbitals does not pose any particular problems and that some existing N-scaling algorithms can, in principle, be adapted to converge density matrices parametrized in the form of Eq. (4).

V. NUMERICAL RESULTS

In the first part of this section, we describe the application of transferability to self-consistent DFT calculations on molecular systems, specifically, to simple hydrocarbon molecules. We illustrate how transferability works in practice, and what the possible efficiency gain is. We find that—at least in this case—transferability allows for the complete elimination of self-consistent cycling.

The second part is devoted to a more theoretical discussion of why the use of *nonorthogonal* localized orbitals appears to be essential for obtaining transferable electronic structure at the level of accuracy required by DFT. To exemplify this, we apply a "conventional" N-scaling algorithm in the limit of strong localization as would be necessary for obtaining a practical transferable electronic structure. We show that this leads to unacceptably large errors because the short range of the orbitals causes the density matrix to be short ranged and rather inaccurate. We then demonstrate that an algorithm based on nonorthogonal orbitals, with the same strong orbital localization, can be made accurate enough for practical DFT applications. As discussed in the preceding section, inclusion of the $\mathcal{D}$ matrix allows for a long-ranged and accurate density matrix in spite of the strongly localized orbitals.

To illustrate the application of transferability in a simple yet realistic situation, we consider the heptane and decane hydrocarbon molecules depicted in Fig. 1. The chemical bond picture suggests bond orbitals as the "natural" choice for the localized nonorthogonal orbitals appearing in our parametrization (4) of the density matrix. On chemical grounds, we would expect that in the central section of the hydrocarbon chains, all the carbon-carbon bonds are fairly similar to each other, and that the same is true for the carbon-hydrogen bonds. Near the chain end, however, the shape of the bond orbitals should display certain distortions as the bonds start to "feel" the disruption of periodicity. Consequently, we would expect that a good starting guess for a long hydrocarbon chain can be generated from a short one by "cutting" the short

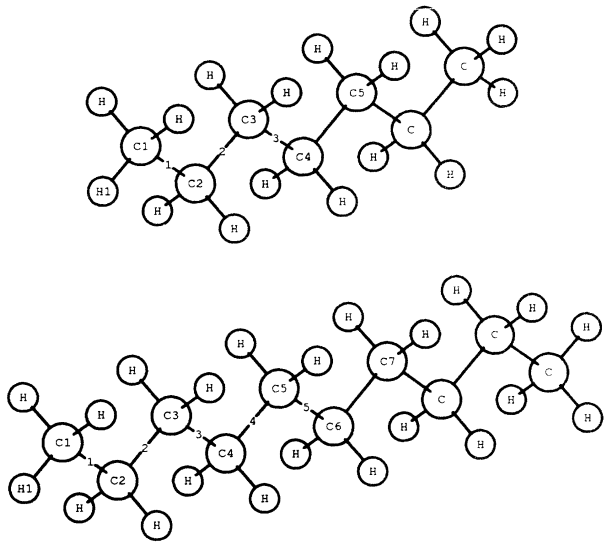

FIG. 1. The heptane (C_7H_{16}) and decane ($C_{10}H_{22}$) molecules with the labels denoting bonds and nuclei appearing in Figs. 2 and 3.

molecule in half and replicating the central section several times. This defines the first test case for transferability that we want to discuss here.

For the hydrocarbon molecules, we assume ideally tetrahedral coordination of the carbon atoms and C-C bond lengths of 1.54 and C-H bond lengths of 1.09 Å. We localize the bond orbitals by restricting their range to the two atoms connected by the bond and the first shell of nearest-neighbor atoms. C-C and C-H bond orbitals thus extend over eight and five atoms, respectively. For these relatively small systems, we impose no cutoffs on the $\mathcal{D}$ matrix because the sparseness in the Hamiltonian and overlap matrices is not sufficient to restrict the ranges of $\mathcal{H}$ and $\mathcal{S}$ significantly.

In all that follows, Hamiltonian and overlap matrices are obtained from the LCAO algorithm by Feibelman[16] in the parallelized version by Sears and Schultz.[17] We use minimal Gaussian basis sets ("single zeta") for carbon and hydrogen. All calculations are fully self-consistent: The only modification of the original LCAO program is the replacement of the diagonalization step with an N-scaling subroutine based on nonorthogonal localized orbitals. For the problems specified in this way, it is possible to converge the total Kohn-Sham energies to errors of $\Delta E_{\mathrm{KS}}=0.555$ meV/bond (heptane) and $\Delta E_{\mathrm{KS}}=0.661$ meV/bond (decane), compared to the self-consistent calculations using full numerical diagonalizations. Because of the absence of a cutoff in $\mathcal{D}$, in this particular case there is no error in the particle number.

In Figs. 2 and 3, we show the carbon-carbon bond orbitals obtained from the heptane and decane calculations, respectively. In Fig. 2, we show contour plots of bonds 1, 2, and 3 of the heptane molecule (cf. Fig. 1), and in Fig. 3 analogous plots of bonds 1–5 of the decane molecule. The bond orbitals are plotted on cross sections of dimensions 4×4 Å that lie in the plane of the carbon atoms. The positions of the nuclei within the plot cross sections

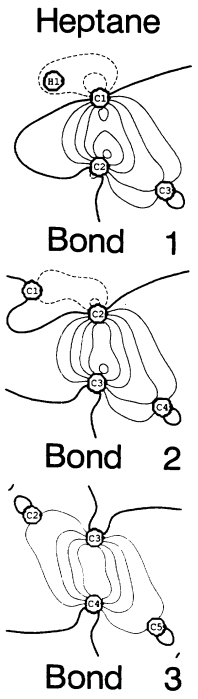

FIG. 2. Converged C-C bond orbitals for heptane.

are indicated. All the bond orbitals are L^2 normalized to unity. The amplitude difference between neighboring contour lines is 0.2, with solid thin contour lines marking positive values, broken contour lines marking negative values, and the thick solid line marking the zero contour. The bond orbitals look very plausible physically. The expected distortions of the bond orbitals close to the molecule ends are clearly visible. In the decane bond orbitals, only bonds 1 and 2 deviate appreciably from the "interior" bonds 3, 4, and 5. Also, the bond orbitals are clearly centered between the bonded atoms and are even more tightly localized than is enforced by the localization constraints. This indicates that the choice of the localization constraints is clearly adequate to the nature of the problem. More importantly, Figs. 2 and 3 also allow some observations relevant to the transferability concept: First, comparison of the heptane with the decane results shows that bonds 1, 2, and 3 of the heptane molecule are almost indistinguishable from the decane bonds 1, 2, and 3, respectively. Second, in the decane molecule, bonds 3, 4, and 5 look almost exactly the same, and are in turn very similar to bond 3 of the heptane molecule. This implies that the decane electronic structure can be generated from the electronic structure of the heptane molecule by threefold replication of the central heptane bond orbital. In other words, the heptane system already contains all the information needed to construct the decane system. This is precisely what is necessary to have orbital transferability.

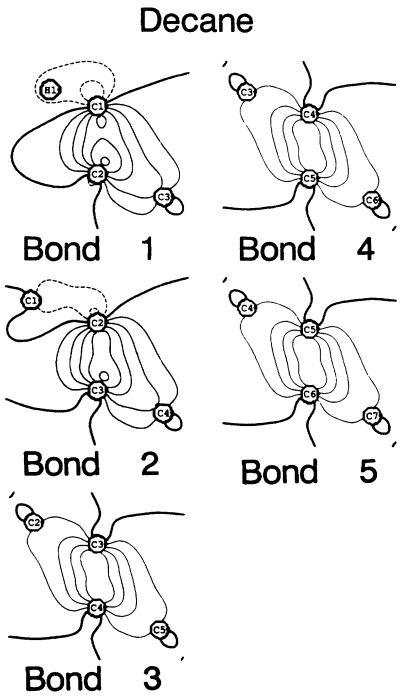

FIG. 3. Converged C-C bond orbitals for decane.

But how well does transferability work in a quantitative sense? To address this question, we transfer the converged orbitals obtained above to the dodecane molecule $C_{12}H_{26}$. Specifically, we first construct the dodecane molecule from decane bonds, then from heptane bonds. The results are summarized in Table I: In the first three rows, we give the error (as compared to full diagonalization) for the heptane, decane, and dodecane molecules; i.e., the residual error in total Kohn-Sham energy after self-consistently converging the molecule, using a nonorthogonal localized orbital N-scaling algorithm. Rows 4 and 5 are the error in the total energy obtained for dodecane, using only information from heptane and decane, respectively. In other words, we have not performed any self-consistent iterations nor any orbital updates. The total energies are very close to those obtained from a fully self-consistent calculation. In order to put the accuracy obtained into perspective, we note the following: The LCAO takes 7–8 iterations to self-consistently converge the above hydrocarbon systems. In the course of the last iteration of the self-consistency cycle, the total Kohn-Sham energy typically changes by an amount that is larger than the differences between the values in rows 3, 4, and 5 of Table I. Hence, the "first guesses" for dode-

TABLE I. Kohn-Sham energy errors for hydrocarbon systems, obtained from self-consistent iteration and/or orbital transfer.

System	ΔE_{KS}/bond	How obtained
$C\text{-}H_{16}$	0.555 meV	sc iteration
$C_{10}H_{22}$	0.661 meV	sc iteration
$C_{12}H_{26}$	0.703 meV	sc iteration
$C_{12}H_{26}$	0.707 meV	first guess (transfer from $C_{10}H_{22}$)
$C_{12}H_{26}$	0.725 meV	first guess (transfer from C_7H_{16})

cane in rows 4 and 5 display strictly the same degree of accuracy as the result in row 3 that was obtained from a full self-consistent calculation. For this particular set of molecules, perfect transferability is possible. The only numerical effort needed is the determination of the $\mathcal{D}$ matrix (which is a fast operation, see the Appendix) and the calculation of the relevant observables (such as charge density, total energy, etc.). This effectively amounts to a reduction in numerical effort by a factor of roughly 7–8. We do not expect that transferability can always enable the complete elimination of self-consistent iterations. However, the existence of realistic examples where self-consistent cycling is eliminated gives us reason to expect that frequently very significant reductions in the number of self-consistent iterations can be achieved.

Now we turn to test calculations that exemplify why nonorthogonal orbitals are important for attaining the simultaneous goal of high accuracy and practical transferability. Specifically, we compare a nonorthogonal orbital N-scaling algorithm to the method of Ordejón *et al.*[11] (The latter method is very similar to the algorithm of Mauri, Galli, and Car[10] for the choice of their parameter $\mathcal{N}=1$.) The explicit relationship between the Ordejón algorithm and the nonorthogonal orbital algorithm that we used here is clarified in the Appendix. Here, we only mention that the Ordejón method also furnishes the density matrix in terms of localized orbitals, but does not include a $\mathcal{D}$ matrix to compensate for the nonorthogonality of the orbitals. We apply the nonorthogonal and the Ordejón algorithm to the same model system, imposing strong localization for the orbitals as would be necessary to have practical transferability. It will be seen that the Ordejón algorithm does not reach a level of acceptable accuracy whereas the nonorthogonal algorithm does. In particular, the total particle number resulting from the Ordejón algorithm is too inaccurate to consider applying that algorithm (with strong localization) to self-consistent LCAO-DFT. For this reason, we choose a TB model as our example system, specifically, the nonorthogonal TB model for silicon from Mattheiss and Patel.[13]

We emphasize that the point of the following discussion is not to show that the Ordejón method is somehow "intrinsically less accurate" than the nonorthogonal scheme, which is certainly not true. Ordejón *et al.* have demonstrated that for sufficiently large localization radii, arbitrarily high accuracy is achievable with their method. Here, we merely show that in the limit of very strong orbital localization, accuracy losses are severe in the Ordejón method, but not for nonorthogonal orbitals. As a consequence, we can infer that a nonorthogonal algorithm can be used to obtain strongly localized and, therefore, transferable orbitals at high overall accuracy. The same is not true for the unmodified Ordejón method.

In the calculations that follow, we use a periodically continued 64 atom silicon supercell. With four sp^3 orbitals per silicon atom, the basis set dimensionality of the problem is 256. 128 fully occupied electronic orbitals are iterated. Their start values are taken to be the 128 bonding orbitals in the system whereas the unity matrix is chosen as the start value for $\mathcal{D}$. In order to facilitate interpretation, we formulate the problem in terms of the bonding and antibonding orbitals of the silicon tight-binding model. Specifically, we confine the orbitals to two adjacent atoms, their connecting bond, and the first shell of neighboring bonds. For the $\mathcal{D}$ cutoff[12] the minimal reasonable value is the range of $\mathcal{S}$ as determined by the cutoff of the localized states and the range of S. (The corresponding real-space distance is approximately 3.5 silicon-silicon bond lengths.) If the $\mathcal{D}$ cutoff were chosen much smaller, we could not expect $\mathcal{D}$ to converge to any meaningful approximation of $\mathcal{S}^{-1}$. The results obtained by our method and the Ordejón algorithm for this set of real-space cutoffs are listed in Table II. See the Appendix for a definition of the modified total-energy functional F, the total band-structure energy $\langle E \rangle$, and the total particle number $\langle N \rangle$ in terms of the localized orbitals. We give the deviations of the converged values of F, $\langle E \rangle$, and $\langle N \rangle$ in the presence of the cutoffs described above, from the respective converged values in the absence of all localization constraints. The errors are normalized with respect to the system size, i.e., divided by the number of iterated orbitals, 128 for this case. Since there are only N eigenvalues of $\hat{\rho}$ that are different from 0, and because these eigenvalues are constrained to be smaller than or equal to 1, $2N$ is a rigorous upper bound to $\langle N \rangle$, and $\Delta\langle N \rangle$ is necessarily negative. The quality of the solution is drastically improved by the inclusion of the $\mathcal{D}$ matrix: The convergence of the functional is improved roughly by a factor of 20, whereas the error in the energy is reduced by a factor of almost 5. Convergence of the particle number is improved by a factor greater than 500.

Still, the few numbers given in Table II do not allow for an entirely informative comparison of the two solutions. Therefore, we give in Fig. 4 the spectral decompositions of the converged density matrices corresponding to the values listed in Table II. These pictures are obtained as follows: We compute the matrix representations of the density matrices in the eigenbasis of the Hamiltonian. The filled dots (●) represent the diagonal elements of this matrix as a function of the corresponding energy eigenvalue, and the empty dots (○) are the averages of the off-diagonal elements in the columns belonging to the respective eigenvalue. The error bars denote maximal and minimal values of off-diagonal elements for a given eigenvalue. The first set of values should approach the Fermi step (the energy scale has been normalized to $\epsilon_F=0$), whereas the second set of values should approach zero. In this picture, the quality difference in the density matrices obtained becomes unambiguous. It appears fair to conclude that in the strong localization limit, the Ordejón solution does not adequately approximate the

TABLE II. Comparison of the converged solutions of the present and the Ordejón algorithms.

	Present	Ordejón *et al.*
ΔF	40.1 meV	758.6 meV
$\Delta\langle E \rangle$	39.1 meV	189.6 meV
$\Delta\langle N \rangle$	-0.90×10^{-4}	-0.49×10^{-1}

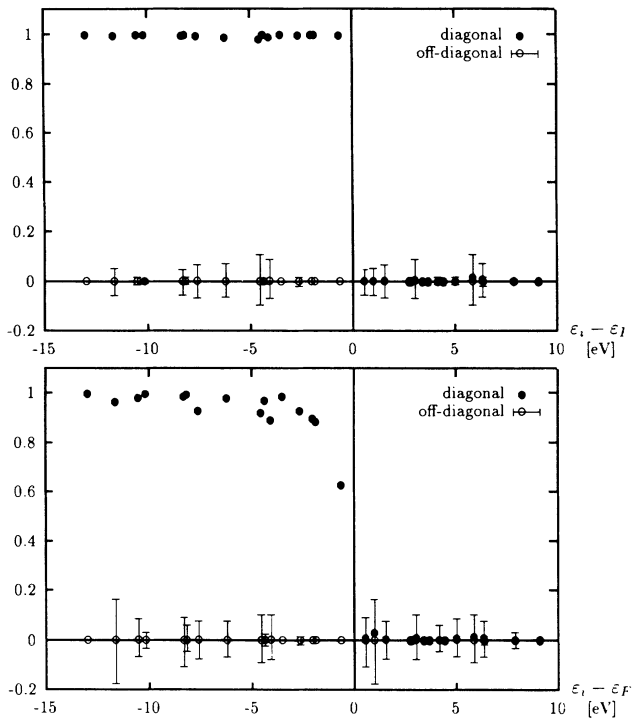

FIG. 4. Comparison of the spectral decompositions of the converged density matrices for the present approach (above) and the Ordejón method (below).

ground state: The occupation number of the highest occupied state is smaller than 0.65, and about 11 electrons are missing in the system. The solution obtained by the present method, however, falls only 0.02 electrons short of the exact electron number. While deviations from the ideal Fermi step are still visible in the spectral composition, the overall picture appears quite accurate.

VI. CONCLUSION

We applied an N-scaling algorithm based on nonorthogonal localized orbitals within a fully self-consistent LCAO-DFT calculation on hydrocarbon molecules. We simultaneously obtained high accuracy of the solutions and optimal transferability. For the systems treated here using minimal basis self-consistent LCAO, transferability could be shown to eliminate the cycle of self-consistent iterations. This was made possible by the use of nonorthogonal localized orbitals. In the strong-localization limit, these orbitals represent compact and conveniently transferable pieces of electronic-structure information. At the same time, the density matrix can have a sufficiently long range in real space to permit high accuracy of the DFT calculation. In an application to a TB model system, we showed that in other N-scaling algorithms that do not contain this separation of length scales, high accuracy and compact transferable subunits are, in general, not simultaneously achievable.

The aim of this paper is to raise attention to the fact that the usefulness of N-scaling algorithms is not limited to the numerical efficiency gains from linear scaling. We demonstrated that in the context of DFT calculations, suitably formulated N-scaling algorithms can bring about qualitatively different savings in computational effort that are related to the transferability of localized electronic-structure information.

ACKNOWLEDGMENTS

We would like to thank Dr. P. J. Feibelman, Dr. J. S. Nelson, and Dr. A. F. Wright for helpful discussions and the critical reading of the manuscript. We thank Dr. M. P. Sears and Dr. P. A. Schultz for help with running the LCAO program and the preparation of I/O facilities. W. Hierse thanks Professor J. Kübler for patient support of this work and Dr. A. R. Williams for many stimulating discussions that contributed to the ideas presented. This work was supported by the United States Department of Energy under Contract No. DE-AC04-94AL85000.

APPENDIX: THE COMPUTATION OF NONORTHOGONAL ORBITALS

There are several ways to obtain electronic structure in terms of nonorthogonal orbitals. We experimented with two different methods, on which we want to comment briefly.

The first method is the algorithm by Stechel, Williams, and Feibelman (SWF)[12] that proceeds by a state-by-state iteration in the course of which the $\phi_i(\mathbf{r})$ states are updated sequentially. Simultaneously, analytical updates on the $\mathcal{D}_{ij}$ matrix are performed. Besides the approach by Galli and Parrinello,[5] these are the only N-scaling algorithms proposed in the literature that have been explicitly constructed to handle nonorthogonal orbitals. The SWF algorithm requires very few sweeps through the set of iterated orbitals to converge. Because this method has been presented in detail elsewhere,[12] we omit an in-depth discussion here.

It is also possible to adapt other N-scaling algorithms for the computation of nonorthogonal orbitals. In some algorithms where (explicitly or implicitly) a density matrix is iterated, substitution of (4) for this density matrix leads directly to a nonorthogonal orbital scheme. To our knowledge, the basic procedure of such a generalization has not yet been reported. For this reason, we describe the fairly straightforward adaptation of the algorithm of Ordejón *et al.*[11] This algorithm is very similar to the method proposed independently by Mauri, Galli, and Car[10] for the choice of their parameter $\mathcal{N}=1$. The adaptation outlined below furnishes the nonorthogonal algorithm with which the results of Sec. V were obtained.

In the simplest case, the original method is based on a modified total energy functional[10,11] that can be written in the form

$$F=\langle E\rangle-\epsilon_{\max}\langle N\rangle\ . \tag{A1}$$

$\epsilon_{\max}$ is a parameter to be defined later. $\langle E\rangle$ and $\langle N\rangle$ are the band-structure energy and the particle number, respectively:

$$\langle E \rangle = 2 \operatorname{tr} \hat{\rho} \hat{H} \ , \tag{A2}$$

$$\langle N \rangle = 2 \operatorname{tr} \hat{\rho} \ . \tag{A3}$$

The factor of 2 accounts for spin degeneracy, as we assume a spin-unpolarized electron system. For the sake of simplicity, we also neglect Hamiltonian self-consistency for the time being. (The corresponding generalizations of what follows are straightforward.) $\hat{\rho}$ is the density operator that is related to the real-space density matrix introduced in Sec. III by

$$\rho(\mathbf{r},\mathbf{r}') = \langle \mathbf{r} | \hat{\rho} | \mathbf{r}' \rangle \ . \tag{A4}$$

Now, $\hat{\rho}$ is reexpressed in terms of an auxiliary density operator $\hat{\rho}'$:

$$\hat{\rho} = 2\hat{\rho}' - \hat{\rho}'^2 \ . \tag{A5}$$

This reparametrization has the effect that $\hat{\rho}$ cannot have eigenvalues larger than 1, which is decisive to give the functional the desired minimal properties.[10,11] With the above substitution, F becomes

$$F = 2 \operatorname{tr}(2\hat{\rho}' - \hat{\rho}'^2)(\hat{H} - \epsilon_{\max}) \ . \tag{A6}$$

In the cited references it is proven that global minimization of (A6) with respect to $\hat{\rho}'$ leads to an approximation of the ground-state density matrix, provided a parametrization of the operator $\hat{\rho}'$ is chosen that constrains it in the following way.

(1) $\hat{\rho}'$ is Hermitian.

(2) All but exactly N eigenvalues of $\hat{\rho}'$ are constrained to zero.

(3) There is an upper bound $\epsilon_{\max}$ to the expectation value $\langle \tilde{\psi} | \hat{H} | \tilde{\psi} \rangle$, where $|\tilde{\psi}\rangle$ may be any of the normalized eigenstates of $\hat{\rho}'$. [This defines the parameter $\epsilon_{\max}$ in (A1).]

Ordejón *et al.* use the following parametrization:

$$\hat{\rho}' = \sum_{i}^{N} |\chi_i\rangle\langle\chi_i| \ , \tag{A7}$$

with localized orbitals $|\chi_i\rangle$, which can be shown to fulfill the above constraints. But so does our ansatz (4), and so we substitute

$$\hat{\rho}' = \sum_{ij}^{N} |\phi_i\rangle \mathcal{D}_{ij} \langle \phi_j| \ , \tag{A8}$$

where the coefficients $\mathcal{D}_{ij}$ are now subject to the Hermiticity constraint $\mathcal{D}_{ij} = \mathcal{D}_{ij}^*$ and the nonorthogonal localized states are represented in the bra and ket notation.

Now, the $\{|\phi_i\rangle\}$ are expanded in a basis set $\{|\alpha\rangle\}$ of dimension $M = O(N)$:

$$|\phi_i\rangle = \sum_{\alpha}^{M} c_{\alpha i} |\alpha\rangle \ . \tag{A9}$$

We reexpress the functional (A6) in terms of the parameters $\{\mathcal{D}_{ij}\}$ and $\{c_{\alpha i}\}$, which will then be used as minimization variables:

$$F = 2 \operatorname{tr}(2 - \mathcal{D}\mathcal{S})\mathcal{D}\mathcal{H} \ , \tag{A10}$$

with

$$\mathcal{H}_{ij} = \langle \phi_i | \hat{H} - \epsilon_{\max} | \phi_j \rangle = \sum_{\alpha\beta} c_{\alpha i}^* H_{\alpha\beta} c_{\beta j} \ , \tag{A11}$$

$$\mathcal{S}_{ij} = \langle \phi_i | \phi_j \rangle = \sum_{\alpha\beta} c_{\alpha i}^* S_{\alpha\beta} c_{\beta j} \ , \tag{A12}$$

where $H_{\alpha\beta} = \langle \alpha | \hat{H} - \epsilon_{\max} | \beta \rangle$ and $S_{\alpha\beta} = \langle \alpha | \beta \rangle$. In (A10), we have used a shorthand notation, denoting all the $N \times N$ matrices with script letters, and all the $M \times M$ matrices with uppercase letters.

In the above form (A10), the partial derivatives of F with respect to the parameters $\{\mathcal{D}_{ij}\}$ and $\{c_{\alpha i}\}$ can be carried out. To conveniently write down the gradient expression, we define an $M \times N$ matrix c containing the expansion coefficients $\{c_{\alpha i}\}$. We represent all $M \times N$ matrices by lowercase letters. In this notation, the gradient with respect to the expansion coefficients $\{c_{\alpha i}\}$ is given by an $M \times N$ matrix g:

$$g = 2[Hc(2\mathcal{D} - \mathcal{D}\mathcal{S}\mathcal{D}) - Sc\mathcal{D}\mathcal{H}\mathcal{D}] \ . \tag{A13}$$

The gradient with respect to $\mathcal{D}$ (again in matrix notation) is given by an $N \times N$ matrix $\mathcal{G}$:

$$\mathcal{G} = 2(2\mathcal{H} - \mathcal{S}\mathcal{D}\mathcal{H} - \mathcal{H}\mathcal{D}\mathcal{S}) \ . \tag{A14}$$

g and $\mathcal{G}$ are the total gradients that can be used for the global minimization of F. As mentioned above, these minimization variables will then converge in such a way that $\hat{\rho}$ and $\hat{\rho}'$ as defined by (A5) and (A8) approach the electronic ground-state solution.

In Sec. V, we said that the c coefficients representing the localized orbitals are transferable, but the $\mathcal{D}$ matrix is not and must be recomputed. From the above equations, it can be shown that this recomputation of the $\mathcal{D}$ matrix at fixed c coefficients is much faster than a complete iteration of c and $\mathcal{D}$. First, the matrices $\mathcal{H}$ and $\mathcal{S}$ are fixed and need not be recomputed in each iteration. Second, Eq. (A14) shows that the computation of the $\mathcal{D}$ gradient $\mathcal{G}$ requires fewer operations than the computation of the c gradient g. Third, the functional (A10) is parabolic in $\mathcal{D}$ so that a conjugate-gradient minimization of the functional with respect to $\mathcal{D}$ alone requires relatively few iterations.

For appropriately imposed real-space cutoffs, the method becomes order N. For localized $|\phi_i\rangle$ states, $\mathcal{H}$ and $\mathcal{S}$ are sparse matrices. If $\mathcal{D}$ is also chosen sparse[12] (e.g., by constraining all elements connecting localized states whose spatial separation is larger than some fixed distance to zero), all the matrix-matrix multiplications among the $N \times N$ matrices scale as $O(N)$. Furthermore, if we choose a spatially localized basis set for $\{|\alpha\rangle\}$, H and S will also be sparse, and each $|\phi_i\rangle$ state will only have $O(1)$ nonzero c coefficients. Because $M = O(N)$, the generation of $\mathcal{H}$ and $\mathcal{S}$ according to (A11) and (A12) will also require a computational effort of $O(N)$, making the entire algorithm scale linearly. We assume that the number of required line minimizations and the number of self-consistent Hamiltonian updates are roughly independent of the system size, although they may vary widely for chemically different systems.

*Permanent address: Institut für Feskörperphysik, Technische Hochschule Darmstadt, 64289 Darmstadt, Germany.

[1]P. Hohenberg and W. Kohn, Phys. Rev. B **136**, 864 (1964).
[2]W. Kohn and L. J. Sham, Phys. Rev. A **140**, 1133 (1965).
[3]W. Yang, Phys. Rev. Lett. **55**, 1438 (1991).
[4]S. Baroni and P. Giannozzi, Europhys. Lett. **17**, 547 (1992).
[5]G. Galli and M. Parrinello, Phys. Rev. Lett. **69**, 3547 (1992).
[6]X.-P. Li, R. W. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10 891 (1993).
[7]M. S. Daw, Phys. Rev. B **47**, 10 895 (1993).
[8]W. Kohn, Chem. Phys. Lett. **208**, 167 (1993).
[9]L.-W. Wang and M. P. Teter, Phys. Rev. B **46**, 12 798 (1992).
[10]F. Mauri, G. Galli, and R. Car, Phys. Rev. B **47**, 9973 (1993).
[11]P. Ordejón, D. A. Drabold, M. P. Grumbach, and R. M. Martin, Phys. Rev. B **48**, 14 646 (1993).
[12]E. B. Stechel, A. R. Williams, and P. J. Feibelman, Phys. Rev. B **49**, 10 088 (1994).
[13]L. F. Mattheiss and J. R. Patel, Phys. Rev. B **23**, 5384 (1981).
[14]O. F. Sankey and D. J. Niklewski, Phys. Rev. B **40**, 3979 (1989).
[15]F. Mauri and G. Galli, Phys. Rev. B **50**, 4316 (1994).
[16]P. J. Feibelman, Phys. Rev. B **33**, 719 (1986); **44**, 3916 (1991).
[17]M. P. Sears and P. A. Schultz (unpublished).
[18]V. Heine, in *Solid State Physics,* edited by H. Ehrenreich, F. Seitz, and D. Turnbull (Academic, New York, 1980), Vol. 35.
[19]R. Car and M. Parrinello, Phys. Rev. Lett. **55**, 2471 (1985); M. P. Teter, M. C. Payne, and D. C. Allan, Phys. Rev. B **40**, 12 255 (1989); M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, Rev. Mod. Phys. **64**, 1045 (1992).
[20]W. Kohn, Phys. Rev. **135**, 809 (1959).
[21]W. Kohn and J. Onffroy, Phys. Rev. B **8**, 2485 (1973).
[22]A. Nenciu and G. Nenciu, Phys. Rev. B **47**, 10 112 (1993).
[23]P. W. Anderson, Phys. Rev. Lett. **21**, 13 (1968), and references therein.