SAND96-2060
Unlimited Release
Printed August 1996

Distribution
Category UC-705

Linear Scaling Algorithms: Progress and Promise

E.B. Stechel

Advanced Materials and Device Sciences
Sandia National Laboratories
Albuquerque, NM 87185

Abstract

The goal of this laboratory-directed research and development (LDRD) project was to develop a new and efficient electronic structure algorithm that would scale linearly with system size. Since the start of the program this field has received much attention in the literature as well as in terms of focused symposia and at least one dedicated international workshop. The major success of this program is the development of a unique algorithm for minimization of the density functional energy which replaces the diagonalization of the Kohn-Sham hamiltonian with block diagonalization into explicit occupied and partially occupied (in metals) subspaces and an implicit unoccupied subspace. The progress reported here represents an important step toward the simultaneous goals of linear scaling, controlled accuracy, efficiency and transferability. The method is specifically designed to deal with localized, *non-orthogonal* basis sets to maximize transferability and state by state iteration to minimize any charge-sloshing instabilities and accelerate convergence. The computational demands of the algorithm do scale as the particle number, permitting applications to problems involving many inequivalent atoms. Our targeted goal is at least 10,000 inequivalent atoms on a teraflop computer. This report describes our algorithm, some proof-of-principle examples and a state of the field at the conclusion of this LDRD.

3S

Acknowledgements:

I gratefully acknowledge help from Arthur R. Williams (IBM), Peter J. Feibelman (Surface & Interface Science Department), Richard M. Fye (Computational Materials Sciences Department), Peter A. Schultz (Computational Materials Sciences Department) and W. Hierse (student) for their contributions during the course of this work. I am also especially grateful to Peter J. Feibelman, J.S. Nelson (Semiconductor Material and Device Sciences) and A.F. Wright (Semiconductor Material and Device Sciences) for helping me learn density functional theory during the course of this LDRD. I also gratefully acknowledge many useful conversations with W. J. Camp (Computational/Computer Sciences & Math Center), J. S. Nelson, A.F. Wright, M.P. Sears (Computational Materials Sciences Department), D. C. Allan (Corning), R. M. Martin (Univ. of Illinois), D. Vanderbilt (Rutgers), W. Kohn (Univ. of CA., San Diego) and P. Ordejón (Univ. of Illinois).

Contents

I. Introduction 1

II. Background 6

- A. Overview of Linear Scaling 6
- B. Perspective 9
- C. Qualitative Discussion: Localized versus Delocalized Solutions 20
- D. Density Functional Theory 25
- E. Conventions and Nomenclature 28

III. Technical Approach 31

- A. Essential Ingredients of the Algorithm 31
- B. Initialization 37
- C. "Inversion" of a Well-Conditioned **S** by Conjugate Gradient Minimization 41
- D. Solving for Extremal Eigenvalues in Non-orthogonal Representations: The Gap 47

IV. Metallic System: FCC Calcium 53

V. Linear Response 57

VI. Conclusion 60

Appendices

A. Update Equations 61

B. Minimization of the Total Band-Structure Energy with Mixing Angle 64

C. Update to the Charge Density 67

D. The Residual Vector 69

References 72

Figures

1. Convergence of average energy in silicon supercells 38
2. Convergence on an expanded energy scale for silicon 39
3. Convergence with cutoff 43
4. Converged NOLOs in Heptane and Decane 45

5. Converged NOLOs from a class of organic molecules 46
6. Convergence of extremum energy levels 52
7. Convergence of average energy in calcium supercells 54

Linear Scaling Algorithms: Progress and Promise

I. INTRODUCTION

Meaningful simulation of the microscopic behavior of condensed matter systems requires a reliable description of inter-atomic forces. More than twenty years of experience have shown that the local density functional approximation (LDA) can accurately predict structure and properties of many classes of materials, including crystals, inter-atomic compounds, surfaces and small molecules. Consequently LDA represents a general framework for the study of important material science issues, e.g., crystal growth, defect properties, catalytic chemistry and materials degradation and failure. Since the seminal 1985 article of Car and Parrinello [1], a great deal of effort has been devoted to making LDA calculations more efficient. However, until relatively recently, all LDA algorithms scaled with N^3 where N is the number of inequivalent atoms in the system. The outlook for treating large systems changed with the publications of Yang [2] and of Baroni and Giannozzi [3] which showed that linear scaling with N is in principle attainable in LDA. Although there are more rigorous, and in general more accurate methods than LDA for the solution of electronic structure problems, none clearly shows the promise of applicability to even modestly complex systems, since they all scale rapidly with system size.

Following the papers of Yang [2] and of Baroni and Giannozzi [3], a number of new algorithms have been suggested [4-21]. As a result of all these efforts it appears certain that the LDA approach to electronic structure calculations can be implemented in a way that scales optimally, i.e., almost linearly (NlnN) with system size. It should be kept in mind, however, that how much computational work is required will, and should be, strongly dependent on the degree of accuracy that is acceptable. Our goal has been to calculate only the physically relevant information which itself scales linearly with N. It was never our goal to achieve a linear plot of CPU vs. system size for all N but to find an algorithm that can be used to study interesting problems that are currently intractable, and to do this on today's computers. It is expected that to achieve reasonable accuracy for small to modest N that the algorithm will scale faster than N. Indeed applying linear scaling algorithms only makes sense when the electronic structure can localize within the unit cell.

In this report we describe a new scheme for the solution of LDA problems that not only asymptotically scales almost linearly with N, but has a reasonably small prefactor. The method is designed to work for insulators and with suitable generalization for metals. The progress reported here represents an important step toward the simultaneous goals of linear scaling, controlled accuracy, efficiency and transferability. The algorithm we propose yields explicitly *non-orthogonal* localized Kohn-Sham orbitals, *NOLO*s [11,12]. The key ingredients are (1) a generalized Jacobi rotation that yields a "state by state" iterative solution for *fully occupied, non-orthogonal localized* orbitals, (2) a conjugate gradient solution for an approximate inverse of a well-conditioned overlap matrix; this determines the complement [22] basis, and (3) for the complement basis an update procedure that enhances the efficiency of working within a non-orthogonal representation. The density matrix is treated so as to guarantee that the band energy is a variational upper bound for a fixed potential. This is accomplished in a manner that is related to the transformation employed in the work of Li, et al. [6]. It is also related to the energy functional methods employed by Mauri, et al. [8] and Ordejón, et al. [9]. Additionally, the space of unoccupied orbitals is never explicitly constructed. Nevertheless, we give a prescription for obtaining the energy gap by explicitly determining (with order N work) the highest (HOMO) and lowest (LUMO) energy level of the occupied and unoccupied space, respectively, or, in a metal, a small N-independent number of eigenstates near the Fermi level. To date, no other linear scaling approach has suggested ways to accomplish this.

In this report, we also give another novel, untested prescription for an efficient approach to treating metals or other open-shell systems, as well as an approach for going beyond LDA to linear-scaling linear-response theory. Current approaches for linear response require the computation of the occupied (and sometimes unoccupied) stationary states, which are unavailable in all linear-scaling algorithms. We show that the stationary states do not need to be constructed and that all linear responses can be written in terms of a unperturbed operator **A**. Furthermore, the operator is only needed implicitly by its operation on a vector **x** (which depends on the perturbation).

The principal aim of what follows is not only to optimize the numerical solution of the LDA variational problem, but also to find an algorithm that, in the course of solving for an electronic structure, will produce transferable, and thus physically meaningful "bond orbitals." If one knows that the bond between atoms A and B is well represented by a particular localized bond orbital for a wide variety of materials in which atoms A and B are nearest neighbors, then making a good first guess for the LDA density will generally be easier and the energy minimization problem can be solved more efficiently. In addition, if our bond orbitals are physically meaningful, then computations may point the way to still more efficient approximate schemes for total energy minimization.

With these ideas in mind, we focus on several issues:

1. Is there a way to optimize the transferability of the set of bond-orbitals that are produced by the algorithm? For example, this concern leads us to frame the algorithm in terms of non-orthogonal orbitals. The oscillatory behavior necessary to produce orthogonality is poorly transferable.
2. Is there any meaning to the concept of optimally localized orbitals in the method, and if so, how does one optimize? Galli and Parrinello [4] construct a set of localized occupied orbitals by adding a very weak localizing potential to the hamiltonian. This "bucket" potential has virtually no effect on the computed energy. We do not use a localizing potential, (other than the cutoffs already built into the algorithm) having found that our iteration scheme does not delocalize an initially localized set. The generality of these promising results remains to be tested or proven. We have demonstrated this on silicon and calcium supercells, alkane chains and on some organic molecules.
3. What are the limitations of a bond-orbital basis for an open shell system, i.e., a metal? Do we need to include some delocalized orbitals in the occupied set for a metal? If so are there order unity or order N of them, and is the work that is necessary to define them order N? Or must we consider partial occupancy?

A related concern is:

4. How do we choose a starting orbital set that does not, by virtue of its symmetry, force the algorithm to converge to an excited electronic state? This question is related to the nature of bond orbitals in metals. There, choosing reasonable orbitals and occupying them symmetrically can easily put the system in a state that is nearly orthogonal to the groundstate and thus converges to an excited state. Once one abandons solving for individual Kohn-Sham eigenfunctions, or explicitly filling density up to the Fermi level (essential features of several of

the order N schemes that have appeared in the literature), then one needs to persuade oneself that only the lowest eigenstates of the system are occupied. If there is no way to know that an excited state orbital is occupied, the meaning of the computed total energy is in question. We provide an order N scheme for dealing with this problem. We do not give a "black-box" prescription for the starting set of orbitals. Nevertheless, in a very large class of systems, relatively simple chemical arguments show unabiguously what the physically relevant localization prescription should be. In metals, this is much less obvious and bears further investigation. Even in some closed shell systems, we expect it won't be obvious and that some experimentation will become necessary. It is our sense that formulating the electronic structure in terms of singly occupied orbitals would provide a lot of flexibility and eliminate much ambiguity. Indeed this becomes essential in local spin density calculations. It also provides an avenue to include local correlation within linear-scaling.

5. As the systems being studied get larger and more complex, self-consistency can lead to a "charge-sloshing" problem which slows or even precludes convergence of the self-consistency loop. This can be avoided by updating the charge density and the resulting effective potential during the iterative process for the orbitals. Thus the potential never changes by a large amount in any update. This leads us to formulate our algorithm as a state-by-state iteration (local relaxation and local updates) rather than global relaxation where all orbitals change before the potential changes. This has been shown to be important to minimize or even eliminate charge-sloshing instabilities [23].

The remainder of this report is structured as follows: Section II contains (A) a brief overview of linear scaling, (B) a perspective of other work in this field as reported at an international workshop dedicated to linear scaling, (C) a qualitative discussion of localized versus delocalized orbitals, (D) background information for the orbital approach to LDA and (E) definitions and conventions for working with non-orthogonal representations. Section III is subdivided into (A) the essential ingredients of the algorithm developed in this LDRD (with details relegated to one of four appendices), (B) the starting set of orbitals or initialization of the iterative scheme, (C) "inversion" of a well-conditioned metric by conjugate-gradient minimization and (D) solving for eigenvalues and eigenfunctions near the Fermi level. In Section III, periodic clusters of silicon in a tight-binding[1] representation and hydrocarbons within minimal basis self-consistent LDA are used as concrete examples. Section IV treats a metallic system and discusses the special issues that arise as well as a novel new approach. Section V discusses a novel approach to linear-response. Section VI is a brief recapitulation. In four appendices we present (A) update equations for the complement basis and relevant matrices (B) the determination of the optimum mixing angle in the generalized Jacobi rotation (C) the update equations for the charge density and (D) the mechanics of determining the residual vector for the generalized Jacobi rotation.

---

[1] Tight binding is a terminology for a fixed (non-selfconsistent) matrix representation of the electron structure problem. The Hamiltonian matrix elements are either determined semi-empirically or fit to accurate self-consistent calculations. The representation may or may not be orthogonal. If it is non-orthogonal the matrix elements of the overlap matrix become additional fitting parameters. The fits usually include first and second neighbor and occasionally third neighbor interactions.

II. BACKGROUND

A Overview of Linear Scaling

We distinguish between algorithms proposed in the 90's that have the focused goal of achieving linear scaling in *accurate* first principles calculations from methods proposed in the 60's and 70's that are in many ways very similar [24-28]. The goals of the early methods were, however, more to justify semi-empirical calculations and to provide for fast but *approximate* methods.

Yang's [2] algorithm relies on the principle of "divide and conquer. Yang solves the Kohn-Sham equations in overlapping small regions of the system using a basis set of atomic orbitals (LCAO representation), thereby solving for local rather than global stationary functions. If each region has $M < N$ atoms and there are $\alpha N/M$ such regions ($\alpha > 1$ accounts for the overlap and scales roughly as the surface area of the region or $M^{2/3}$ in three dimensions), the computational effort then scales as αNM^2m^3 or $NM^{8/3}m^3$ where m is the number of basis functions per atom. (The number of regions times a diagonalization of order $(mM)^3$ for each region.) Accuracy is improved by increasing both M and α, but the variational principle is not rigorously preserved, i.e., the total energy is not necessarily an upperbound. Global communication in the system occurs only through the definition of a common Fermi energy (the energy that separates the occupied from the unoccupied states). Yang [2] has some very impressive results on a large number of molecules and is one of the very few and definitely the first to have achieved full self-consistency while working with a linear scaling algorithm.

Baroni and Giannozzi's [3] method represents the hamiltonian on a real-space grid and solves directly for the charge density, using a Lanczos algorithm [29] to determine $\langle r|G(z)|r\rangle$ where G(z) is the Green's function for the Kohn-Sham hamiltonian (See Eq. 2 in Section II.D below). The solution at any point is independent of the solution at any other point given a fixed potential. The charge density and band-structure energy are then determined via contour integration on z. The hamiltonian is very sparse in this representation (provided a local approximation to the kinetic energy is used). The sparseness of the representation and the Lanczos algorithm allows convergence of each $\langle r|G(z)|r\rangle$ with work that is independent of the size of the system. The Lanczos algorithm [29] or recursion [30] is a natural way to achieve linear scaling, in that it involves evaluating the repeated action of a sparse matrix on a localized vector. In a Lanczos algorithm linear scaling is obtained only with an orthogonal representation, making a real space grid most natural. It should be noted, however, that some proposals have been made that utilize non-orthogonal tight-binding representations [31] without destroying the favorable scaling.

In both the divide and conquer and recursion methods explicit knowledge of the occupied and the unoccupied subspaces is obtained. In particular, in Baroni and Giannozzi's approach [3], the orthogonal representation is a real space grid and because there are so many more grid points than occupied states, most of the information acquired concerns the unoccupied states. Since all information about the total energy and the force is embodied in the occupied electronic states it seems likely that excess computational effort is being expended.

A third type of approach has been proposed by Li, et al. [6] and by Daw [7], based on solving directly for the density matrix, $\hat{\rho}$, in a sparse localized orthogonal fixed representation such as in tight-binding calculations. In this approach the idempotency constraint, $\hat{\rho}^2=\hat{\rho}$, is replaced by a transformation that constrains the eigenvalues of the density matrix to fall between 0 and 1 consequently the variational principle is preserved, i.e., the total energy is a rigorous upperbound. Minimization of the energy then drives the eigenvalues to be all either 0 (unoccupied) or 1 (occupied), idempotency, provided the zero of energy is adjusted to lie between the occupied and unoccupied states. Again explicit knowledge of both the occupied and unoccupied subspaces is obtained, implying that the method's prefactor is likely to be unnecessarily large. This is an iterative global algorithm which is to say that all the parameters (density matrix elements) are updated simultaneously. Linear scaling is obtained because the representation of the density matrix is sparse (hence order N non-zero elements) and the algorithm involves matrix multiplications of only sparse matrices.

Galli and Parrinello [4] introduced a localized orbital formulation of the Kohn-Sham equations which, in principle, solves for orbitals that are fully occupied. No information about unoccupied states is determined. This formulation is very closely related to the chemical pseudopotentials and chemically invariant orbitals that were proposed by Adams [24], Gilbert [25] and Anderson [26,28] and extensively used by Bullett [27]. This point was not recognized in the Galli and Parrinello [4] paper. The density operator requires knowledge of the inverse overlap matrix which is approximated as a Taylor series expansion of $(1-x)^{-1}$. Recent variational implementations of the localized orbital formulation by Mauri, Galli and Car [8] (in a tight-binding representation) and by Ordejón, et al. [9] (in a non-orthogonal minimal atomic orbital basis and no self-consistency), achieve minimization of the total energy only for orbitals that are orthogonal or, in other words, for generalized Wannier orbitals [32-34]. Constraining the occupied orbitals to be localized makes this class of algorithms scale linearly roughly as NMm where M is the number of basis functions per orbital and m is the number of orbitals that physically interact. Refs. [8,9] achieve orthogonality and energy minimization simultaneously, without imposing additional constraints. These iterative localized orbital algorithms, like the density matrix algorithms, are global, i.e., all orbitals are updated at once. Kohn [10] has also given a prescription for determining generalized Wannier orbitals with linear scaling.

Another approach, related to the generalized Wannier orbitals and to the chemical pseudopotentials, explicitly solves for *non-orthogonal* localized occupied orbitals. No explicit knowledge of the unoccupied states is required or obtained. It can be solved either in a state-by-state algorithm [11] (described herein) or in a global algorithm [12] much like that for the generalized Wannier orbitals of Mauri, Galli and Car [8] and Ordejón, et al. [9]. The localized orbital approaches (orthogonal or not) could be considered algorithms in search of the ultimate minimal basis set (a concept introduced by Frost [35]), N orbitals to represent N occupied stationary states.

B. PERSPECTIVE

In this section, we discuss in greater detail methods now in the literature. We do this in the context of a synopsis of a timely and relevant international workshop at which we presented the work from this LDRD in an invited presentation. The overall objective of the CECAM Workshop, entitled “Linear Scaling Methods in Electronic Structure Theory” was to exchange technical information regarding the relatively new and exciting field of linear scaling electronic structure theory. Twenty invited presentations and eleven short discussion papers were presented in three days. This international workshop was valuable in assessing the current level of algorithmic development. Much progress over the preceeding three years has been realized and enough experience has been gained. For tight-binding applications the issues are conservation of energy and of charge and the slow convergence. In LDA applications these same issues exist plus the additional issue of what is the best real-space basis or representation. A true basis retains the variational principle. However, most real-space grid (the ultimate in localized basis functions) methods sacrifice the variational principle. There are two themes that are pervasive in the field of linear scaling. The first is W. Kohn’s [20] concept of “near-sightedness” which is the physical concept that allows for localization regions or finite range spheres of influence, localized orbitals, finite-range density matrices, etc. Without this there would be no field of linear scaling. The second concept is less fundamental but nevertheless quite practical and that is “purification” which is a convenient way to handle the nasty constraint of idempotency. With basically these two concepts an amazing array of approaches has been and continues to be developed.

The specific purpose of the workshop was to discuss: (1) strengths and weaknesses of existing methods; (2) relationships among existing methods; (3) new approaches and methods; (4) progress towards practical and flexible linear-scaling LDA (local density approximation) calculations; (5) Choice of localized basis sets; (6) treatment of screening self-consistency; (7) convergence rates and (8) parallel computing implementations. The intended emphasis was to be more on methods than on applications.

R. Haydock, who probably has the longest history of using and promoting the local nature of electronic structure [30], presented work entitled “Electronic Structure Calculations for Infinite Systems: Problems and Some Solutions". He introduced an error theory analysis. He explained that the error should be an exponentially decreasing function of the number of machine cycles with a theoretical limiting value that goes as $\frac{1}{\sqrt{N}}$ where N is a measure of the memory of the computer. To achieve this ideal performance, one needs a problem with sparse coupling, localized basis sets and one must also use a dynamic basis, that is a basis which changes during the course of the calculation. The elementary operation in the recursion method is $\mathbf{v} = \mathbf{H}\mathbf{u}$, with **u** and **v** vectors and **H** an operator. Sparsity implies that each component of **u** affects only a finite (few) number of components of **v**. The way to look at this is to use graph theory, beginning with one (or a small number) element in **u**. Then one needs a growth rule (for **v**) and a way of killing off elements of **u** so **v** does not grow too large. In essence the idea is very simple but very powerful. He explained that, despite killing elements off, errors do not accumulate. The error is always determined only by the largest element ignored. Several types of problems lend themselves to this way of thinking. Local density of states in tight-binding electronic structure is perfectly suited. Other applications include classical Liouvillian, functional derivative on an energy functional, and commutation of the hamiltonian with an operator.

S. Goedecker [13] introduced a method based on a Chebychev polynomial expansion of the finite temperature density matrix. This is one of a class of methods which could be considered moment methods. Reference [21] describes a related method. In Goedecker's method [13], the order of the polynomial expansion goes as (emax-emin)/de. However, if the basis is such that emax >> E_f which is likely to lead to an unphysical scaling. A similar problem arises for similar reasons in the method of Ref. [3] (see below). Goedecker introduces a localization region and emphasizes that the size of the region and the degree of the polynomial are independent quantities. The size of the localization region is proportional to the electronic correlation length and is physically determined, a point that is often overlooked. Goedecker's localization region seems closely related to the buffer region used in Yang's [2] method (see below). Apparent advantages of this approach include: (1) intrinsically parallel, (2) low crossover with existing N^3 methods, (3) perfect energy conservation in molecular dynamics, (4) dividable - basic quantities are sum of local ones, (5) can handle finite electronic temperature naturally and (6) numerically stable.

Baroni and Giannozzi's [3] method represents the hamiltonian on a real-space grid and solves directly for the charge density, using a Lanczos algorithm [29] to determine <r|G(z)|r> where G(z) is the Green's function. The charge density and band-structure energy are then determined via contour integration on z. The hamiltonian is very sparse in this representation (provided a local approximation to the kinetic energy is used). The sparseness of the representation and the Lanczos algorithm allows convergence of each <r|G(z)|r> with work that is independent of the size of the system. The Lanczos [29] algorithm or recursion [30] is a natural way to achieve linear scaling, in that it involves evaluating the repeated action of a sparse matrix on a localized vector. In a Lanczos algorithm linear scaling is obtained only with an orthogonal representation, making a real space grid most natural. It should be noted, however, that some proposals have been made that utilize non-orthogonal representations [31]. Baroni gave a short presentation on his method entitled "what is good; what is (still) bad". While, as his title suggests, there are some short-comings to the method, it is fair to say that his method did create a lot of excitement and could be credited with getting the field started. The largest problem that I see with a simplistic implementation is the unphysical scaling. As in the graph theory approach that Haydock talked about, the first vector starts off with one element (basically a delta function at a point r). At each successive iteration the radius of elements grows as h (the spacing between points). Hence for L iterations the work goes as $L^{d+1} \sim (1/h)^{d+1}$ where d is the dimension of space, and this is for each point. L goes as 1/h because the range that needs to be spanned is determined by the physics. Furthermore the total number of points goes as $(1/h)^d$, hence the method scales as $(1/h)^{2d+1}$ which is unphysical. This unphysical scaling is related to that in Goedecker's [13] method. It has more to do with emax (the maximum energy which depends sensitively on the choice of basis). As h is decreased, emax increases without bound. Some of this is physical but much of it is not. The Green's function being computed involves the entire energy spectrum but the charge density only "cares" about the energy spectrum below the Fermi level. As h decreases, the useful information becomes a smaller and smaller fraction of the total computed information. Baroni also talked about embedding in a free particle space outside of some radius R_0. This effectively mitigates a problem of oscillatory convergence. The embedding also provides for faster convergence; but curious enough, embedding after 40 iterations was only marginally better than after 20 iterations. In other words, improving the accuracy is difficult. In conclusion Baroni noted that (1) real-space local implementation of LDA is good - needed for O(N) (2) O(N) scaling is trivially achieved - not great crossover but okay (3) independent points are calculated independently - good for parallelization but a large waste of information; i.e., no communication but no intelligence (4) short-sightedness is automatically accounted for (5) too much detail; i.e., should be able to put in different levels of accuracy for different distances and (6) it is not variational. He emphasized the need in most methods to take advantage of different levels of accuracy, in this case for different distances. This concept is also valuable in the localized orbital methods and in Yang's [2,18] method (in the buffer region).

W. Yang's [2,18] approach relies on the principle of "divide and conquer". It solves the Kohn-Sham equations in overlapping small regions of the system using an LCAO representation, thereby solving for local rather than global eigenfunctions. If each region has M atoms and there are $\alpha N/M$ such regions ($\alpha > 1$ accounts for the overlap and scales roughly as the surface area of the region or $M^{2/3}$); the method then scales as $\alpha N M^2$. Accuracy is improved by increasing both M and α. Within density functional theory, this is the first true linear scaling approach; it dates back to 1991. Yang introduces his method with the relationship to Thomas-Fermi Theory, which is clearly O(N). Originally he introduced a partitioning of real space (no approximation). Within each partition he approximates the hamiltonian with a local cluster hamiltonian. He shows that this is accurate as long as the "cluster" is large enough (has enough buffer atoms); i.e., the charge density is only used far enough from the surface region. This again is the near-sighted nature of electronic structure and this buffer region is closely related to the localization region that Goedecker uses. In his new approach he still partitions space but partitions the density matrix. $\rho(r,r') = \varphi_i(r)\rho_{ij}\varphi_j(r')$ (summed over repeated indices). The partitioning is such that $P^{\alpha}{}_{ij}=1$ ($i \varepsilon \alpha$), 1/2 ($i \varepsilon \alpha$ but j not $\varepsilon \alpha$ or vice-versa), and 0 (i not $\varepsilon \alpha$, j not $\varepsilon \alpha$,). The α region is typically an atom. He points out that the original method with the real space partitioning is a Hirshfeld-type population analysis while the second method with the density matrix partitioning is more like a Mulliken population analysis. The advantages of the second method are no integrals over the partition function. Furthermore it can be used in other applications such as tight-binding, Hartree-Fock and semi-empirical methods. He also showed that his method can obtain forces. One curious thing was: he tried "purifying" the density matrix. However, this made the results worse. One final point that I would like to note is that one diagonalization is replaced by N smaller diagonalizations. The one scales as $(mN)^3$ (where m is the number of basis functions per atom and N is the number of atoms). Each of the N diagonalizations scale as $(mM)^3$ where M is independent of the size of the system and is the number of atoms in the localization region. The cross-over should then be roughly $N^2 = M^3$. For M=25, N=125 but for M=100, N=1000. If Baroni's point of using different levels of accuracy at different distances were used, however, this cross-over could be improved. For example less basis functions per atom could be used in the region further than one neighbor away.

F. Mauri discussed a variant [15] on localized orbitals[8]. He indicated that the open problems involved local minima and slow convergence. While the original algorithm [8] solved only for occupied states, the newer algorithm [15] uses more orbitals than occupied states. In other words, it includes some empty states. They find that the additional freedom solves the "local minimum" problem and improves convergence. By adding Fermi surface smoothing they also find that the convergence rate is improved slightly. In their method they replace the inverse of the overlap matrix by a Taylor's series expansion of $(1+x)^{-1}$. Mauri showed results for a 256 atom carbon slab with a reconstructed C(111)-2x1 surface on each side. With a localization region up to 2nd neighbor, (presumably nearest-neighbor orthogonal tight-binding hamiltonian), random input, the system gets stuck in a local minimum with charge transferred between the surfaces. The claim is this is a bad feature in that one supposedly needs to know something about the solution. Actually I would argue that random starts are not a good idea., especially when knowing nothing more than very simple chemistry, reasonable starts are trivial. This is not to say that one will not get stuck in a local minima. However, it is easy to determine if this is the case. We show below that in a local minimum state the lowest unoccupied level falls below the highest occupied level. We also show below that the energies of these two levels are easily determined in O(N). One could then simply try a different starting point. This is not unlike the problem of trying to find the minimum energy geometric structure. However, for atomic positions, it is typically impossible to distinguish between a local and a global minimum, except by trial-and-error. Mauri and his coworkers see the local minima problem as a lemon. Nevertheless, that lemon is easily turned into lemonade. Mauri also discussed the slow convergence of the iterative methods for localized orbitals and made some conjectures but unfortunately could not offer real solutions.

W.M.C. Foulkes [36] discussed the Chemical Pseudopotential approach, an idea that dates back to the early 1970's [28]. The only person to do anything with the concept, to my knowledge was Bullett [27] and it never quite caught on. Foulkes points out that the problem with LCAO (linear combination of atomic orbitals) is that the basis is not fundamental, it is overcomplete, systematic improvement is difficult and on top of that one has to do lots of "boring" integrals. He suggests that useful criteria for a basis set might be (1) simple (2) fundamental, general and (3) capable of systematic improvement to completeness. A "clever" basis set is efficient, fundamental but maybe not simple. He then introduced what he called the "perfect" basis set that is (1) localized, (2) "perfect" or minimal, (3) not overcomplete. With such a basis one could achieve O(N) by using a grid to determine the "Wannier" (his "perfect") basis to form the tight-binding hamiltonian and then use one of the O(N) methods being developed for tight-binding.

D. Vanderbilt, also relatively early on, developed an O(N) method (Li, et al. [6]). The method is a density matrix variational approach that is closely related to one developed by Murray Daw [7] while he was at Sandia Livermore. Vanderbilt showed the generalization to non-orthogonal tight-binding and indicated that it is about 4x slower than for orthogonal tight-binding. For LCAO calculations he pointed out that with the overlap matrix typically being nearly singular this could be a problem. He also talked about the scaling of the convergence. He pointed out that his method has the advantage over Daw's that one can use a previous solution whereas Daw has to always start over from "infinite" temperature. However, he also pointed out that Daw's method seems to have the advantage that finite temperature can he handled naturally. Being able to work at finite temperature is considered valuable. With regards to the "local minimum" problem in orbital based methods, he pointed out that $\Omega(\psi)$ is a quartic function of the degrees of freedom whereas $\Omega(\rho)$ is cubic. A quartic function can have multiple minimum. He also pointed out that local minima are not necessarily bad (as I indicated above). Local minima could be indicative of metastable states. In comparison to Mauri, Galli and Car (MGC) [8] (1) both methods are variational (2) the number of variational degrees of freedom in Li, et al [6] is m^2NM versus mN_eM in MGC, where as above (Yang [2,18]) M is the number of atoms in the localization region (near-sightedness). $N_e < mN$ is the number of electron orbitals (3) Li, et al. [6] has no k-points where in MGC [8] each k-point is treated independently (this is not necessary). (4) There is a danger of run-away solutions in the density matrix approach but no danger in MGC [8] (5) Li, et al. [6] works at a given chemical potential whereas MGC [8] does not (the new version [15] does) and finally (6) no local minima in Li, et al.[6] and there is in MGC [8] (but not in the new method [15]).

K-M Ho reported on tight-binding Molecular dynamics using the variational density matrix approach that Vanderbilt discussed. The substantial development work of this Iowa group on this method should not be underestimated just because the basic method was introduced elsewhere. This group has extensive amount of results on nanovoids, grain boundaries, dislocations, impurities, polymers and bio-molecular systems. Ho showed that the cross-over from traditional methods varied substantially with type of system. (This was a point that would also be emphasized by Mauri, Galli and coworkers). For crystalline carbon the cross-over occurs at about 60 atoms (where only 2-3 iterations per time step are needed); for amorphous carbon at about 120 atoms and for liquid carbon at about 230 atoms. The biggest problem with the liquid is that because the system moves so quickly, information from the previous time step is of no value. He also emphasized that the slope for the linear scaling depends on the localization radius (which is typically about 5th neighbor) and also on the number of iterations (which is very system dependent). He also finds that working at constant chemical potential does not work well, but it is necessary to work at constant N (number of electrons). On a 512 node Intel Paragon on amorphous carbon (a problem of interest at Sandia) the calculations run about 200 steps/hr for 20 atoms/node. One additional, very interesting part of his presentation, was a linear-scaling solution to the problem of defining the electronic entropy. The electronic entropy formally involves the logarithm of the density matrix. The functional form has infinite slope at both $\rho \rightarrow 0$ and $\rho \rightarrow 1$. This is problematic. But he showed that a quartic function of the density matrix can reproduce the Fermi-Dirac function quite well. In fact, the minimum free energy reproduces the infinite slope at $\rho=0$ and $\rho=1$.

P. Ordejon discussed a linear scaling method to obtain phonons with 3N computations each scaling with N. Hence an N^2 method to compute the dynamical matrix. He indicated that this would be advantageous over a method using the vibrational power spectrum especially if one wanted to get the low frequency modes. To get O(N) one could restrict the relaxation to only a few orbitals; this would make each of the 3N calculations O(1). One also would truncate the dynamical matrix (this is typical) and finally one could use the Maximum Entropy method. He showed a calculation of a 216 atom cube of diamond and only the q=0 phonons. He used a 4Å cutoff and noted that the cutoff primarily affects the highest frequencies.

R. Martin then talked about how to put finite temperature into linear scaling; this was similar to what K-M. Ho discussed. One of the motivations behind finite temperature is that the density matrix localizes even in metals at finite temperature. The essence of linear scaling is in the locality of the density matrix; near sightedness. Metals become near-sighted only at finite temperature. He did point out that conductivity is inherently a non-local phenomenon and argued that maybe one needs to deal explicitly with some states near the Fermi level. This is a viewpoint also taken herein.

G. Galli discussed the strengths and weaknesses of the orbital based O(N) methods [8,15] within tight-binding. In the newer method [15] they include some empty orbitals (this is to be distinguished from partially occupied orbitals in Gillan's method [16] below). Hence one solves for $M > N_e/2$ orbitals. At solution $N_e/2$ are occupied and $M-N_e/2$ are empty. They explicitly showed that in 2D graphite with 31 atoms in the localization range, the number of iterations to converge is independent of size. They solved for three orbitals per atom and found that the number of iterations needed decreased as they increased the localization range. They do not get perfect energy conservation and finite electronic temperature does not make for a dramatic improvement. They systematically loose energy from the system with time.

M. Gillan [16] talked about practical considerations in taking linear scaling algorithms to LDA. Most of the work in linear scaling has dealt with tight-binding. It is not often recognized that this is not true of the work at Sandia [12] or of the work of Yang [2,18], both of which have used LCAO-LDA. Gillan's method is the most general of all the orbital and variational density matrix approaches. It is the generalization of the density matrix approach [6] to a variationally determined minimal basis. It is also the generalization of Hierse/Stechel [12] to partially occupied orbitals. Like Li, et al.[6], Gillan uses a cubic purification of the density matrix. Gillan centers his localization regions on the atoms and he uses a real-space finite difference method.

O. Sankey talked about getting the total energy for a large system by using what he introduced as the "impartial" vector which has equal amplitude (up to a phase) on each eigenvector of the system. A random vector satisfies this statistically, hence a suitable average will work. Since only the occupied spectrum is of interest he biases his random vectors to fit exactly the first few moments. He uses maximum entropy with some number of moments, L, to get the full spectrum. The test system that he used was 64 atoms of GaAs in tight-binding. He was able to achieve 1meV accuracy with 200-300 random vectors and L=40-50 moments. However, he indicated that getting forces with such a method is problematic.

W. Kohn (one of the fathers of density functional theory) began his presentation with several disclaimers (1) no physics (2) no numerical results - yet and (3) no claims for practical usefulness. Despite these disclaimers his presentation was, as anyone might expect, quite thought provoking. The goal and the challenge is to calculate a density matrix that is hermitian, normalized (for constant N, but could be constant chemical potential) and idempotent (at T=0). If not for the criterion of idempotency, the problem would be easy. The orbital- density matrix-based methods already talked about (and yet to be talked about) use some form of "purification" and functional minimization to achieve idempotency. In contrast Kohn defined a "Penalized" Energy Functional which adds a term P[n] that vanishes for an idempotent density matrix, n, and increases as the matrix deviates from idempotency. He then proceeded to show that the variational functional minimum, with respect to n, of the Functional $Q[n,\alpha] = E[n] + \alpha P[n]$ is equal to E_0 (the energy minimum of E[n]) for $\alpha > \alpha_c$. α must be real and positive, n must be non-negative and it may not matter if the spectrum is bounded. He proceeded to describe the proof.

In my own presentation, my focus was on the full scaling of the problem not just the scaling with the size. I defined complexity scaling as an algorithm for which the computational effort scales with the "useful, unknown information" in the system. One point that has come up repeatedly is the "near-sightedness" of electronic structure. This, of course, is what makes linear scaling even a possiblility. However, I suggest that there are really two length scales in the problem. One is the range of the density matrix (this is closely related to the localization range in several of the above talked about methods) and the other is how localized can an occupied orbital be. If one demands that the orbitals be orthogonal then these two ranges are one and the same. For reasonable accuracy this localization range can easily incorporate 100 atoms. Using less atoms begins to seriously degrade the accuracy. By going to non-orthogonal orbitals one can solve for very localized (one or two neighbors) orbitals and a relatively long range (still independent of N) without giving up any accuracy and without a large cost in efficiency. This method easily achieves constant N_e which can be a problem in other methods. Another viewpoint taken was that possibly more important than the linear scaling is the possibility of attaching physical significance to the orbitals. Since they are very localized it seems they should be able to be transferred between chemically similar systems. In fact in self-consistent LDA calculations with Hierse using Sandia's parallelized LCAO- QUEST code, localized orbitals were computed in heptane (a hydrocarbon with 7 carbons). In full diagonalization the variational degrees of freedom scale with the square of the number of atoms. With localized orbitals, the number of variational degrees of freedom scale only linearly with the number of atoms. As a consequence, we were able to reduce the number of degrees of freedom by an order of magnitude with no sacrifice in accuracy. For a longer hydrocarbon (decane, 10 carbon atoms) the variational degrees of freedom were reduced to NONE by using the results from heptane. Again there was no sacrifice in accuracy. I went on to talk some about our original algorithm which to date still is the only one without the slow convergence problem. In conclusion I emphasized that the density matrix is only quasi-localized but the occupied orbitals can be truly localized if one goes to non-orthogonal orbitals. This is not 6 of 1, 1/2 dozen of the other since it is the total number of degrees of freedom that matter. This becomes more important in LDA calculations than in tight-binding calculations and especially if one is solving on a grid. I re-emphasized that more attention needs to be given to the problem of slow convergence and the solution probably lies with a good pre-conditioner. For the lowest overall complexity algorithm I placed my bet on non-orthogonal localized orbitals on a real space, variable spaced mesh with either the state-by-state updates or with a pre-conditioned global update of all the orbitals. One final point I tried to make was that there is no obvious answer to balance of complexities especially since it will be system dependent..

P. Ordejon presented his LCAO approach to self-consistent LDA calculations with linear scaling. It is based on Sankey's minimal basis set approach, but replaces the Harris functional with a FFT Poisson solver and true self-consistency. The matrix elements are calculated in O(N) and account for 90% of the work. The Poisson solver is $O(N_G \ln N_G)$ and accounts for 10% of the work. The "virtues" of this approach include (1) the tight-binding language is retained, the hamiltonian and overlap matrices are sparse and the extra work involves the determination of the Hartree potential and matrix elements of the Hartree potential and the exchange-correlation potential. (2) the matrices are small (minimal basis) (3) except for the FFT all operations on the grid are local and (4) the kinetic energy and non-local pseudo-potentials are handled exactly. The problems have to do with convergence. An MD calculation on 64 atoms with a time-step of 2 fs. With a 5 Å cutoff for the orbitals, the number of CG (conjugate gradient) iterations was 30 and the number of self-consistent (SCF) cycles was 6. In contrast with a 20 Å cutoff only 6 CG iterations and 3 SCF cycles were needed. As a measure of accuracy the O(N) 5 Å (20 Å) cutoff got a bulk modulus of 94 (96) GPa. A PW-PP calculation got 98 GPa. Experiment is 99 Gpa and the Harris-Functional Sankey approach gets 91 GPa.

D. Vanderbilt also presented a short discussion paper on the relationships between several of the variational linear-scaling electronic structure algorithms. The ones under discussion all express a trial density matrix as $\sigma = L_{ij} |\varphi_i\rangle\langle\varphi_j|$ where $\{i,j=1,....,M)$. The density matrix is expressed as a polynomial in σ. The polynomial acts as a "purification" which constrains ρ to have eigenvalues between 0 and 1. The minimization then drives them to be 0 and 1, thereby achieving idempotency. In Li, et al (LNV) [6] the L_{ij} are varied and the $\{\varphi_i\}$ are fixed. The purification is cubic $\rho= 3\sigma^2-2\sigma^3$. M is greater than the number of occupied orbitals; it is equal to the number of basis orbitals in the problem. In Gillan's [16] method the L_{ij} are varied and the $\{\varphi_i\}$ are also varied. The purification is again cubic $\rho= 3\sigma^2-2\sigma^3$. M is greater than the number of occupied orbitals but less than the number of basis orbitals in the problem. In Kim, et al. [15] the L_{ij} are fixed and equal to δ_{ij} but the $\{\varphi_i\}$ are varied. The purification is quadradic $\rho= 2\sigma-\sigma^2$. The quadradic is possible because σ is positive definite which assures that $\sigma=0$ is a minimum without going to the cubic. Like in Gillan's [16] method M is greater than the number of occupied orbitals but less than the number of basis orbitals in the problem. The NOLO [11,12] method doesn't fix L_{ij} but it isn't freely varied. It is effectively a dependent quantity. The $\{\varphi_i\}$ are varied. The purification is again quadradic $\rho= 2\sigma-\sigma^2$ and M is equal to the number of occupied orbitals. Vanderbilt talked then about the relationship between the purification and the approaches to finite temperature by looking at the inverse of the purification operation.

C. Qualitative Discussion: Localized versus Delocalized Solutions

Stationary states will transform as irreducible representations of the total symmetry group. In contrast, in real space methods, such as localized orbitals the solutions transform as reducible representations in which a symmetry operation of the group will transform one real-space orbital into another. As an example consider the methane molecule, CH_4. Carbon has six electrons and each hydrogen has one. Thus methane has 10 electrons and five occupied states. Each state can be doubly occupied, with a spin up and a spin down electron. One of the occupied levels is the 1s core level of carbon. For the purposes here it is sufficient to describe the electronic structure of this molecule as the occupied core atomic state and four occupied valence states, described as linear combinations of 2s and 2p orbitals on carbon and 1s orbitals on each hydrogen[2]. The molecule has tetrahedral symmetry and the four stationary states transform as an A_1 and a triply degenerate T. We can write these functions schematically as follows:

$$|A_1\rangle = \alpha_A|C_{2s}\rangle + \beta_A[|H_1\rangle + |H_2\rangle + |H_3\rangle + |H_4\rangle]$$

$$|T_x\rangle = \alpha_T|C_{2p_x}\rangle + \beta_T[|H_1\rangle + |H_2\rangle - |H_3\rangle - |H_4\rangle]$$

$$|T_y\rangle = \alpha_T|C_{2p_y}\rangle + \beta_T[|H_1\rangle - |H_2\rangle + |H_3\rangle - |H_4\rangle]$$

$$|T_z\rangle = \alpha_T|C_{2p_z}\rangle + \beta_T[|H_1\rangle - |H_2\rangle - |H_3\rangle + |H_4\rangle]$$

where the positions of C, H_1, H_2, H_3 and H_4 are at (0,0,0), (1,1,1), (1,-1,-1), (-1,1,-1) and (-1,-1,1), respectively. These positions are in units of $r_0/\sqrt{3}$ where r_0 is the CH bond length. Two normalization conditions determine two of the parameters $\{\alpha_A, \beta_A, \alpha_T, \beta_T\}$ and energy minimization determines the other two.

Chemically, however, we think of this molecule as four equivalent C-H bonds. No one bond has the symmetry of the molecule; the symmetry operations of the tetrahedral point group transform one bond into another. But where is the bond orbital in the above denumerated stationary

---

[2] These atomic orbitals are not necessarily identical to the free atom but in the most accurate representation would be chemically modified by the presence of other atoms. An even more accurate treatment would include p orbitals on hydrogen and d orbitals on carbon.

states? The bond orbitals are defined to be that linear combination of the four occupied stationary states (and therefore is not, itself, a stationary state) that most localizes along one bond. For this system we can completely eliminate any weight on the three hydrogens not participating in the bond. That is we can write four equivalent localized occupied orbitals as follows:

$$|\phi_1\rangle = \alpha_L|C_{2s}\rangle + \beta_L|C_{2p_{111}}\rangle + \gamma_L|H_1\rangle$$

$$|\phi_2\rangle = \alpha_L|C_{2s}\rangle + \beta_L|C_{2p_{1-1-1}}\rangle + \gamma_L|H_2\rangle$$

$$|\phi_3\rangle = \alpha_L|C_{2s}\rangle + \beta_L|C_{2p_{-11-1}}\rangle + \gamma_L|H_3\rangle$$

$$|\phi_4\rangle = \alpha_L|C_{2s}\rangle + \beta_L|C_{2p_{-1-11}}\rangle + \gamma_L|H_4\rangle$$

The subscript on the p orbitals indicates the direction. (It points directly at the hydrogen atom in the corresponding bond). These orbitals are determined by a set of three parameters $\{\alpha_L, \beta_L, \gamma_L\}$ one of which is arbitrarily determined by normalization, the others by minimizing the energy. We should point out that these orbitals are necessarily *non*-orthogonal; $S_{ij} = \langle\phi_i|\phi_j\rangle \neq \delta_{ij}$ and nonstationary; $H_{ij} = \langle\phi_i|\hat{H}|\phi_j\rangle \neq \varepsilon_i\delta_{ij}$. However, because the set of four bond orbitals are linearly independent, the four "bond" orbitals span the same Hilbert space as the occupied stationary states. The bandstructure minimum energy is $Tr\{\mathbf{S}^{-1}\mathbf{H}\}$. The relationship between the two sets of parameters is:

$$\frac{\gamma_L}{4\alpha_L} = \frac{\beta_A}{\alpha_A} \qquad \frac{\sqrt{3}\gamma_L}{4\beta_L} = \frac{\beta_T}{\alpha_T} \qquad \frac{\beta_L}{\alpha_L} = \sqrt{3}\frac{\alpha_T\beta_A}{\alpha_A\beta_T}$$

Consequently, the stationary state representation or the bond orbital representation are equivalent, are equally valid and carry the same information content. For this simple system, both sets are described by the same number of independent parameters. For this molecule, being as small as it is, there is no obvious advantage to either solution. We presented it here to illustrate the difference in the types of solution.

We now consider the hydrocarbon molecule, decane. This is a case in which the "bond" orbitals have an obvious advantage over the stationary states. The "bond" orbitals are localized over a relatively small fraction of the total system. For example, the total information content in decane, $C_{10}H_{22}$ with 62 valence electrons, 62 atomic orbitals in the basis and 31 occupied states requires approximately 930 parameters to describe the electronic system in terms of delocalized stationary states. This is after taking account of symmetry and normalization. In contrast the bond orbital description has three *chemically* inequivalent CC bonds and four *chemically* inequivalent CH bonds. Each bond orbital is described by 8-15 parameters depending on the orbital. The total parameter set thus has less than 105 parameters. More importantly, this same set of parameters can describe C_nH_{2n+2} for all $n \geq 7$. Quantitatively, the same set of "bond" orbitals can describe heptane and decane with a precision error of 10^{-2} meV and an accuracy (within LDA and the chosen basis) of 2 meV. These calculations are described below. The intrinsic accuracy limit is determined by the chosen range for the "bond" orbitals (nearest neighbor to the bond being formed). Thus, when chemically invariant subunits exist on length scales smaller than the system size (or the unit cell for periodic systems), a formulation of the electronic structure in real space localized terms can be very advantageous. Naturally, this approach relates well to chemical thinking. The benefits of a real space approach to condensed matter systems has been emphasized in the works of Haydock, Heine and Kelly beginning in the 1970's [30]. The emphasis in those works, however, are on obtaining local density of states of atomic valence orbitals within non-self consistent tight-binding representations. In contrast, the emphasis here is on chemically invariant subunits that are smaller than the size of the system and on solving for bond-orbitals that depend only on the local chemical environment not on the global system and to do this within an accurate, self-consistent formulation. The *two* advantages in changing how we represent the electronic structure of a system from a delocalized to a localized representation are: (1) by determining localized orbitals directly (i.e., not first determining delocalized stationary functions[3]) we compute the electronic structure with an algorithm that scales linearly with the size of

---

[3] Localized orbitals in the framework of HF theory have been obtained, typically as a unitary transformation of the already computed delocalized eigenfunctions (or molecular orbitals). See for example, Refs. [37,38]. The method of Edmiston and Rudenberg [37] is the method most often used.

the system (see Section 4.2 and 4.3 below) and (2) by determining localized orbitals we can get very accurate first guesses for the electronic structure of large systems by breaking it into chemically invariant subsystems for which a smaller calculation can be performed. As the above example illustrated, the electronic structure of any C_nH_{2n+2} (n>7) could be guessed by doing one calculation on a smaller (C_7H_{16}) system. In essence this latter point is sublinear scaling.

Cast in this form localized orbital theories can be thought of as rigorously bridging the gap between the intuitive chemical "building block" picture of a molecule, surface or solid (well established chemical concepts) and the delocalized nature of the quantum mechanical solution of the Schroedinger-like equation in the independent-electron approximation. While stationary states are sensitive to perturbations anywhere in the system, localized orbitals, in contrast, are sensitive to perturbations only within its range of localization and insensitive to perturbations far away. In localized orbital formulations, we seek a solution in which analogous chemical environments have similar bond orbitals (e.g., if two different molecules contain parts that for chemical or physical reasons are considered similar in character, then the quantities defined should also look similar). Hydrocarbon chains of different lengths should have terminal carbon-carbon bonds that look alike and internal carbon-carbon bonds that also look alike. These type of results are exactly what we look for in a localized orbital formulation. They are invariant in going between chemical environments that are similar even if not rigorously related by symmetry. In essence, having solved for the electronic structure of heptane, we can accurately "guess" the electronic structure of decane as is consistent with chemical intuition but inconsistent with traditional delocalized molecular orbital solutions.

While not used extensively in large scale computations the concept of the Wannier [32-34] and generalized Wannier orbitals are fairly well-known. These are linear combinations of different k-stationary states or Bloch waves (in a translationally invariant system) that are localized in real space. They have the property that a Wannier function localized in one unit cell is orthogonal to its translational replica in another unit cell. For our purposes we define generalized Wannier functions as real space, localized, orthogonal, occupied (i.e., linear combinations of the occupied stationary states of a system) functions. Thus they are related to the occupied stationary state representation by a unitary transformation.

The asymptotic behavior of the Wannier orbitals at large distances, r is determined by $\exp(-\kappa_b|r|)$ where κ_b is a branch point in complex k space, and is, in magnitude, of the order of the smallest gap involving the given band [32,33]. Because the range is related to the gap it is clearly physically significant. However, in real systems this range can be rather large (semi-global) precluding chemical invariance for all practical purposes. As pointed out by Anderson [26], it is trivial to show that the Wannier functions are not the most localized linear combination of the Bloch waves. However, since the range is physical it cannot be removed by making a qualitative change. It is our contention that this semi-global behavior should not be ignored. It is, however, a property of fermions and Pauli exclusion and thus does not necessarily enter the definition of chemically invariant localized orbitals that are non-orthogonal. All the relatively long-range information, that is sensitive to fairly distant perturbations in the system can be embodied in the density matrix of the bond-orbital representation. In an orthogonal occupied representation the density matrix is unity. Thus all physical content must be contained in the orbitals. For non-orthogonal occupied representations, however, the density matrix **D** is formally the inverse of the overlap matrix and thus can carry relatively long range information that is sensitive to semi-global aspects of the system while the individual bond orbitals are insensitive. At this point, this is largely conjecture. Nevertheless, the significance of this point seems to have been missed in the early literature.

D. DENSITY FUNCTIONAL THEORY

In density functional theory (DFT) the basic independent element is the density operator, $\hat{\rho}$ which itself is a trace over the outer product of occupied orbitals $\{\psi_k\}$;

$$\hat{\rho} = \sum_k |\psi_k\rangle\langle\psi_k|. \tag{1}$$

The diagonal element of the density matrix in a spatial representation is the charge density, $\rho(r) = \langle r|\hat{\rho}|r\rangle$.

For a paramagnetic ground state with 2N electrons, the occupied orbitals are the N lowest energy eigenfunctions of the one-electron, self-consistent Kohn-Sham equation:

$$H_{KS}\left|\psi_k\right\rangle = \left(-\frac{\hbar^2}{m_e}\nabla^2 + V_{DF} + V_{NL}\right)\left|\psi_k\right\rangle = \varepsilon_k\left|\psi_k\right\rangle \tag{2}$$

where V_{DF} includes the local potential and the self-consistent Coulomb and exchange-correlation potentials. V_{NL}, if used, is the non-local (but short-range) pseudo-potential used to eliminate the core electrons from the calculation. The solution of this Schroedinger-like equation, whether it is achieved by direct diagonalization or by an iterative method, is the source of the "N^3" scaling. In the iterative methods this scaling arises from orthogonality constraints [23].

The density operator is, however, a projection operator onto the occupied subspace of H_{KS}. No physical quantities in the density functional formalism depend on individual occupied orbitals but only on a trace over the occupied subspace. The trace operation is invariant to all similarity transformations. Hence, any set of $\left|\phi_j\right\rangle$ such that

$$\left|\phi_j\right\rangle = \sum_k T_{jk}|\psi_k\rangle \tag{3}$$

can equally well represent the occupied subspace, provided the inverse of $\underline{\underline{T}}$ exists. We exploit this freedom to overcome the N^3 computational bottleneck. The $\left|\phi_j\right\rangle$, in general are non-orthogonal, with an overlap matrix $\underline{\underline{S}} = \underline{\underline{T}}\,\underline{\underline{T}}^\dagger$ Thus only a unitary $\underline{\underline{T}}$ implies an orthogonal set of $\left|\phi_j\right\rangle$. Orthogonal or not, the density operator can always be written as

$$\hat{\rho} = \sum_j |\overline{\phi}_j><\phi_j| \tag{4}$$

where the $\left|\overline{\phi}_j\right\rangle$ are the bi-orthogonal complement of the $\left|\phi_j\right\rangle$ as explained further below.

The basic *non-linear* equation that each $\left|\phi_j\right\rangle$ must satisfy is

$$\left(1-\hat{\rho}\right)\hat{H}\left|\phi_j\right\rangle = 0. \tag{5}$$

To cast Eq. 5 in a projection operator language, we refer to the occupied and unoccupied subspaces as P and Q, respectively where Q=1-P. The basic equation for the $\left|\phi_j\right\rangle$ is then QHP = 0. In other words, it is only necessary for the hamiltonian to be *block* diagonalized into two blocks (occupied and unoccupied) subject to the criterion that the band structure energy $Tr\left\{P\hat{H}\right\}$ is an absolute minimum (for a fixed one-electron potential).

$$Tr\left\{P\hat{H}\right\} = \sum_j \left\langle\overline{\phi}_j\middle|\hat{H}\middle|\phi_j\right\rangle = \sum_k \varepsilon_k. \tag{6}$$

This essentially defines the problem. It is both necessary and sufficient that Eq. 5 is satisfied and Eq. 6 is an absolute minimum for a fixed potential. This follows because of the invariance of the trace. The only way to minimize the sum of the one electron energies is to be a representation of the lowest N eigenfunctions of Eq. 2. We emphasize, however, that Eq. 5 while minimizizing Eq. 6 does not have a unique solution. Indeed our aim is to use the freedom corresponding to this non-uniqueness to try to find a set of $\left|\phi_j\right\rangle$ that not only solves Eq. 5 but which is also optimally "localized", or equivalently, that makes the matrix representations of **H** and **S** optimally sparse. In the work reported here we find a solution which is localized but is not necessarily optimally localized.

E. CONVENTIONS AND NOMENCLATURE

Here we set forth definitions and conventions that are convenient when working with non-orthogonal basis sets. (cf. Ref 22 & 39 for more information) Consider a finite Hilbert space, $\mathcal{H}$, of dimension N, spanned by a linearly independent set $\left\{\phi_k\right\}$, k=1,2,..N. The inner product of the functions in this set defines a matrix **S**

$$S_{jk} = \left\langle \phi_j \middle| \phi_k \right\rangle \tag{7a}$$

which is the *metric* for the space. Without loss of generality, assume that the $\left\{\phi_k\right\}$ are normalized to unity. The metric is a positive definite, hermitian matrix. This means that all its eigenvalues are greater than zero and that $\mathbf{S}_{jk} = \mathbf{S}^*_{kj}$. There also exists a uniquely determined set, $\left\{\overline{\phi}_j\right\}$, j=1,2,...N called the "bi-orthogonal complement" which satisfies the bi-orthogonality relationship,

$$\left\langle \overline{\phi}_j \middle| \phi_k \right\rangle = \delta_{jk}. \tag{7b}$$

The two sets are related by the inverse of the metric; i.e.,

$$\left| \overline{\phi}_j \right\rangle = \sum_k \left| \phi_k \right\rangle \left[S^{-1} \right]_{kj}. \tag{8}$$

The inverse of the metric is the metric in the complement space. We define

$$\left[S^{-1}\right]_{kj} \equiv D_{kj} = \left\langle \bar{\phi}_k \middle| \bar{\phi}_j \right\rangle. \tag{9}$$

For an efficient algorithm, it is convenient to approximate the matrix **D** and thus in general we will not know the inverse of **S**. When the **D** we use is identically $\mathbf{S}^{-1}$ we will refer to it as "full" inversion, although we note that this matrix cannot generally be computed in order N.

On $\mathcal{H}$ the unit operator is defined as:

$$1_{\mathcal{H}} = \sum_k \left|\bar{\phi}_k\right\rangle\left\langle \phi_k\right|. \tag{10}$$

One way to see this is to substitute Eq. 9 into Eq. 8. When operating on a general function, $1_{\mathcal{H}}$ projects that function onto $\mathcal{H}$.

We will refer to the condition number of the metric. It can be defined in several ways, for example as the ratio of the largest to the smallest eigenvalue of **S**. We find it more useful to define it as the largest diagonal element of **D**. The better the condition of the metric the closer (from above) the diagonal elements will be to one. (Without loss of generality, we assume that the diagonal elements of **S** are unity.) Another definition is the norm of the metric or of the complement metric where the norm of a matrix, **A,** of dimension N is defined as

$$\mathrm{Norm}(\underline{\underline{\mathbf{A}}}) = \frac{\sqrt{\mathrm{Tr}\left[\underline{\underline{\mathbf{AA}}}\right]}}{\mathrm{N}}. \tag{11}$$

The smaller the norm, the better the condition.

In addition to defining the metrics in $\mathcal{H}$, it is also necessary to define matrix representations of the hamiltonian. We adhere to the convention that the hermitian matrix of an operator in the basis $\left\{\phi_k\right\}$ is the "standard" representation, denoted with a tilde, for example,

$$\tilde{\mathbf{H}}_{jk} = \left\langle \phi_j \middle| \hat{\mathrm{H}} \middle| \phi_k \right\rangle. \tag{12a}$$

A "proper" [39] matrix is one represented in the dual basis as follows:

$$\mathbf{H}_{jk} = \left\langle \bar{\phi}_j \middle| \hat{\mathrm{H}} \middle| \phi_k \right\rangle. \tag{12b}$$

The relationships between the two representations are

$$\underline{\underline{\mathbf{H}}} = \underline{\underline{\mathbf{D}}}\,\underline{\underline{\tilde{\mathbf{H}}}} \quad \text{and} \quad \underline{\underline{\tilde{\mathbf{H}}}} = \underline{\underline{\mathbf{S}}}\,\underline{\underline{\mathbf{H}}}. \qquad (13).$$

As mentioned, we choose representations for these matrices that render them sparse. For sparse matrices we refer to an X-range, where X can be any one of the five relevant matrices, $\mathbf{S}, \tilde{\mathbf{H}}, \mathbf{D}, \mathbf{H}$ or $\mathbf{H}^{\dagger}$. The set of elements in the range of k is the non-zero matrix elements in the kth column. Each of the five matrices will, in general, have different, but comparable, ranges.

III. TECHNICAL APPROACH

A. ESSENTIAL INGREDIENTS OF THE ALGORITHM

We describe here an iterative solution for a set of "occupied" orbitals $\left|\phi_k^{(n)}\right\rangle$. As the iterations (n) proceed, the quantity $\left[1-P^{(n)}\right]\hat{H}\left|\phi_k^{(n)}\right\rangle$ approaches zero for k=1,..,N. Thus at convergence, $n \rightarrow \infty$, $Q^{(n)}\,\hat{H}\,P^{(n)} \rightarrow 0$ where $Q^{(n)} = \left[1-P^{(n)}\right]$. $P^{(n)}$ is the projection operator onto the subspace spanned by the $\left|\phi_k^{(n)}\right\rangle$ and converges to the density operator. In each "ϕ-update" exactly one $\left|\phi_k\right\rangle$ changes [16]. However, all $\left|\overline{\phi}_j\right\rangle$ in a D-range of k also change. (Recall that "D-range" of k means all j such that $\mathbf{D}_{jk} \neq 0$.) With a full inversion or a poorly conditioned metric, the D-range of k would be essentially all j. However, in a later section we show that it is sufficient for the D-range of k to be comparable to that defined by the metric (the S-range), provided the metric is well-conditioned. We return to this below. With "ϕ-update" defined this way, we have a non-standard iterative scheme in which the number of ϕ-updates is linear in N, but the work per each ϕ-update is N-independent. Reference below to the number of "iterations" means to the number of ϕ-updates divided by N.

Our algorithm is a sequence of ϕ-updates. First we describe the "in principle" algorithm, by which we mean no approximations. We try to indicate where the in principle will differ from the in practice by explicitly including the phrase "in principle". An orbital is updated by optimally mixing in the *residual,* $\left|\chi_k^{(n)}\right\rangle$ which is defined as

$$\mathrm{Q}^{(n-1)}\hat{\mathrm{H}}\left|\phi_k^{(n-1)}\right\rangle = \eta_k\left|\chi_k^{(n)}\right\rangle. \tag{14}$$

The mixing is a "generalized Jacobi rotation" (where "generalized" will be defined momentarily); i.e.,

$$\left|\phi_k^{(n)}(\theta)\right\rangle = \cos(\theta)\left|\phi_k^{(n-1)}\right\rangle - \sin(\theta)\left|\chi_k^{(n)}\right\rangle. \tag{15}$$

The η_k is chosen so that $\left|\chi_k^{(n)}\right\rangle$ is normed to unity. At convergence the η_k vanish, and the $\left|\chi_k^{(n)}\right\rangle$ are undefined. Prior to convergence, the residual is orthogonal (at least in principle) to *all* the $\left|\phi_j^{(n-1)}\right\rangle$; i.e., $\left\langle\chi_k^{(n)}\middle|\phi_j^{(n-1)}\right\rangle = 0$, for j = 1,2,..,N. In other words, it is an element in $Q^{(n-1)}$ space. Furthermore, it can be shown that any other direction in $Q^{(n-1)}$, that is orthogonal to the residual, has no hamiltonian matrix element with the orbital being updated. The mixing angle θ is chosen to minimize the total band structure energy $E^{(n)}(\theta)$ [4] where

$$E^{(n)}(\theta) = \mathbf{H}_{kk}^{(n)}(\theta) + \sum_{j\neq k} \left\langle\bar{\phi}_j^{(n)}(\theta)\middle|\hat{\mathbf{H}}\middle|\phi_j^{(n-1)}\right\rangle. \tag{16a}$$

We have left the (n-1) superscript label on $\left|\phi_j\right\rangle$ (j≠k) to emphasize that it does not change during the update. We define the kth state contribution to the band structure energy, as the real part of the diagonal element in the proper representation, $\mathbf{H}_{kk}^{(n)}(\theta)$,

---

[4] Generalizing the energy minimization to that of the total energy rather than the total band-structure energy; i.e., eliminating the double counting of the coulomb energy and including the exchange-correlation energy can be accomplished in a similar manner to that described in Ref. 23.

$$\mathbf{H}_{kk}^{(n)}(\theta)=\left\langle \overline{\phi}_k^{(n)}(\theta)\right|\hat{\mathbf{H}}\left|\phi_k^{(n)}(\theta)\right\rangle \tag{16b}$$

We make this definition to point out the analogy to the band-by-band conjugate-gradient iterative schemes [23] wherein a band refers to a specific type of occupied orbital, i.e., an eigenstate. Here a state refers to a general "occupied" orbital.

At first glance, approximations to minimize $E^{(n)}(\theta)$ may appear necessary, because with an update to $\left|\phi_k^{(n-1)}\right\rangle$, all $\left\langle\overline{\phi}_j^{(n-1)}\right|$ in the D-range of k must also be updated. Nevertheless, algebra shows that no approximations are necessary. Appendix B gives the details for determining θ_k, the optimal mixing angle. It involves evaluating a 2x2 standard hamiltonian matrix with the functions $\overline{\phi}_k^{(n-1)}$ (note this is the complement function) and $\chi_k^{(n)}$. However, θ_k is *not* the rotation angle for diagonalizing this 2x2 matrix, nor is it the rotation angle for diagonalizing the standard hamiltonian matrix for the two directions $\left(\phi_k^{(n-1)}\right.$ and $\left.\chi_k^{(n)}\right)$ that are being rotated (Eq. 15). Because θ_k is not the diagonalizing rotation angle, we refer to Eq. 15 as a "generalized Jacobi rotation". Nevertheless, despite the change in many of the complement functions - and hence many individual state energies - the change in the *total* band structure energy depends only on matrix elements involving the residual and the current *complement* function.

After determining θ_k, the kth row and column of $\underline{\underline{\mathbf{S}}}$ and $\underline{\underline{\tilde{\mathbf{H}}}}$ are trivially updated. Updating the standard representation of the hamiltonian, $\underline{\underline{\tilde{\mathbf{H}}}}$, involves determining the inner product of $\hat{\mathrm{H}}\left|\chi_k^{(n)}\right\rangle$ with the $\left|\phi_j^{(n-1)}\right\rangle$ for all j in the $\tilde{\mathbf{H}}$-range of k, as follows:

$$\tilde{\mathbf{H}}_{kk}^{(n)}=\tilde{\mathbf{H}}_{kk}^{(n-1)}-\sin(\theta_k)\left[\cos(\theta_k)h_k^{(n)}-\sin(\theta_k)\left(h_{\chi\chi}^{(n)}-\tilde{\mathbf{H}}_{kk}^{(n-1)}\right)\right] \tag{17a}$$

$$\tilde{\mathbf{H}}_{kj}^{(n)} = \left[\tilde{\mathbf{H}}_{jk}^{(n)}\right]^{*} = \left[\cos(\theta_k)\tilde{\mathbf{H}}_{kj}^{(n-1)} - \sin(\theta_k) h_j^{(n)}\right] \quad for\ j \neq k \tag{17b}$$

$$\tilde{\mathbf{H}}_{jl}^{(n)} = \tilde{\mathbf{H}}_{jl}^{(n-1)} \quad for\ j,l \neq k \tag{17c}$$

where

$$h_j^{(n)} = \left\langle \chi_k^{(n)} \middle| \hat{\mathrm{H}} \middle| \phi_j^{(n-1)} \right\rangle \qquad h_{\chi\chi}^{(n)} = \left\langle \chi_k^{(n)} \middle| \hat{\mathrm{H}} \middle| \chi_k^{(n)} \right\rangle. \tag{17d}$$

In principle the metric would update as follows: without loss of generality, we move the kth function to the top of the basis, and define the remaining (N-1)x(N-1) states as q-space, then

$$\mathbf{S}^{(\mathbf{n-1})} = \begin{bmatrix} \mathbf{1} & \left[s_q^{(n-1)}\right]^{\dagger} \\ s_q^{(n-1)} & S_{qq}^{(n-1)} \end{bmatrix} \tag{18a}$$

where $s_q^{(n-1)}$ is a vector of length N-1 and $S_{qq}^{(n-1)}$ is a matrix of dimension (N-1) x (N-1). Because the residual is orthogonal to all the "occupied" orbitals, the update to the metric is

$$\mathbf{S}^{(\mathbf{n})} = \begin{bmatrix} 1 & \cos(\theta_k)\left[s_q^{(n-1)}\right]^{\dagger} \\ \cos(\theta_k) s_q^{(n-1)} & S_{qq}^{(n-1)} \end{bmatrix} \tag{18b}$$

The important implication of this result is that matrix elements of $\mathbf{S}^{(\mathbf{n-1})}$, which are zero, *stay* zero, and any non-zero off-diagonal elements can only *decrease* in magnitude, because

$|\cos(\theta)| \leq$. Hence, the norm of **S** decreases, which implies that its condition improves. Furthermore, the metric is updated trivially and involves only m multiplications, where m is the number of non-zero elements in the kth row of S. In addition, due to this form as the update of **S**, the update of **D** is also relatively simple. More important, the complement basis can be updated directly. Consequently, the proper representation of the hamiltonian can be updated without multiplying **D** and $\tilde{\mathbf{H}}$ together, which would be order m^3. The key step in making our algorithm scale linearly with a relatively small pre-factor is the non-obvious, non-trivial update of the set of $\left\{\bar{\phi}_j^{(n)}\right\}$ in the D-range of k and the direct update of the order m^2 elements (where m is independent of N) of **D** and **H** that change. We show this in Appendix A. The total number of operations to accomplish this update is order m^2.

We emphasize that (1) the "inversion" of the metric is *not* necessary at every ϕ-update - in principle it is necessary only at the first step; and (2) the norm of **D** and **S** decrease; therefore the condition of the metric improves with each ϕ-update. Actually it is not very significant that the condition improves, but it *is* meaningful that the condition cannot get worse, thus assuring stability of the algorithm. Note that, while the metric *approaches* unity (orthogonality), it does *not*, in general, become unity. Hence, our algorithm does not iterate to orthogonal or generalized Wannier orbitals as in the methods of Mauri, et al [8] or Ordejón, et al [9]. Additionally we emphasize that the algebra of the updates relies only on the residual vector being orthogonal to the current set of "occupied" orbitals.

After updating $\left|\phi_k^{(n)}\right\rangle$ and the relevant matrices, a new k is chosen. It would be most natural to choose the orbital with the largest residual (largest η_k). However, since many η_k change with each update, to literally do this would take order N work. Instead, one approach is: at the beginning of each iteration a set of η_k's are calculated to fix the order. In other words, for the purpose of ordering the updates the change in a η_j with an update to a $\left|\phi_k^{(n)}\right\rangle$ is ignored. Other orderings are possible including highest or lowest state energies, $\mathrm{Re}\left\{\mathbf{H}_{kk}^{(n-1)}\right\}$, or random.

In practice, after each iteration (or few iterations), we recompute all the matrices given the $\left|\phi_j^{(n)}\right\rangle$ to eliminate cumulative errors due to (1) round-off, (2) **D** not identically $\mathbf{S}^{-1}$, and (3) $\left|\chi_k^{(n)}\right\rangle$ restricted to be more spatially localized than dictated by Eq. 14. We should note that this localization is imposed with cutoff functions as opposed to using bucket potentials [4]. We are investigating other methods of assuring localization with the goal to approach optimal localization by a quantitative measure; i.e., to minimize the ranges of the relevant matrices and maximize transferability.

One final point is: as noted, **D** as computed is not identically $\mathbf{S}^{-1}$; consequently Tr**H** (where Tr is the sum of the diagonal elements) is not necessarily a variational upper bound to the band structure energy. In principle, this is a problem. However, $\mathrm{Tr}\{(2-\mathrm{DS})\mathrm{H}\}$ can be shown to be an upper bound, provided the zero of energy is taken so that **H** is negative definite. With this formulation, the effective density matrix is $2\mathrm{D}-\mathrm{DSD}$; thus, the particle number is $\mathrm{Tr}\{(2-\mathrm{DS})\mathrm{DS}\}$, which can be shown to be less than or equal to N. The transformation to the effective density matrix guarantees that its eigenvalues are all nearly - but less than or equal to - one. Equality arises if **D** is identically $\mathbf{S}^{-1}$. The correction term to the band structure energy, $\mathrm{Tr}\{(1-\mathrm{DS})\mathrm{H}\}$, is similar to that used in Refs. [6,8,9].

B. INITIALIZATION

The goal of our algorithm is to determine a final set of fully occupied orbitals. This means that if we energy resolve or determine a local density of states for each orbital, the local state density will fall entirely below the Fermi level. It is advantageous to start as close to the final answer as feasible. For a system such as silicon with directional bonds, a natural starting point exists. We believe that this will be the case for most closed shell molecules and insulating solids. We proceed here as if silicon is the rule rather than the exception. In a later section, we use fcc calcium as an example, and find that for metallic systems a natural starting point is less obvious. However, this is only an issue at the first geometry or the first time step of a MD simulation. Provided we have good transferability then the result at one geometry will be a good starting point for another geometry. This will similarly be the case for any system made up of subunits that are the same as a system for which we have a solution from a prior calculation.

We assume that we have a LCAO or some other spatially localized representation of our system. Here we use a minimal basis set, non-orthogonal tight-binding representation. For concreteness we use silicon as an example. Silicon has four valence electrons; this means the minimal basis consists of four orbitals, an s and three p's for each atom. For a supercell with n Si's, the basis has 4n functions and the number of occupied orbitals is N=2n. We first transform from the s, and three p basis at each silicon site to a sp^3 basis, where each sp^3 orbital points directly at another silicon atom. In the diamond lattice structure the four sp^3 orbitals per silicon are orthogonal to each other but not to the orbitals at the neighboring Si atom. In fact, in general, the overlap between two nearest neighbor orbitals pointing directly at each other (along a Si-Si bond direction) is rather large. This is typical of an ab initio LCAO calculation. If we energy resolve each sp^3 orbital, then by symmetry half the spectral weight falls below the Fermi level and half falls above. However, making a transformation to symmetric ("bonding") and anti-symmetric ("anti-bonding") combinations of two nearest neighbor sp^3 orbitals pointing along a Si-Si bond not only eliminates the largest overlap in the metric but also places the spectral weight of the symmetric combination (a valence bond orbital) primarily below the Fermi level. For example, using the non-orthogonal tight-binding parameters with up to third neighbor interactions of Mattheis and Pattel [MP] [40], the nearest-neighbor overlap of directed sp^3 orbitals is ~0.6 and the spectral weight of the symmetric combination that falls below the Fermi level is ~0.95. Hence, the iterative procedure needs only to refine out the remaining 5%. This 5% reflects itself in the average energy. For example, with the MP tight-binding parameters, a valence bond orbital in a 96 atom Si cluster has an average energy (determined with a full inversion of **S** and sampling the Brillouin zone at the Γ-point only) of -4.726 eV (see Figure 1). This energy is in error by only 0.3 eV or 6%, at the zeroth iteration.

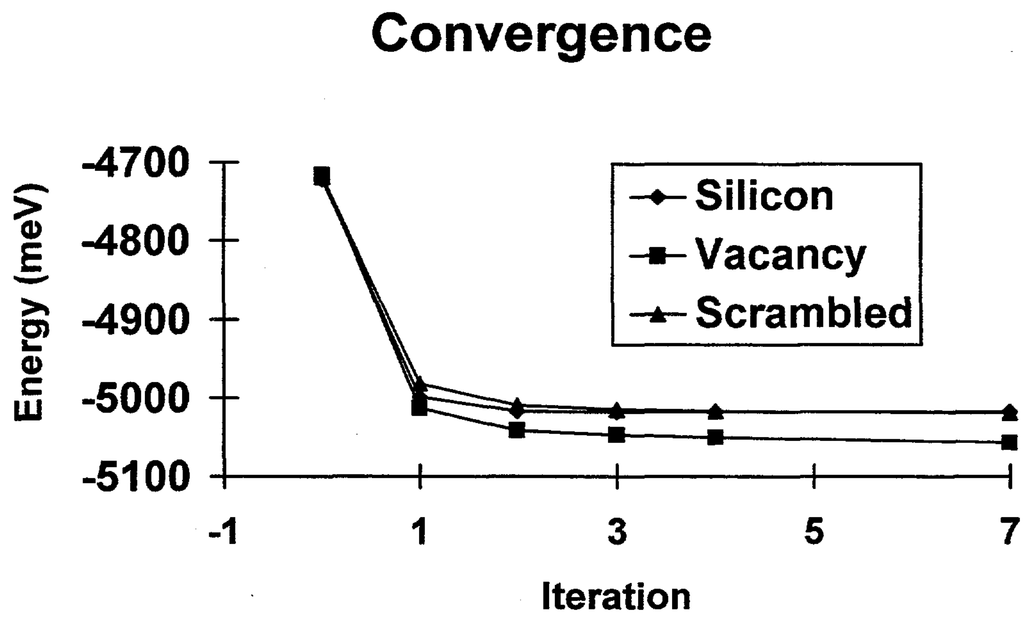

Fig. 1. Convergence of the average energy per occupied orbital for a supercell (96 atoms) of crystalline silicon, randomized ("scrambled") silicon (10% randomization from the crystalline positions) and crystallline silicon with one vacancy (an open-shell system). The final point is the exact answer (as determined by full diagonalization).

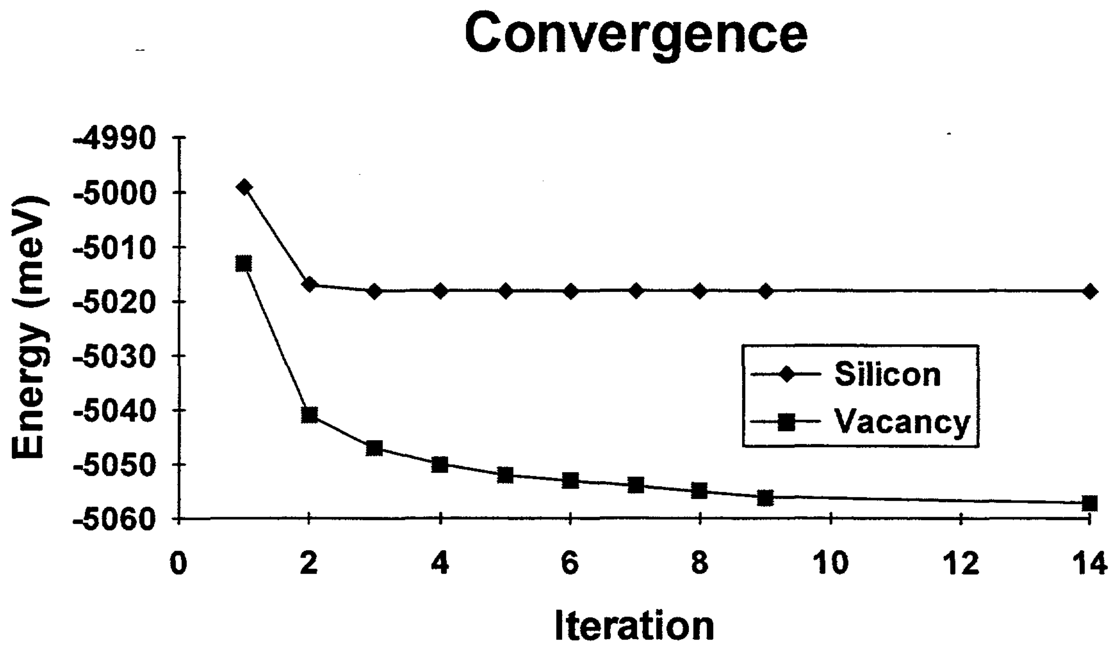

Fig. 2. Convergence of the average energy per occupied orbital for a supercell (64 atoms) of crystalline silicon with and without one vacancy The final point is the exact answer (as determined by full diagonalization). Note the total range of the energy scale is only 7 meV compared to 400 meV in Fig. 1.

Figures 1 and 2 show the convergence of our algorithm for a 96 atom Si cluster in the diamond lattice, slightly randomized about the crystalline positions and with a vacancy. What is plotted is the average energy per occupied orbital after each iteration. Note that after two iterations the error is only 3-5 meV which is sufficient for many purposes. After five iterations the energy for the non-defected Si cluster is converged to better than 1 meV. However, close to convergence improving the accuracy is slow as can be seen in Fig. 2. Should increased accuracy be necessary a different iterative scheme (probably one that updates all the orbitals at once[5], and/or one which uses conjugate directions) should improve this situation. For example, if we update all the orbitals at once the overlap matrix no longer will have a form which can be analytically inverted.

---

[5] Changing one ϕ at a time is advantageous far from convergence. We expect that this will alleviate problems of charge sloshing in large systems. However, near convergence we expect that it might be advantageous to switch to a global iterative scheme that relaxes all the orbitals simultaneously.

However, near convergence the *change* in the overlap matrix is very small and thus it will be sufficient to approximate $D^{(n)}$ either as

$$D^{(n)} = D^{(n-1)} - D^{(n-1)}\left[S^{(n)} - S^{(n-1)}\right]D^{(n-1)}$$

or

$$D^{(n)} = 2D^{(n-1)} - D^{(n-1)}S^{(n)}D^{(n-1)}$$

If necessary, one conjugate gradient iteration could be used to improve D.

For completeness, we have also started the iterations from atom centered orbitals (as in Ref. 4 and 8) and found that convergence is achieved with six or seven iterations. Hence the valence bond orbitals are not a necessary starting point, but they are physically motivated and likely to be more transferable. Note that the number of iterations to reach convergence is independent of the cluster size. The results shown in Figs. 1 and 2 were determined with essentially full inversion; (m=N or infinite range) at the zeroth iterate and thus were not computed in order N. This is still significant because it shows that numerical accuracy was not degraded through the ϕ-updates. We next consider a conjugate gradient determination of **D** with a cutoff radius on the complement basis.

C. "INVERSION" OF A WELL-CONDITIONED S BY CONJUGATE GRADIENT MINIMIZATION

In this section we discuss the nature of the bi-orthogonal complement functions. We consider a column by column conjugate gradient (CG) determination of **D** with various cutoff radii on the complement basis in the non-orthogonal tight-binding silicon system. The kth column of **D** satisfies the following system of equations

$$\sum_j \mathbf{S}_{lj} \mathbf{D}_{jk} = \delta_{lk}. \tag{19a}$$

We choose a cutoff radius of R_c to define the D-range of k and simply set all $\mathbf{D}_{jk} = 0$ for j outside this range; i.e., if the center of the jth bond is a distance greater than R_c from the center of the kth bond. If there are m functions in the D-range of k then the system of equations involves an mxm sub-matrix of **S**. The $\mathbf{D}_{jk}$ for j within the D-range of k can then be computed with a CG algorithm to minimize the function $\mathrm{f}_k(\mathbf{D}) = \sum_j \lambda_j^{(k)} \left(\delta_{jk} - [SD]_{jk}\right)^2$ for each k. The $\lambda_j^{(k)}$ are weight factors (weighted least-squares), for example, $\lambda_j^{(k)} =$ for j in the D-range of k and 0 otherwise or $\lambda_j^{(k)} = \left[S_{jk}\right]^2$. If we have no other estimate of $\mathbf{D}_{jk}$ from prior calculations, we start the iteration with $\mathbf{D}_{jk} = \delta_{jk}$ or $\mathbf{D}_{jk} = 2\delta_{jk} - \mathbf{S}_{jk}$.

We do, however, find it useful to make two corrections to the **D** determined by the straightforward CG solution [41]. First, we guarantee that **D** is symmetric, either by post-symmetrization or by solving for the lower triangle only. The latter method changes the system of equations (Eq. 19a) to

$$\sum_{j \ge k} \mathbf{S}_{lj} \mathbf{D}_{jk} = \delta_{lk} - \sum_{j<k} \mathbf{S}_{lj} \mathbf{D}_{kj}^{*}. \tag{19b}$$

Solving for the lower triangle works well in a sequential scheme since then the $j<k$ terms are already evaluated. The second correction that we find useful fixes $\mathbf{D}_{kk} = [\mathbf{DSD}]_{kk}$ to guarantee that the diagonal elements of **D** are the diagonal elements of the complement metric. Recall that **D** serves a dual role. It is supposed to be the inverse metric and it is also supposed to be the transformation from the basis to the complement basis. The second correction guarantees that at least the diagonal elements exactly satisfy this dual role. It can also be seen (in Appendix C) that this is a necessary condition to assure charge conservation. Note that with $\mathbf{D}_{kk} = [\mathbf{DSD}]_{kk}$ then the diagonal elements are invariant to the transformation $\mathbf{D} \rightarrow 2\mathbf{D}\text{-}\mathbf{DSD}$, the transformation to guarantee that the band structure energy is greater than $\mathrm{Tr}[\mathbf{S}^{-1}\tilde{\mathbf{H}}]$. Our approach has been to add a correction to the diagonal elements of **D**; i.e., $\mathbf{D}_{kk} \rightarrow \mathbf{D}_{kk} + d_k$. Then each diagonal correction, d_k, is independent of the others and satisfies a simple quadratic equation.

Future work will still need to investigate whether alternative solutions are advantageous, e.g., minimizing, for each k, either $f_k(\mathbf{D}) = \sum_j \lambda_j^{(k)} \left(D_{jk} - [DSD]_{jk}\right)^2$ or $f_k(\mathbf{D}) = \frac{1}{2}[DSD]_{kk} - D_{kk}$. Another appealing possibility is to minimize the band structure energy $\mathrm{Tr}\{(2\mathbf{D} - \mathbf{DSD})\tilde{\mathbf{H}}\}$.

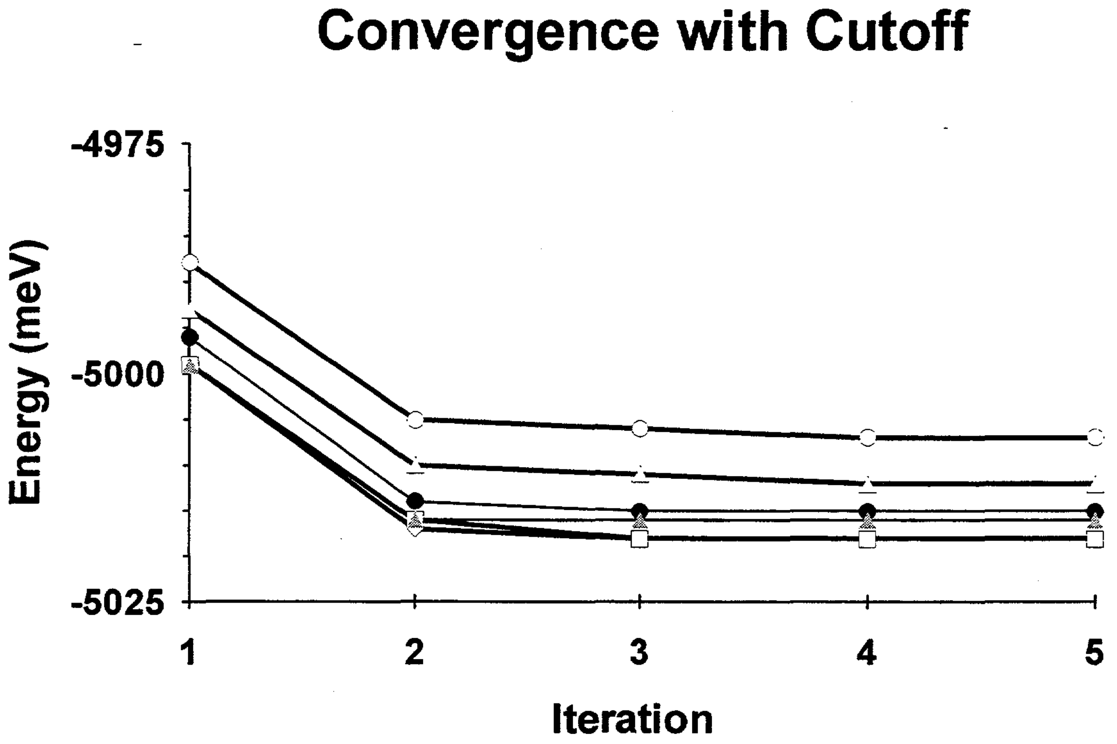

Fig. 3. Convergence as a function of iteration number and cutoff for a 96 atom supercell of crystalline silicon. Open circles: $A_c = 18$, Opened Triangles: $A_c = 24$, Opened Squares: $A_c = 32$ and Opened Diamonds: $A_c = \infty$ (indistinguishable from open squares). Filled circles: Iterating with $A_c = 18$ and relaxing the cutoff after convergence. Filled Triagonals same as Filled circles except $A_c = 24$. See text for further description.

Figure 3 shows results for the 96 atom cluster of Si. The metric is for the space spanned by the 192 valence bond orbitals without any further preconditioning. We compare the energy per orbital (Tr{**(2-DS)H**} for several values of the cutoff radius. At any given iteration, the open markers reflect energies as computed with the approximate D. We will elaborate on the filled markers momentarily. Note that even with a root-mean-square error (rmserr) of 10^{-3} and a cutoff of $R_c=2.52r_0$ (where r_0 is the nearest neighbor distance; in the figure this is labeled $A_c=18$[6]), the

---

[6] The unitless parameter A_c is related to the square of the cutoff distance R_c as $A_c=3[R_c/r_0]^2$.

computed energy is in error by only 9 meV. This cutoff is actually shorter than the range of **S** which has matrix elements to R_c=2.83r_0 (A_c=24). With a cutoff at the same range as **S** the energy is -4.722 eV (in error by only 4 meV). The results for a range of R_c=3.27r_0 (A_c=32) are indistinguishable from full inversion. Decreasing the rmserr below 10^{-3} has no effect on these energies.

Figure 3 also shows the results of iterating with a less than accurate **D** and improving the accuracy only after stopping the iterations. These are the filled markers shown at iteration # 2, 4 and 5. It can be seen that very little in overall accuracy has been lost. This implies that for the early iterations, there is no need to have a very accurate **D**, and accuracy (by extending the range) can be improved when approaching convergence.

It is important to realize that a sparse **D** is not a general result. In other words, we have not found a method which can invert an arbitrary sparse matrix in order N. It should work, however, for any well-conditioned *metric* where the orbitals are spatially localized. We emphasize that this does not restrict the generality of our method. It simply restricts the set of possible starting orbitals. This is a fairly trivial matter since completeness of the basis is not an issue. In an LCAO calculation very poorly conditioned metrics arise primarily because of the polarization and or diffuse functions. The metric here is for the ultimate minimal basis set. Recall there are only as many orbitals as half the number of electrons. It is easy to select a starting set of orbitals that do not have very large overlaps, and hence will have a well-conditioned metric. What is not easy, however, is to choose a set that is, in general, orthogonal. There is one exception: in an orthogonal tight-binding model the valence bond orbitals can be chosen to form an orthogonal set and consequently produce metrics, **S** and **D**, with a range of zero. We do not consider this to be physical nor representative of an ab-initio calculation.

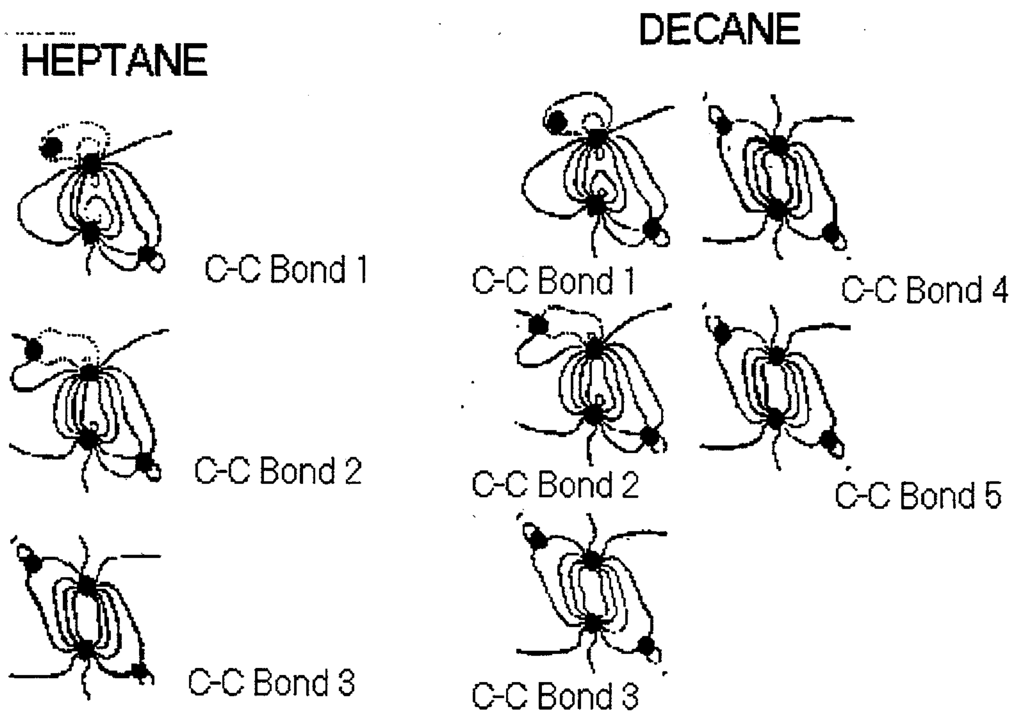

Fig. 4: Converged carbon-carbon bonds for Heptane and for Decane. Bond 1 (the terminal C-C bond) in Heptane is virtually indistinguishable from Bond 1 in Decane, similarly for Bond 2. Bonds 3-5 in Decane are indistinguishable from each other as well as from Bond 3 in Heptane.

To complete this section we show some pictures of converged "bond" orbitals from self-consistent minimal basis density functional calculations. The calculations use Sandia developed software, QuEST [7] but replaces the diagonalization with an algorithm that solves for NOLOs instead of stationary states. The first example is two different length alkane chains, heptane

---

[7] QuEST is a parallel LCAO-DFT code developed by M.P. Sears and P.A. Schultz. The code is based on the linear combination of atomic orbitals. The atomic-orbital basis consist of contracted gaussians. The code exploits sparseness of the hamiltonian and overlap matrices to achieve linear-scaling in the construction of the matrices. The solution of Poisson's equation to obtain the Hartree potential is done very efficiently by separating the charge into a fast and a slow term. The fast term is a sum of spherical neutral atom terms for which Poisson's equation is trivially solved and is short-range. For the slow term, Poisson's equation is solved by a fast Fourier transform on a relatively coarse grid.

(C_7H_{16}) and decane ($C_{10}H_{22}$). Chemically these are very similar as is clearly reflected in the converged "bond" orbitals in Fig. 4. These orbitals are localized on the atoms comprising the bond and the nearest neighbor atoms. Hence a carbon-carbon bond uses the atomic orbitals of four carbons and four hydrogens. The terminal bond uses the atomic orbitals of three carbons and three hydrogens. This enforced localization represents reduced variational freedom and hence the localization error must be positive. However, despite this strong localization the error is ~0.6 meV per occupied orbital. Indeed, if we create the electronic structure for decane using converged orbitals from heptane there is no additional error. Transfering the orbitals from heptane to decane represents a calculation with zero variational freedom. As another example we show, in Fig. 5, some converged orbitals from a more complicated class of molecules, namely C_6H_{13}-CONH-$(CH_2)_n$-CONH-C_6H_{13} with n varying from two to five. The average localization error in these systems is still less than 2meV.

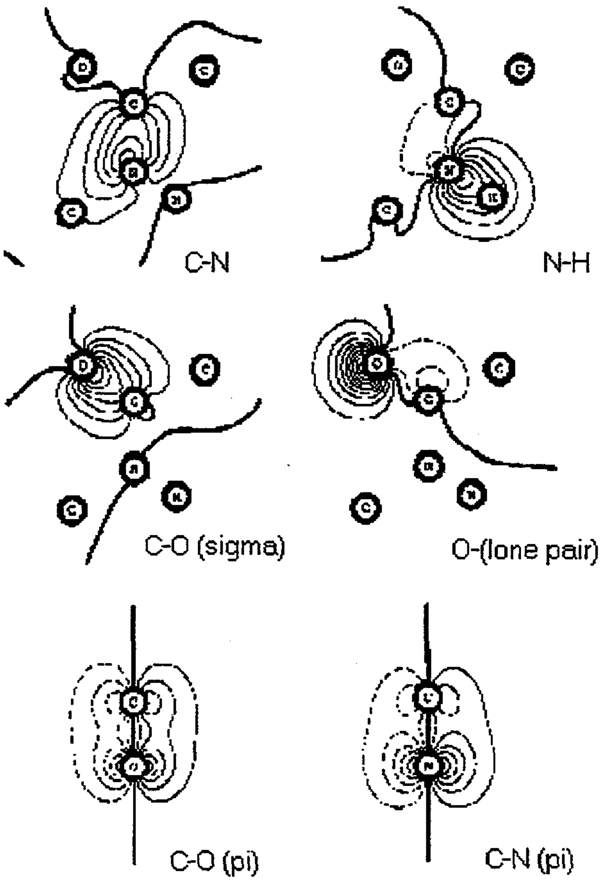

Fig. 5 - Converged NOLOs from molecules containing amide and di-amide fragments.

D. SOLVING FOR EXTREMAL EIGENVALUES IN NON-ORTHOGONAL REPRESENTATIONS: THE GAP

After each full iteration, the Hilbert space of the primitive basis $\left\{\xi_{\mu}\right\}$ has been divided into a P and a Q space. The primitive basis is the full basis in which the problem is being represented, for example, the sp^3 orbitals of the silicon system. P and Q space are approximations to the occupied and unoccupied subspaces, respectively. The eigenfunctions of the hamiltonian in P space are an approximation to the lowest N eigenfunctions. At convergence they *are* the lowest N eigenfunctions. While we have no cause to diagonalize the P space hamiltonian fully, there is reason to determine the largest eigenvalue and possibly the corresponding eigenfunction. By the variational principle this eigenvalue is necessarily greater than the Nth eigenvalue of H (the highest occupied level, HOMO). Similarly, we may want to determine the lowest eigenvalue of Q space which is necessarily lower than the N+1 eigenvalue of H (the lowest unoccupied level LUMO). There are three reasons why we might want to do this. First, we may be interested in determining the gap in an insulator. The difference between these two levels is then necessarily a *lower bound* to the gap (for a fixed one-electron potential). Second, with our algorithm (or any iterative algorithm) there is no guarantee that we won't converge on an excited electronic state. If this is the case then the energy of the supposed HOMO will be greater than that of the LUMO and at least we will know either to change the starting set of orbitals or to swap some levels between P and Q space. Third, in a metal, in general, one does not know how many levels are occupied, at a given point in the Brillouin zone. It will then be necessary to “depopulate” the highest one, two or more "occupied" levels, or to populate a few of the lowest "unoccupied" levels. This can be done within the present algorithm, if we can explicitly solve for just those levels which also has the effect of explicitly including the declocalized character in a metal in an order 1 number of states. We now present an algorithm which is order N for determining the HOMO and the LUMO and is easily generalized to more than one level from each subspace. This method is different from that proposed by Wang and Zunger [42] which would work in order N only for an orthogonal representation. We are unaware of any method in the literature to project out levels in the middle of a band, from a non-orthogonal representation and to do it in order N.

In order to determine the HOMO for P space we need only the sparse N-dimensional metric S and standard representation of the hamiltonian, $\tilde{\mathbf{H}}$ in the "occupied" orbital basis. We take advantage of the fact that the HOMO is an extremal eigenvalue in this Hilbert space. We use a conjugate gradient algorithm to solve a system of equations for ψ_j

$$\sum_j \left[\tilde{\mathbf{H}}_{kj} - \langle E \rangle \mathbf{S}_{kj}\right] \psi_j = 0 \tag{20}$$

where

$$\langle E \rangle = \frac{\sum_{jk} \psi_j^* \tilde{\mathbf{H}}_{jk} \psi_k}{\sum_{jk} \psi_j^* \mathbf{S}_{jk} \psi_k}. \tag{21}$$

We define an unnormalized residual vector, $\mathbf{v}^{(n)}$, with components

$$\mathbf{v}_k^{(n)} = \sum_j \left[\tilde{\mathbf{H}}_{kj} - \langle E^{(n-1)} \rangle \mathbf{S}_{kj}\right] \psi_j^{(n-1)} \tag{22}$$

Conjugate gradient dictates that we add to the residual a portion of the prior residual. If we do not add a portion of the prior residual, the method would be steepest descents. However, we find that CG cuts down on the number of iterations required. Thus we define

$$\lambda_k^{(n)} = \mathbf{v}_k^{(n)} + \frac{\beta^{(n)}}{\beta^{(n-1)}} \lambda_k^{(n-1)} \tag{23}$$

where $\beta^{(n)}$ is the square of the *vector* norm of $\mathbf{v}^{(n)}$. We use the qualifier, *vector*, to contrast this from the *physical* norm which involves the metric.

$$\beta^{(n)} = \sum_k \left| \mathbf{v}_k^{(n)} \right|^2 . \tag{24}$$

For the first iterate $\lambda_k^{(1)} = \mathbf{v}_k^{(1)}$ and $\psi_k^{(0)} = \delta_{kj}$ where j is any one of the occupied orbitals or the one with the largest diagonal energy. Other choices are obviously possible, including random. To iterate, the two vectors ψ and λ are mixed,

$$\psi_k^{(n)} = a_1^{(n)} \psi_k^{(n-1)} + a_2^{(n)} \lambda_k^{(n)} . \tag{25}$$

The $a_1^{(n)}$ and $a_2^{(n)}$ are chosen so as to *maximize* <E> *and* to normalize ψ with respect to the metric. Small enough $\beta^{(n)}$ signals convergence. We refer to the converged ψ as $\psi^{(\mathrm{HOMO})}$. We should point out that if the basis is orthogonal and if we minimize rather than maximize <E>, then this is similar (but not identical) to the modified Lanczos method for determining the groundstate employed in the work of Moreo and Dagatto [43]. As nearly as we can tell their method is equivalent to steepest descents. As mentioned, we find that changing from steepest descents to conjugate gradient significantly cuts down on the number of iterations required to achieve convergence.

Figure 4 shows the convergence for the zeroth, first and second iterate of P-space for the 96 atom MP tight-binding system. We note that convergence for the energy of the HOMO is not uniform with the iterations for P-space.

Determining the LUMO is similar except we do *not* have an explicit representation for Q space. It is necessary to use the implicit representation using the projector 1-P, but with one caveat. It is necessary to have a reasonably accurate representation for P and not to push too hard on the convergence. Otherwise, the ground state of the full hamiltonian can be projected from the noise. The differences in the algorithm from that which determines the HOMO are: (1) The hamiltonian and the metric are represented in the starting "primitive" basis. (2) Suppose we have iterated so that P space is represented by the $\left|\phi_j^{(n)}\right\rangle$. Then we start the iteration scheme for the

LUMO by choosing an arbitrary orbital $\left|\alpha\right\rangle = \sum_{\nu}\alpha_{\nu}\left|\xi_{\nu}\right\rangle$ and projecting out any P space components; i.e.,

$$Q_{\mu\nu} \equiv \left\langle \bar{\xi}_{\mu}\right|\left[1-P\right]\left|\xi_{\nu}\right\rangle, \tag{26a}$$

$$\psi_{\mu}^{(0)} = \sum_{\nu} Q_{\mu\nu}\alpha_{\nu}. \tag{27a}$$

(3) The residual must be projected into Q space; i.e.,

$$\mathbf{v}_{\mu}^{(n)} = \sum_{j} Q_{\mu\nu} u_{\nu}^{(n)}, \tag{28a}$$

$$u_{\mu}^{(n)} = \sum_{\nu}\left[\widetilde{\mathbf{H}}_{\mu\nu} - \left\langle E^{(n-1)}\right\rangle \mathbf{S}_{\mu\nu}\right]\psi_{\nu}^{(n-1)}. \tag{28b}$$

(4) The mixing is the same as with the HOMO except now the energy is *minimized.* Again a small vector norm for the residual signals convergence. If the norm of the residual is not allowed to get too small, dropping into P-space can be avoided for a reasonably accurate representation of P. We refer to the converged ψ as $\psi^{(\mathrm{LUMO})}$. Figure 5 shows the convergence for the LUMO again for the zeroth, first and second iterate for P of the 96 atom MP tight-binding system. Note the difference in the LUMO and HOMO energies is a variational *lower bound* to the gap. We now show that we can also obtain an improved estimate of the gap which is a variational *upper bound.*

Prior to convergence of P-space these estimates for the HOMO and LUMO energies may not be very good; i.e., when $Q^{(n)}\hat{\mathrm{H}}P^{(n)} \neq 0$ as seen in Figs 4 and 5. However, an improved estimate for the HOMO and LUMO can be obtained, relatively easily. For the LUMO we can make the N+1 dimensional space from $P^{(n)}$ plus the LUMO of $Q^{(n)}$. The HOMO of the resulting N+1-dimensional Hilbert space can be determined by first evaluating the overlap and the standard hamiltonian matrix element of the LUMO with each orbital in $P^{(n)}$.; i.e., adding an N+1st row and column. The iteration is then naturally started with the LUMO which in this representation is simply $\psi_k^{(0)} = \delta_{k,N+1}$. The energy of the HOMO of the N+1 dimensional system is an *upper bound* to the N+1 eigenenergy (LUMO).

Similarly $Q^{(n)}$ space can be augmented with the HOMO of $P^{(n)}$ and the LUMO of the resulting system is a *lower bound* to the Nth eigenenergy (HOMO). To augment, $Q^{(n)}=1-P^{(n)}$ is replaced by

$$Q_{\mu\nu}^{(n)} \to Q_{\mu\nu}^{(n)} + \left\langle \bar{\xi}_\mu \middle| \psi^{(HOMO)} \right\rangle \left\langle \psi^{(HOMO)} \middle| \xi_\nu \right\rangle. \tag{26b}$$

To start the iteration

$$\psi_\mu^{(0)} = \left\langle \bar{\xi}_\mu \middle| \psi^{(HOMO)} \right\rangle. \tag{27b}$$

The difference between the energies of the HOMO and the LUMO from the augmented calculations gives a variational *upper bound* to the gap. Figure 6 summarizes the very fast convergence for the upper and lower bounds to the HOMO, LUMO and gap energies.

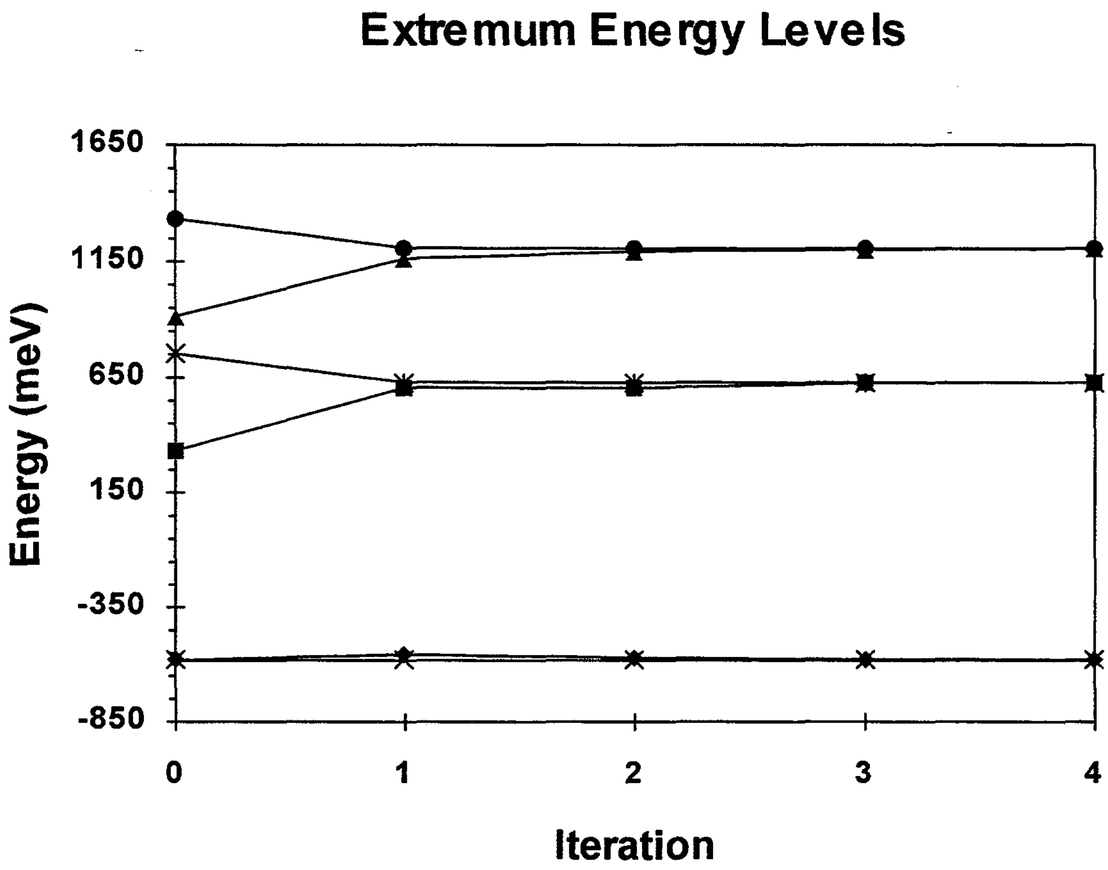

Fig. 6. Upper and lower estimates of the HOMO, the LUMO and the gap for the 96 atom supercell of crystalline silicon as a function of full iteration number. Circles, x's (upper curve) and diamonds are upper bound estimates of the LUMO, the gap and the HOMO energies, respectively. Triangles, squares and x's (lower curve) are lower bound estimates of the LUMO, the gap and the HOMO energies, respectively.

IV. METALLIC SYSTEM: FCC CALCIUM

For a metallic system there is a conceptual issue to address. In particular, can the density matrix be represented in terms of order N localized orbitals and order unity delocalized orbitals? For a large unit cell all the bands will be relatively flat and the majority of bands will be occupied across the entire Brillouin zone. Any band that is fully occupied should be localizable. The relatively few bands that cross the Fermi level will embody the delocalized character of the metal. Alternatively, the density matrix will localize at finite temperature -- but this will require a generalization to partially occupied orbitals. We will consider this below.

In this section we explicitly consider a fcc metallic system with two electrons and thus one occupied orbital per atom, e.g., calcium. Again we use a minimal basis set, non-orthogonal tight-binding model. We take the parameters from Papaconstantopoulos [43]. In choosing a starting set of orbitals we look for a set which has a well-conditioned metric and which satisfies the octet rule for s and p electrons; i.e., each atom should be surrounded by four orbitals. Unlike the case of silicon, this cannot be accomplished with two-centered valence bond orbitals. Nonetheless, it can be accomplished with four-center orbitals centered at tetrahedral interstitial sites [45,46]. A tetrahedral interstitial site is a site in the lattice which is surrounded by four atoms in a tetrahedral arrangement. In the fcc lattice each atom is surrounded by eight such sites. We can thus pair two of them as one orbital or center orbitals at half the tetrahedral interstitial sites. The latter is conceptually simpler. For the latter case the tetrahedral interstitial orbital is a symmetric linear combination of four orbitals one from each of the four nearest neighboring atoms. The orbital from each atom is a s-p hybrid which points its positive lobe directly at the tetrahedral interstitial site. The s-p mix is arbitrary, but 25% s, i.e., an sp^3 hybrid, leads to a metric which is best conditioned. The iteration adjusts the mix and convergence seems independent of the starting mix. Whether this is still the case when self-consistency is added remains to be investigated.

Figure 7 shows the convergence for this case with one variation. Rather than a symmetric combination of the four orbitals pointing at a tetrahedral interstritial site the mix of the four has been randomized[8]. The periodic cluster is for 32 atoms of calcium which at the Γ point has only

---

[8] The parameter r (for degree of randomization) is used as follows: each of the four orbitals is multiplied by a random factor, $\cos(rx\pi)$ where x is a uniform random number between 0 and 1.

30 occupied states. (For the present discussion we assume that somehow we have determined the Fermi level). Thus it is necessary to depopulate the highest two. The slow-down around between the third and seventh iteration (with no randomization) is due to finding a local minimum. The problem arises because there is a set of occupied orbitals (four-fold degenerate) which is missed by symmetry. It is this "missed" state that slows convergence. We find this in itself a very intriguing result. Indeed, a similar result occurs with atomic-centered s-orbital initialization. We return to this below. Randomization mitigates the local minimum and convergence (to 0.1 mRy accuracy) can be achieved with relatively few iterations, 15-25. In these calculations the unit cell is too small to investigate whether the orbitals can be localized. Additionally, a k=0 only calculation is insufficient.

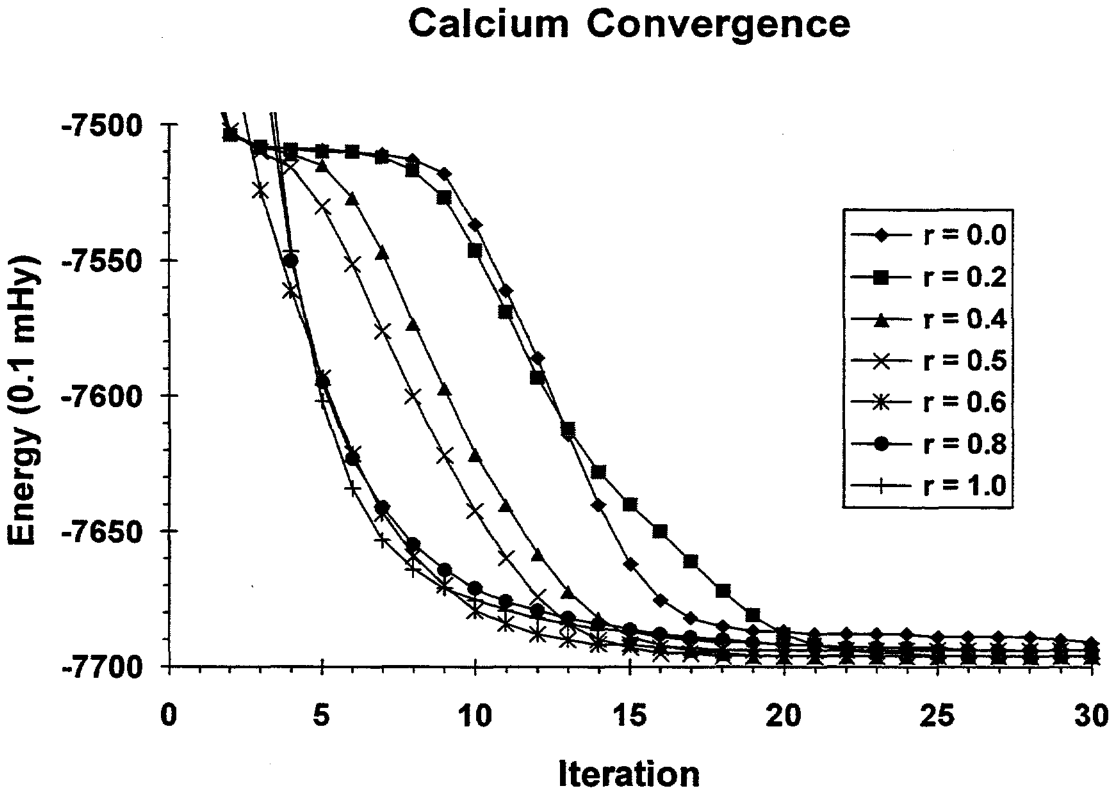

Fig. 7. Convergence of the average energy per occupied orbital in a 32 atom supercell of crystalline fcc calcium. The different curves represent a different degree of randomization in the initial starting set of orbitals[9].

---

[9] The parameter r (for degree of randomization) is used as follows: each of the four orbitals is multiplied by a random factor, $\cos(rx\pi)$ where x is a uniform random number between 0 and 1.

We have also started from atom centered orbitals with arbitrary sp mix[10]. We find that this starting set converges in slightly more iterations than the tetrahedral interstitials, but nonetheless it does converge in relatively few iterations. Hence the algorithm is robust and can converge from a variety of starting points including random as long as the initial guess is localized. One thing to note is: when we use full inversion, the iterative scheme does not find its way out of a local minimum. Hence, it appears errors in **D** allow the possibility of escaping from a local minimum. However, we do not believe that using "noise" to get out of a local minimum is a satisfactory solution to the problem of converging to an excited state.

The surprising result here was that the algorithm converged quickly (but to a local minimum) for some starting guesses, such as A_1 symmetry about the tetrahedral interstitial or A_1 symmetry about the atom center. What we think this means is that there are closed-shell like overlapping “invariant-subspaces” in metals. This concept can be used to project out a minimal basis set for metals. If P_0 is the occupied subspace and we solve for P such that $P_0 \subset P$ then the Baroni and Giannozzi [3] method can be suitably adapted; i.e.,

$$\langle r|g_{pp}(z)|r\rangle = \langle rP|\frac{1}{H-z}|\mathrm{Pr}\rangle.$$

By having the projection operator P the bandwidth or emax problem is overcome. The charge-density and band-structure energy are still obtainable by a contour integral which crosses the real axis at the Fermi energy. The starting vector for the Baroni & Giannozzi method, rather than being the unit vector, is a column of the projection operator P; i.e,

$$\langle r'|\mathrm{Pr}\rangle = \Sigma_j \bar{\phi}_j(r')\phi_k(r) = \Sigma_{jk}\phi_j(r')D_{jk}\phi_k(r)$$

The Fermi level is determined to conserve charge. An advantage of this approach over, for example that of Hernandez and Gillan [16] method, is that the invariant subspaces are defined separately and any subspaces that fall entirely below the Fermi level are out of the picture. For example, suppose we are looking at aluminum. It has 3 electrons per unit cell. Like calcium it will

---

[10] If the sp mix has no p-character, i.e., pure s orbitals, the algorithm converges to an excited state.

have an invariant subspace with A_1 symmetry about the tetrahedral interstitials. However, unlike calcium it will now be fully occupied and hence closed. We then must find another invariant subspace which includes the remaining occupied states and determine the partial occupancy as in Ref. 3. However, now the partial occupancy will always be a finite and invariant fraction of the total whereas in Ref. 3 it is a decreasing fraction with grid-size.

V. LINEAR RESPONSE

The most important thing to recognize about linear scaling Kohn-Sham density functional theory is that there are no stationary states of any hamiltonian. This contrasts traditional approaches that directly solve for the lowest N stationary states of the Kohn-Sham hamiltonian, which is the functional derivative of the energy with respect to the density. Furthermore, all existing approaches to linear response require at least the occupied stationary states. Therefore, to go beyond LDA to linear-response theory in large systems, it becomes necessary to develop a new approach.

Recall that in our linear scaling algorithm the basic non-linear equation that the orbitals satisfy is

$$(1-P)H_{KS}\phi_n = 0$$

It is from these basic quantities, the $\{\phi_j\}$, that we now derive linear response.

Linear Response with no stationary states can be derived from perturbation theory. We expand all quantities in the basic equation through first order changes

$$(1-P-\lambda\,\delta P)(H_{KS} + \lambda\delta V_{ext} + \lambda\,\delta V_{SCF})|\phi_n + \lambda\,\delta\phi_n\rangle = 0.$$

After expanding δP and δV_{SCF} in terms of ϕ_n and $\delta\phi_n$, we arrive at an equation of the form Ax=b.

$$(1-P)(H_{KS})|\delta\phi_n\rangle+(1-P)(\delta V_{ext} + \delta V_{SCF})|\phi_n\rangle-\delta P H_{KS}|\phi_n\rangle = 0$$

The variation δP is

$$\delta\mathbf{P} = \sum_n|\bar{\phi}_n\rangle\langle\delta\phi_n| + \sum_n|\delta\phi_n\rangle\langle\bar{\phi}_n|$$

The variation δV_{ext} is the external perturbation and the operation of δV_{SCF} is the change in the screening potential. The operation of δV_{SCF} on $|\phi_n\rangle$ is

$$\langle r|\delta V_{SCF}|\phi_n\rangle = \int d^3r' \frac{\delta^2 E}{\delta\rho(r)\delta\rho(r')}\langle r'|\delta\mathbf{P}|r'\rangle = 2\phi_n(r)\sum_m \int d^3r' \frac{\delta^2 E}{\delta\rho(r)\delta\rho(r')}\phi_m(r')\delta\phi_m(r')$$

Finally

$$\langle r|(1-P)(H_{KS})|\delta\phi_n\rangle - \sum_m \langle\phi_m|H_{KS}|\phi_n\rangle\delta\phi_m(r) + 2\phi_n(r)\sum_m \int d^3r' \frac{\delta^2 E}{\delta\rho(r)\delta\rho(r')}\phi_m(r')\delta\phi_m(r') = -(1-P)\delta V_{ext}|\phi_n\rangle$$

A is thus an operator in the space which is a cross-product of occupied and unoccupied subspaces; **x** and **b** are vectors in this same space. The dimensionality of this space scales as N^2. The first term is an electron term; it is diagonal in P space. The second term is a hole term; it is diagonal in Q-space and the last term is a coupling term, off-diagonal in both P and Q. The components of **x** are $\delta\phi_n(r)$. The components of **b** are $-(1-\mathbf{P})\delta V_{ext}\phi_n(r)$. A representation of the operator **A** is

$$\langle n,r|A|m,r'\rangle = \langle r|(1-P)(H_{KS})|r'\rangle\delta_{nm} - \langle\phi_m|H_{KS}|\phi_n\rangle\delta(r-r') + 2\phi_n(r)\frac{\delta^2 E}{\delta\rho(r)\delta\rho(r')}\phi_m(r').$$

If the ϕ_n and δV_{ext} are localized the vectors **x** and **b** are sparse. Indeed, the number of non-zero elements in the vectors scale as order 0. Similarly the operation of the operator **A** on the sparse vector can be performed in order 0. To get second derivatives, with respect to atom positions, for example, the difficult term in $\delta^2E/\delta R_{i\alpha}\delta R_{j\beta}$ is Tr $(\delta\rho/\delta R_{i\alpha})(\delta V_{ext}/\delta R_{j\beta})$ which in this formalism becomes $\mathbf{b}^{i\alpha\dagger}\,\mathbf{A}^{-1}\,\mathbf{b}^{j\beta}$ where $\mathbf{b}^{i\alpha}{}_n(r) = -(1-\rho)\delta V_{ext}/\delta R_{i\alpha}\,\phi_n(r)$. The required matrix element of the inverse of **A** can be computed directly using Krylov space methods, such as Lanczos [29]. By analogy with fixed hamiltonians we interpret **A** as the effective electron-hole interaction or evolution operator. It does consist of the hamiltonian for the free electron minus the hamiltonian for a free hole plus a coupling term that scatters electron-hole pairs. A representation of the operator requires only having a representation of the occupied subspace. It does not require having stationary states. The major technical challenge to implement this formalism will be to perform the operation of **A** on a localized vector very efficiently. Extending this to frequency dependent responses requires determining matrix elements of the form -Im $\mathbf{b}^\dagger\ [\mathbf{A}-\omega]^{-1}\ \mathbf{b}$. Different responses involve different vectors **b**, but the basic operation and the major part of the effort is in the operation **Ab**.

VI. CONCLUSION

In this report we have presented a linear scaling density functional algorithm that iteratively solves for localized, fully occupied, non-orthogonal orbitals (NOLOs) that form a basis for the occupied subspace of the Kohn-Sham hamiltonian. These orbitals are potentially transferable as we have demonstrated in a few cases. The linear scaling arises from keeping all matrices sparse. We have also demonstrated that by controlling the condition of the overlap matrix of the starting set of orbitals (first guess), it is possible to use a sparse, approximate inverse of the overlap matrix and analytic inverses during the iterative updates. The only requirement is that the "residual" be orthogonal to the current guess for the occupied sub-space. We cast the algorithm in terms of explicitly non-orthogonal functions so that the computed occupied orbitals will have physical meaning and will thus be transferable between similar systems. The updates (at least far from convergence) are state by state rather than global with the future goal in mind of controlling any charge-sloshing instability [23] of the Kohn-Sham self-consistent energy. The required numerical effort to update a single state is order m^2, where m is independent of N for large N.

Appendix A: Update Equations

In this Appendix, we give the non-obvious and non-trivial update equations which involve the complement basis. Given the mixing angle, θ_k. Appendix B gives details for determining the mixing angle that minimizes the total band structure energy. We find that the generalized Jacobi rotation for the kth orbital rotates the kth complement function as follows:

$$\left\langle \bar{\phi}_k^{(n)} \right| = g_k(\theta_k)\left[\cos(\theta_k)\left\langle \bar{\phi}_k^{(n-1)} \right| - \mathbf{D}_{kk}^{(n-1)}\sin(\theta_k)\left\langle \chi_k^{(n)} \right|\right] \quad \text{(A.1a)}$$

The change in the complement functions for j≠k is

$$\left\langle \bar{\phi}_j^{(n)} \right| = \left\langle \bar{\phi}_j^{(n-1)} \right| - \sin(\theta_k)\, g_k(\theta_k)\, \mathbf{D}_{jk}^{(n-1)}\left[\sin(\theta_k)\left\langle \bar{\phi}_k^{(n-1)} \right| + \cos(\theta_k)\left\langle \chi_k^{(n)} \right|\right]$$

(A.1b)

Notice that the change in $\left\langle \bar{\phi}_j \right|$ is proportional to the overlap of $\bar{\phi}_j$ with $\bar{\phi}_k$ and, hence, only a $\bar{\phi}_j$ for j in the D-range of k changes. The bi-orthogonality relationship is verifiably preserved with these updates.

Also the inverse matrix is easily updated as follows:

$$\underline{\underline{\mathbf{D}}}^{(n)} = \begin{bmatrix} D_{kk}^{(n)} & \left[d_q^{(n)}\right]^{\dagger} \\ d_q^{(n)} & D_{qq}^{(n)} \end{bmatrix} \quad \textbf{(A.2)}$$

where $D_{kk}^{(n)}$ is a scalar equal to

$$D_{kk}^{(n)} = g(\theta_k)\, D_{kk}^{(n-1)} \qquad \text{(A.2a)}$$

The N-1 dimensional vector $d_q^{(n)}$ is

$$d_q^{(n)} = \cos(\theta_k)\, g(\theta_k)\, d_q^{(n-1)}. \qquad \text{(A.2b)}$$

Finally the (N-1)-dimensional matrix $D_{qq}^{(n)}$ is

$$\underline{\underline{\mathbf{D}}}_{qq}^{(n)} = \underline{\underline{\mathbf{D}}}_{qq}^{(n-1)} - g_k(\theta_k)\sin^2(\theta_k)\, \underline{\mathbf{d}}_q^{(n-1)} \left[\underline{\mathbf{d}}_q^{(n-1)}\right]^\dagger. \qquad \text{(A.2c)}$$

Note that the total operation count for these updates is order m^2 where m is number of elements in the D-range of k. We can easily verify that $\mathbf{D}^{(n)}$ is the metric for the updated complement space *and* that it is the transformation that relates the updated basis and complement basis.

Eliminating the work to invert **S** is a major step toward efficiency of our algorithm, but a major part of the remaining work can also be eliminated in updating the proper representation of the hamiltonian. This reduces the operation count from order m^3 (for matrix multiplication) to order m^2.

That update proceeds as follows: The off-diagonal elements of the kth row where only the bra (complement) function changes are

$$\mathbf{H}_{kj}^{(n)} = g_k(\theta_k)\left[\cos(\theta_k)\,\mathbf{H}_{kj}^{(n-1)} - \mathbf{D}_{kk}^{(n-1)}\sin(\theta_k)\, h_j^{(n)}\right] \quad \textit{for } j \neq k$$

(A.3a)

where the $h_j^{(n)}$ were defined in Eq. 17d. The qq block (l,j ≠k) where again only the bra (complement) function changes is

$$\mathbf{H}_{jl}^{(n)} = H_{jl}^{(n-1)} - g_k(\theta_k)\sin(\theta_k)\mathbf{D}_{jk}^{(n-1)}\left[\sin(\theta_k)H_{kl}^{(n-1)} + \cos(\theta_k)h_j^{(n)}\right] \tag{A.3b}$$

Finally, the kth column, where both the bra (complement) and the ket functions change, can easily be accomplished with a matrix times a vector multiplication.

$$H_{jk}^{(n)} = \sum_l D_{jl}^{(n)} \tilde{H}_{lk}^{(n)} \tag{A.3c}$$

We should point out that this does generalize to finite k-point sampling. However, NOLOs only make sense if they can localize within the unit cell, which implies that the overlap with its image in another unit cell vanishes. The equations then easily generalize with all images updating simultaneously.

Appendix B: Minimization of the Total Band-Structure Energy with Mixing Angle

This Appendix first explains how we can analytically minimize the real part of $\mathbf{H}_{kk}^{(n)}(\theta)$, then shows that a small change and essentially no additional work allows us to minimize $E^{(n)}(\theta)$ directly. The state contribution is:

$$\mathrm{Re}\left\{\mathbf{H}_{kk}^{(n)}(\theta)\right\} = \mathrm{Re}\left\{\mathbf{H}_{kk}^{(n-1)}\right\} + g_k(\theta)\sin(\theta)\left[\sin(\theta)\,\Delta_k - \cos(\theta)\,\widetilde{\beta}_k\right] \tag{B.1}$$

where

$$g_k(\theta) = \frac{1}{\cos^2(\theta) + \mathbf{D}_{kk}^{(n-1)}\sin^2(\theta)} \tag{B.2a}$$

$$\Delta_k = \mathbf{D}_{kk}^{(n-1)}\left[h_{\chi\chi}^{(n)} - \mathbf{H}_{kk}^{(n-1)}\right] \tag{B.2b}$$

$$\widetilde{\beta}_k = \mathrm{Re}\left\{\left\langle\overline{\phi}_k^{(n-1)}\middle|\hat{H}\middle|\chi_k^{(n)}\right\rangle + \mathbf{D}_{kk}^{(n-1)}h_k^{(n)}\right\} \tag{B.2c}$$

where $h_{\chi\chi}^{(n)}$ and $h_k^{(n)}$ were defined in the text (Section III.A Eq. 17d). Recall that $\mathbf{D}_{kk}^{(n-1)}$ is the norm of $\left|\overline{\phi}_k^{(n-1)}\right\rangle$, and $\mathbf{H}_{kk}^{(n-1)}$ is the kth diagonal of the proper representation of the hamiltonian. Algebra shows that minimization is achieved with an angle θ_k such that

$$\tan(2\theta_k) = \frac{\beta_k c_k - \Delta_k \gamma_k}{\Delta_k c_k + \beta_k \gamma_k} \tag{B.3}$$

where

$$\beta_k = \frac{1}{2}\left(D_{kk}^{(n-1)} + 1\right)\tilde{\beta}_k \tag{B.4a}$$

$$\gamma_k = \frac{1}{2}\left(D_{kk}^{(n-1)} - 1\right)\tilde{\beta}_k \tag{B.4b}$$

$$c_k = \sqrt{\Delta_k^2 + \mathbf{D}_{kk}^{(n-1)} \tilde{\beta}_k^2} \tag{B.4c}$$

Clearly the standard Jacobi rotation (2x2 diagonalization) is recovered for the orthogonal case, $\mathbf{D}_{kk}^{(n-1)} = 1$ (.$\gamma_k = 0$). We call this more general rotation a generalized Jacobi rotation.

Given the results in Appendix A, the total band structure energy, as a function of the mixing angle, is

$$E^{(n)}\left(\theta_k\right) = E^{(n-1)} + \sin\left(\theta_k\right) g_k\left(\theta_k\right)\left[\sin\left(\theta_k\right)\Delta_k - \cos\left(\theta_k\right)\tilde{\beta}_k\right] \tag{B.5}$$

where

$$\Delta_k = \mathbf{D}_{kk}^{(n-1)} h_{\chi\chi}^{(n)} - \left\langle \bar{\phi}_k^{(n-1)} \middle| \hat{\mathbf{H}} \middle| \bar{\phi}_k^{(n-1)} \right\rangle \tag{B.6a}$$

$$\tilde{\beta}_k = \left\langle \overline{\phi}_k^{(n-1)} \middle| \hat{H} \middle| \chi_k^{(n)} \right\rangle + \left\langle \chi_k^{(n)} \middle| \hat{H} \middle| \overline{\phi}_k^{(n-1)} \right\rangle \tag{B.6b}$$

The change in the minimization from minimizing the kth state energy to minimize the total band energy amounts to only a redefinition of the two parameters $\tilde{\beta}_k$ and Δ_k.

Appendix C: Update to the Charge Density

In problems for which self-consistency is necessary it may be useful to update the charge density at the same time as the updates. Given the update equations of Appendix A we can derive the *change* in the charge density. Recall that the charge density is the diagonal elements of the density matrix in a spatial representation:

$$\left\langle r\middle|\hat{\rho}^{(n)}\middle|r\right\rangle = \sum_j \left\langle r\middle|\bar{\phi}_j^{(n)}\right\rangle\left\langle \phi_j^{(n)}\middle|r\right\rangle. \tag{C.1}$$

Inserting the updated functions and complement functions (Recall that only one function and many complement functions get updated. The sum is eliminated using the relation

$$\left\langle r\middle|\bar{\phi}_k^{(n-1)}\right\rangle = \sum_j D_{jk}^{(n-1)}\left\langle r\middle|\phi_j^{(n-1)}\right\rangle. \tag{C.2}$$

The final result is

$$\left\langle r\middle|\hat{\rho}^{(n)}\middle|r\right\rangle = \left\langle r\middle|\hat{\rho}^{(n-1)}\middle|r\right\rangle + g(\theta_k)\sin\theta_k\left[\sin\theta_k f_1(r) - \cos\theta_k f_2(r)\right] \tag{C.3}$$

where

$$f_1(r) = D_{kk}^{(n-1)}\left|\chi_k^{(n)}(r)\right|^2 - \left|\bar{\phi}_k^{(n-1)}(r)\right|^2 \tag{C.4a}$$

$$f_2(r) = \left\langle r\middle|\bar{\phi}_k^{(n-1)}\right\rangle\left\langle \chi_k^{(n)}\middle|r\right\rangle + \left\langle r\middle|\chi_k^{(n)}\right\rangle\left\langle \bar{\phi}_k^{(n-1)}\middle|r\right\rangle. \tag{C.4b}$$

Note that both $f_1(r)$ and $f_2(r)$ are functions whose integral over all space vanishes, (this assures charge conservation) provided that

$$D_{kk}^{(n-1)} = \int dr \left| \overline{\phi}_k^{(n-1)}(r) \right|^2 = \left\langle \overline{\phi}_k^{(n-1)} \middle| \overline{\phi}_k^{(n-1)} \right\rangle \qquad \text{C.5a}$$

and that the residual $\left| \chi_k^{(n)} \right\rangle$ and the complement function $\left| \overline{\phi}_k^{(n-1)} \right\rangle$ are othogonal;

$$\left\langle \overline{\phi}_k^{(n-1)} \middle| \chi_k^{(n)} \right\rangle = 0. \qquad \text{C.5b}$$

For full inversion these identities definitely hold. As we have implemented our algorithm, we simply assure that these identities hold independent of any approximations.

The scaling to update the charge density is only linear in the number of grid points for which the density is changing, assuming this is calculated on a grid. Determining the charge density at each full iteration scales as order Nm^2.

Appendix D: The Residual Vector

In this Appendix, we discuss the mechanics of determining the residual vector and its matrix elements. The kth residual is the "leakage" out of the space defined by the "occupied " orbitals when operating $\hat{\mathrm{H}}$ on the kth "occupied" orbital . We put the word, occupied, in quotes because until there is no leakage the orbitals are not 100% occupied.

We need the residual represented in the primitive and in the complement to the primitive basis; i.e., we need define two vectors $\left\langle \bar{\xi}_\mu \middle| \chi_k^{(n)} \right\rangle$ (covariant representation) and $\left\langle \xi_\mu \middle| \chi_k^{(n)} \right\rangle$ (contravariant representation). The dimension is that of the primitive basis. They are related by the primitive metric. When referring to the primitive basis we use Greek indices. The $\left\langle \bar{\xi}_\mu \middle| \phi_j^{(n)} \right\rangle$ are determined by the update Eq. 15. We, thus, apparently have the option of computing $\left\langle \bar{\xi}_\mu \middle| \chi_k^{(n)} \right\rangle$ by solving a system of equations to determine $\left\langle \bar{\xi}_\mu \middle| \hat{\mathrm{H}} \middle| \phi_k^{(n-1)} \right\rangle$ or alternatively we can evaluate $\left\langle \xi_\mu \middle| \chi_k^{(n)} \right\rangle$ and then solve a system of equations to obtain $\left\langle \bar{\xi}_\mu \middle| \chi_k^{(n)} \right\rangle$. Either way, solving the resulting system of equations is the dominant work in a ϕ-update for non-orthogonal primitive representations. An orthogonal primitive representation (such as a grid based method) would avoid having this system of equations to solve.

We prefer solving the system of equations for $\left\langle \bar{\xi}_\mu \middle| \hat{\mathrm{H}} \middle| \phi_k^{(n-1)} \right\rangle$ rather than directly for the convariant represention of the residual, $\left\langle \bar{\xi}_\mu \middle| \chi_k^{(n)} \right\rangle$, because then we can use $\left\langle \bar{\xi}_\mu \middle| \hat{\mathrm{H}} \middle| \phi_k^{(n-2)} \right\rangle$ as a starting point in a conjugate gradient iterative solution of the system of equations. It is also the most efficient method if we use the approximate η_k determined at the beginning of each iteration to fix the order of ϕ-updates. The $\left\langle \bar{\xi}_\mu \middle| \hat{\mathrm{H}} \middle| \phi_k^{(n-1)} \right\rangle$ would then already be available. Hence

$$\eta_k \left\langle \bar{\xi}_\mu \middle| \chi_k^{(n)} \right\rangle = \left\langle \bar{\xi}_\mu \middle| \hat{\mathrm{H}} \middle| \phi_k^{(n-1)} \right\rangle - \sum_j \left\langle \bar{\xi}_\mu \middle| \phi_j^{(n)} \right\rangle \mathbf{H}_{jk} \tag{D.1a}$$

To quarantee charge conservation in practice (see Appendix C) it is necessary to explicitly orthogonalize $\chi_k^{(n)}$ to $\bar{\phi}_k^{(n-1)}$ which is done at this point in the calculation. The work involved is order m^2. In principle this does not change $\chi_k^{(n)}$, hence we do not indicate this step in the remaining equations. The contravariant representation is obtained from the relation:

$$\left\langle \xi_\mu \middle| \chi_k^{(n)} \right\rangle = \sum_\nu \left\langle \xi_\mu \middle| \xi_\nu \right\rangle \left\langle \bar{\xi}_\nu \middle| \chi_k^{(n)} \right\rangle \tag{D.1b}$$

The normalization, η_k is chosen so that the residual has a physical norm of 1.

$$\sum_\nu \left\langle \chi_k^{(n)} \middle| \xi_\nu \right\rangle \left\langle \bar{\xi}_\nu \middle| \chi_k^{(n)} \right\rangle = 1 \tag{D.1c}$$

Given $\left\langle \bar{\xi}_\mu \middle| \chi_k^{(n)} \right\rangle$ a matrix vector multiply determines $\left\langle \xi_\mu \middle| \hat{\mathrm{H}} \middle| \chi_k^{(n)} \right\rangle$

$$\left\langle \xi_\mu \middle| \hat{\mathrm{H}} \middle| \chi_k^{(n)} \right\rangle = \sum_\nu \tilde{\mathbf{H}}_{\mu\nu} \left\langle \bar{\xi}_\nu \middle| \chi_k^{(n)} \right\rangle \tag{D.2}$$

and simple dot products determine the remaining matrix elements, $h_j^{(n)}$ and $h_{\chi\chi}^{(n)}$.

$$h_{\chi\chi}^{(n)} = \sum_\mu \left\langle \chi_k^{(n)} \middle| \bar{\xi}_\mu \right\rangle \left\langle \xi_\mu \middle| \hat{\mathrm{H}} \middle| \chi_k^{(n)} \right\rangle \tag{D.3a}$$

$$h_j^{(n)} = \sum_\mu \left\langle \phi_j^{(n-1)} \middle| \bar{\xi}_\mu \right\rangle \left\langle \xi_\mu \middle| \hat{\mathrm{H}} \middle| \chi_k^{(n)} \right\rangle \qquad \text{(D.3b)}$$

References

[1] R. Car and M. Parrinello, Phys Rev. Lett. 55 2471 (1985).

[2] W. Yang, Phys Rev. Letts. 66 1438 (1991); Phys. Rev. A 44 7823 (1991); C. Lee and W. Yang, J. Chem. Phys. 96 2408 (1992).

[3] S. Baroni and P. Giannozzi, Europhys. Letts. 17 547 (1992).

[4] G. Galli and M. Parrinello, Phys. Rev. Letts. 69 3547 (1992).

[5] L.-W. Wang and M.P. Teter, Phys. Rev. B 46 12798 (1992).

[6] X.-P. Li, R.W. Nunes and D. Vanderbilt, Phys. Rev. B 47 10891 (1993).

[7] M. S. Daw, Phys. Rev. B 47 10899 (1993).

[8] F. Mauri, G. Galli and R. Car, Phys. Rev. B 47 9973 (1993).

[9] P. Ordejón, D.A. Drabold, M.P. Grumbach and R.M. Martin, Phys. Rev. B 48 14646 (1993).

[10] W. Kohn, Chem. Phys. Letts. 208 167 (1993).

[11] E.B. Stechel, A.R. Williams and P.J. Feibelman, Phys. Rev. B 49 10088 (1994).

[12] W. Hierse and E. B. Stechel, Phys. Rev. B, 50 17811 (1994).

[13] S. Goedecker and L.Colombo, Phys. Rev. Lett. 73 122 (1994).

[14] A. Alavi, J. Kohanoff, M. Parrinello and D. Frenkel, Phys. Rev. Lett. 73 2599 (1994).

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

[16] E. Hernàndez and M.J. Gillan, Phys. Rev. B 51 10157 (1995).

[17] Y. Wang, G.M. Stock, W.A. Shelton and D.M.C. Nichelson, Phys. Rev. Lett. 75 2867 (1995).

[18] W. Yang and T.-S. Lee, J. Chem. Phys. 103 5674 (1995).

[19] J.J.P. Stewart, Int. J. Quant. Chem. 58 133 (1996).

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

[21] A.F. Voter, J.D. Kress and R.N. Silver, Phys. Rev. B 53 12733 (1996).

[22] See for example, M.C. Pease, III Methods of Matrix Algebra, Mathematics in Science and Engineering Vol 16, Academic Press, NY (1965).

[23] I. Stich, R. Car, M. Parrinello and S. Baroni, Phys. Rev. B 39 4997 (1989); M.P. Teter, M.C. Payne and D.C. Allan, Phys. Rev. B 40 12255 (1989); M.C. Payne, M.P. Teter, D.C. Allan, T.A. Arias and J.D. Joannopoulos, Rev. of Mod. Phys. 64 1045 (1992).

[24] W.H. Adams, J. of Chem. Phys. 34 89 (1961); J. of Chem. Phys. 37 2009 (1962); Chem. Phys. Letts. 11 71 (1971); W.H. Adams, Chem. Phys. Letts. 12 295 (1971); Chem. Phys. Letts. 11 441 (1971); Phys. Rev. Letts. 13 1093 (1974).

[25] T.L. Gilbert, in Molecular Orbitals in Chemistry, Physics and Biology, eds. P.-O. Lowdin and B. Pullman (Academic Press, N.Y. 1964), pg. 405.

[26] P.W. Anderson, Phys. Rev. Letts. 21 13 (1968Phys. Rev. 181 25 (1969); J.D. Weeks, P.W. Anderson and A.G.H. Davidson, J. of Chem. Phys. 58 1388 (1973).

[27] D.W. Bullett, J. Phys. C: Solid State Phys. 8 2695, 2707 (1975).

[28] For a brief review of chemical pseudopotentials see, P.W. Anderson, Phys. Reps. 110 311 (1984).

[29] C. Lanczos, J. Res. Nat. Bur. of Standards, Sect. B 45 255 (1950). See also J. K. Collum and R.A. Willoughby Lanczos Algorithms for Large Symmetric Eigenvalue Computations, Progress in Scientific Computing Vol. 3, Birkhäuser (1985).

[30] Solid State Physics edited by H. Ehrenreich, F. Seitz and D. Trumbull (Academic Press, NY) 35 (1980) and references therein.

[31] R. Riedinger, et al Phys. Rev. B 39 13175 (1989) and references therein.

[32] W. Kohn, Phys. Rev. 115, 809 (1959).

[33] E.I. Blount, in Advances in Solid State Physics, ed. F. Seitz and D. Trumbull, Academic Press, N.Y., Vol 13, (1962).

[34] G.H. Wannier, Rev. of Mod. Phys. 34 645 (1962).

[35] A.A. Frost, in Methods of Electronic Structure Theory, ed. H.F. Schaefer III, (Plenum Press, N.Y. and London, 1977), pg. 29 and references therein.

[36] W.M.C. Foulkes, Phys. Rev. B 48 14216 (1993).

[37] K. Ruedenberg, Rev. of Mod. Phys. 34 326 (1962); C. Edmiston and K. Ruedenberg, Rev. of Mod. Phys. 35 457 (1963).

[38] S.F. Boys, Rev. of Mod. Phys. 32 296 (1960); J.M. Foster and S.F. Boys, Rev. of Mod. Phys. 32 300 (1960).

[39] E.B. Stechel, T.A. Schmalz and J.C. Light, J. Chem. Phys. 70 5640 (1979).

[40] L.F. Mattheis and J.R. Pattel, Phys. Rev. B 23 5384 (1981).

[41] W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vetterling, Numerical Recipes Cambridge University Press (1986).

[42] L.W. Wang and A. Zunger, J. Chem. Phys.,(1994).

[43] E. Dagatto and A. Moreo, Phys. Rev. D 31 865 (1985); E.R. Gagliano and S. Bacci, Phys Rev. D 36546 (1987); E.R. Gagliano, E. Dagotto, A. Moreo and F. Alcarez, Phys. Rev. B 34 1677 (1987).

[44] D.A. Papaconstantopoulos, Handbook of the Band Structure of Elemental Solids Plenum Press, NY (1986). For simplicity we have included only the s and p orbitals in the basis.

[45] M.H. McAdon and W. A. Goddard, III, Phys. Rev. Lett. 55 2563 (1985); J. Non-Cryst. Sol 75 149 (1985); J. Phys. Chem. 91 2607 (1987).

[46] R.P. Messmer, R.C. Tatar and C.L. Briant, in Alloying, ed. J.L. Walter, M.R. Jackson and C.T. Sims, ASM International (1988) p. 29.

Distribution

1	MS-1436	C.E. Meyers, 1007
1	MS-1427	G.A. Samara, actg. 1100
10	MS-1421	E.B. Stechel, 1153
1	MS-9018	Central Technical Files, 8523-2
5	MS-0899	Technical Library, 4414
2	MS-0619	Review and Approval Desk, 12630
96		For DOE/OSTI: UC-705