

Order- N Multiple Scattering Approach to Electronic Structure Calculations

Yang Wang, G. M. Stocks, W. A. Shelton, and D. M. C. Nicholson
Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831-6114

Z. Szotek and W. M. Temmerman
Daresbury Laboratory, Daresbury, Warrington WA4 4AD, United Kingdom
(Received 19 June 1995)

A new approach to the calculation of the electronic structure of large systems within the local density approximation is outlined. The electronic structure problem is formulated using real space multiple scattering theory. Employing a *compute-node* \leftrightarrow *atom* equivalence, the method has been implemented on a massively parallel processing supercomputer. The method is naturally highly parallel and ideal order- N scaling is obtained. The convergence of the method is demonstrated by comparison with the result of conventional electronic structure calculation for elemental metals and through calculation of the ordering energy of β brass.

PACS numbers: 71.10.+x, 71.20.Ad

It is now generally accepted that the local density approximation (LDA) to density functional theory [1,2] provides a basis for treating the energetics of condensed many-electron systems with sufficient relative precision to allow for the study of many interesting physical phenomena. In addition, methods of solution of the LDA equations have progressed, thanks to a constructive interplay between computational power and algorithmic development, to the point that *ab initio* calculations can now be performed routinely for systems where the essential physics can be encapsulated by studying from a few tens to a few hundred atoms. However, there are important classes of problems that involve interactions between large numbers of atoms, several hundred to thousands, which remain beyond the reach of LDA methods using existing algorithms and conventional computational technology. Given the availability of a new generation of massively parallel processing supercomputers (MPP), it is opportune to investigate their potential for allowing the extension of LDA methods to this regime.

Extending conventional reciprocal space band structure algorithms for solving the LDA Schrödinger equation to systems or unit cells consisting of a number of atoms, N , in the range of hundreds or thousands is nontrivial due to an order- N^3 [$O(N^3)$] divergence in the amount of work associated with the linear algebra operations of inverting, diagonalizing, or orthogonalizing a matrix whose size is proportional to N . Recently, a number of algorithms [3] for which the computational effort increases only linearly with system size [$O(N)$ scaling] have been proposed. Of these methods, however, the ones that have so far found practical applications to large systems depend on either the use of a pseudopotential approximation to the full (all-electron) electron-ion interaction or the use of a model tight-binding Hamiltonian.

In this Letter we outline an all-electron, real space, multiple scattering theory approach that leads naturally to an $O(N)$ algorithm. We will refer to the new technique

as the locally self-consistent multiple scattering (LSMS) method. The problem of solving the LDA equations for a N -atom system is made tractable by decomposing into N linked locally self-consistent subproblems, and by exploiting the analytic properties of the single particle Green's function and the variational properties of the LDA free energy functional. Importantly, the new approach is specifically designed for application to MPP architectures, such as the Intel MP Paragon XP/S systems, and to heterogeneous networks of computing resources using virtual parallelism packages such as Parallel Virtual Machine. The result is a method that significantly enlarges the domain of applicability of all electron methods such as the Korringa-Kohn-Rostoker (KKR) band structure method [4,5], to which it is closely related.

At the heart of the LSMS method is the observation that a good approximation to the electron density and the density of states (DOS) on a particular atom within a large condensed system, and thereby the total energy of that system, can be obtained by considering only the electronic multiple scattering processes in a finite spatial region centered at that atom. The cluster of M atoms included in the region is referred to as the local interaction zone (LIZ) of the atom. Every atom in the system is considered to be at the center of its own LIZ, and one can use this fact to obtain an approximation to the electron density and DOS, $\rho_M^i(\mathbf{r})$ and $n_M^i(\epsilon)$ ($i = 1, 2, \dots, N$), associated with every atom in the system. The potential to be used in the next iteration of the self-consistent field (SCF) calculation is then reconstructed by solving Poisson's equation for a crystal electron density, $\rho(\mathbf{r})$, made up of a sum of the single site densities:

$$\rho(\mathbf{r}) = \sum_i \rho_M^i(\mathbf{r}) \sigma^i(\mathbf{r}), \quad (1)$$

where $\sigma^i(\mathbf{r})$ is the truncation function [6] for the Voronoi polyhedron that confines atom i . Clearly, this simple procedure is particularly well suited to a parallel computing paradigm if one associates one or more atoms with either

a compute node, in the case of one atom per node scheme, or a task, in the case of several atoms per node scheme.

Given the above algorithm, the electronic structure problem is reduced to that of calculating the single particle Green's function at the central atom of a finite cluster of sites from which $\rho_M^i(\mathbf{r})$ and $n_M^i(\epsilon)$ can be obtained in the standard way. For this, the multiple scattering theory (MST) [4,7] is a particularly useful approach. In the MST one of the important quantities of interest is the scattering path matrix [8], which in real space is given by

$$\tau_M(\{\xi\}_i; \epsilon) = [\underline{T}_M^{-1}(\{\xi\}_i; \epsilon) - \underline{G}_M(i; \epsilon)]^{-1}, \quad (2)$$

$$\rho_M^i(\mathbf{r}) = -\frac{2}{\pi} \text{Im} \int_{-\infty}^{\epsilon_F} d\epsilon \left\{ \sum_{LL'} Z_L^i(\mathbf{r}; \epsilon) [\tau_M(\{\xi\}_i; \epsilon)]_{LL'}^i Z_{L'}^i(\mathbf{r}; \epsilon) - \sum_L Z_L^i(\mathbf{r}; \epsilon) J_L^i(\mathbf{r}; \epsilon) \right\}, \quad (3)$$

where the factor of 2 is for spin, ϵ_F is the Fermi energy, and L represents a combination of angular and azimuthal quantum numbers, l and m . The functions $Z_L^i(\mathbf{r}; \epsilon)$ and $J_L^i(\mathbf{r}; \epsilon)$ are regular and irregular solutions of the single site Schrödinger equation. Clearly, this real space method is highly scalable on MPP's since each node can be assigned the calculation of the τ -matrix elements, the electron density, and DOS for the atom mapped onto it. Typically, angular momenta are kept up to $l_{\max} = 3$ or 4. Construction of the τ matrix at each site requires that each node invert a matrix of dimension $(l_{\max} + 1)^2 M$. Thus, the overall scaling is proportional to $M^3 N$, i.e., $O(N)$ since the prefactor M^3 is independent of system size. If $M \ll N$ this is clearly advantageous compared to the corresponding reciprocal space KKR method since it requires inversion of a matrix of dimension $(l_{\max} + 1)^2 N$, for which the scaling is $O(N^3)$, at a sufficient number of \mathbf{k} points to converge necessary Brillouin zone integrations.

A schematic representation of the mapping of the LSMS algorithm onto a MPP platform is displayed in Fig. 1, where, in the interests of clarity, we assume a one atom per node case. Consider, for example, the i th atom of an N -atom system and assume, although this is not necessary, periodic boundary conditions, and that the LIZ of atom i is confined to its nearest neighbor shell (atoms j, k, l, m in the illustration). The LSMS algorithm proceeds as follows. An initial guess of the potential $v^i(\mathbf{r})$, and electron density $\rho_M^i(\mathbf{r})$, for the i th site, and the positions of all atoms in the system, is loaded onto the i th node. The node determines which other nodes are within its LIZ. It then calculates the t matrix, $\underline{t}^i(\epsilon)$, corresponding to its own potential and requests and receives the t matrices of other atoms (nodes), i.e., $\underline{t}^j(\epsilon)$, $\underline{t}^k(\epsilon)$, $\underline{t}^l(\epsilon)$, and $\underline{t}^m(\epsilon)$, within its LIZ; in addition it sends its own t matrix to the nodes for which it is one of the atoms (nodes) in the remote atoms' LIZ. At this point, the i th node has sufficient information to calculate the real space τ -matrix elements using Eq. (2), and hence the local electron

density using Eq. (3), for the atom it is associated with. Since all nodes execute this procedure, the total electron density is now known [Eq. (1)], and the potential for the next SCF iteration can be constructed. Calculation of ϵ_F and the Madelung potential apart, this is essentially a local process. Calculating the Madelung potential requires knowledge of the boundary conditions of the large cell, as well as the electric multipole moments on the remote sites which have to be exchanged among nodes. The complete process is then repeated until charge self-consistency is achieved.

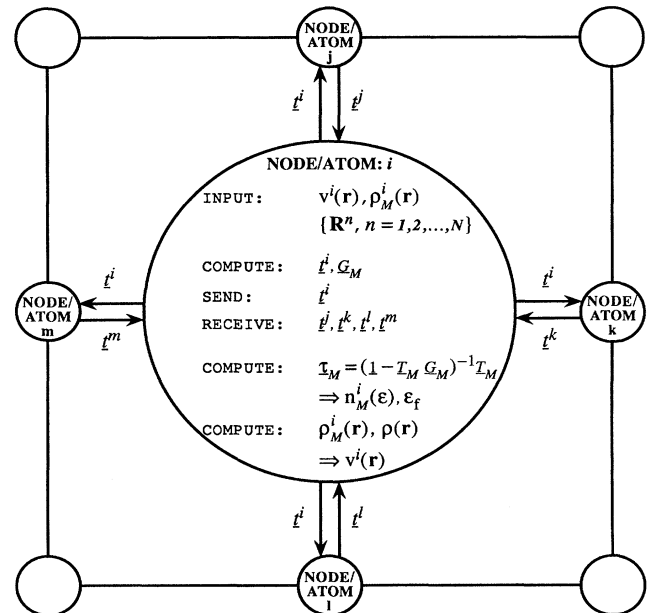


FIG. 1. Schematic of a parallel SCF algorithm: We envision a basic unit cell consisting of N atoms labeled $1, \dots, N$ and about each atom, for example the i th, we define a local interaction zone consisting of M atoms including itself, for example, atoms labeled i, j, k, l, m .

The above LSMS algorithm has been implemented on a distributed memory single-tasking massively parallel message passing platform. To demonstrate its scalability, we applied the calculation to various fcc Cu systems having different numbers of atoms (N) in the unit cell. These systems were constructed by repeating an underlying fcc unit cell along the X , Y , and Z directions until a desired number of sites has been generated. In Fig. 2, we show how the execution time scales with respect to system size up to 1024 atoms. The data were gathered on the 1024-node Intel MP Paragon XP/S-150 MPP supercomputer in the Center for Computational Sciences at Oak Ridge National Laboratory. The solid line represents an ideal $O(N)$ scaling. The open circles in the figure show the logarithm of total execution time per SCF iteration versus the logarithm of N . The fact that they fall onto the solid line illustrates the $O(N)$ scaling property of the LSMS algorithm. Of course, the wall clock times for a four atom per cell run, utilizing four nodes of the machine, and a 1024 atom per cell run, using 1024 nodes, are identical, approximately 117 sec per SCF iteration, illustrating linear scale-up on a MPP supercomputer.

Having established the $O(N)$ scaling behavior, it is now necessary to demonstrate that the convergence of the total energy with respect to the size of the LIZ is sufficient to make the crossover between the conventional KKR method and LSMS acceptable. Clearly, the LSMS method is advantageous for system sizes $N > M_{\min}$, where M_{\min} is the minimum size LIZ that is required to obtain the acceptable precision on the total energy. As was noted recently by Nicholson *et al.* [10] an ex-

pression for the thermodynamic free energy can be devised that is stationary with respect to the variations in the electron density, one electron potential, Fermi function, chemical potential, and temperature. This stationary feature allows use of an approximate electron density, potential and chemical potential for the fully SCF calculated quantities without making a significant error in the energy. As was observed by Nicholson *et al.* [10] it is sufficient to use quite small LIZ's, typically one or two neighbor shells in close packed systems, during the calculation of $\rho_M^i(\mathbf{r})$. However, calculating the total energy using entirely real space methods is more problematic because the total energy is more sensitive to the convergence of the band structure contribution, $\int_{-\infty}^{\infty} \epsilon n(\epsilon) f(\epsilon - \mu) d\epsilon$, where $f(\epsilon - \mu)$ is the Fermi-Dirac distribution function and μ is the chemical potential, than it is to the convergence of $\rho_M^i(\mathbf{r})$.

The convergence of the total energy as a function of LIZ size is shown in Fig. 3 where we show the error in the free energy as a function of LIZ radius, R_M , for elemental fcc Cu and bcc Mo for which energies obtained from conventional \mathbf{k} space KKR calculations can serve as a reference. For Cu better than 0.5 mRy accuracy can be obtained even at zero electron temperature $T = 0$ K by utilizing a LIZ of 87 sites ($R_M = \sqrt{3}a_0$, where a_0 is the lattice spacing). This error can be further reduced to about 0.1 mRy by increasing the LIZ to 135 atoms ($R_M = \sqrt{14}a_0/2$). For Mo the convergence is much worse because the chemical potential falls within the d -band complex. This makes it difficult, at $T = 0$ K, to converge the total energy in real space. However, a large electron temperature in the range $T = 4000 - 8000$ K can be used to ensure tenths of mRy convergence using LIZ's as small as 65 sites ($R_M = 2a_0$) while still allowing extrapolation back to $T = 0$ K in the manner suggested previously [10]. The reason for the improved convergence at very high electron temperature lies in the fact that only energies far off in the complex energy plane, $\text{Im}\{\epsilon\} \geq \pi k_B T$, are involved in the contour integrations used to calculate the band energy and $\rho_M^i(\mathbf{r})$, and for these energies the structure constant matrices $\underline{G}_M(i; \epsilon)$ appearing in Eq. (2) are damped in real space.

The preceding results are monitors of the absolute convergence; more important is the convergence of some quantities of direct physical interest. We have applied the LSMS method to $B2$ CuZn and disordered bcc solid solution $\text{Cu}_{0.5}\text{Zn}_{0.5}$, also known as β brass. The solid solution is modeled by a 256 atom supercell having equal numbers of Cu and Zn atoms which are distributed on the sites of the supercell in such a way that the short range order parameters match the experimental values [11] (this is a trivial matter in the LSMS method). The ordering energy, defined as the energy difference between the ordered and disordered phases, is found to be rapidly convergent in the LIZ size. For a LIZ containing the first two neighboring shells (15 atom cluster) the ordering

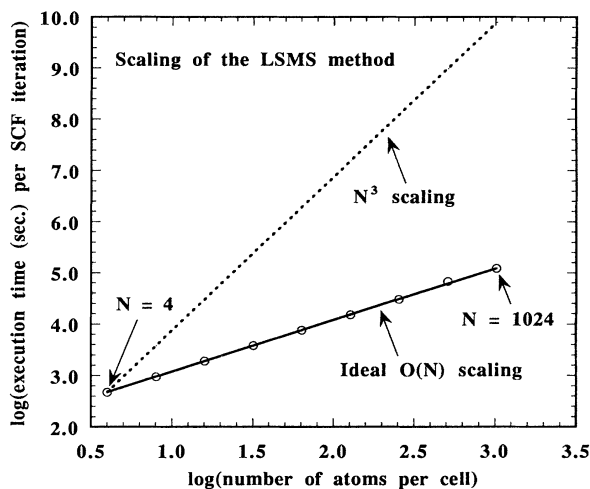


FIG. 2. Scaling behavior of the LSMS algorithm. The calculation is for a large cell model (see text) of fcc Cu. The data were gathered for a LIZ consisting of 13 atoms, sufficient for determining the SCF charge densities. The open circles are the logarithm of the execution time per SCF iteration vs the logarithm of system size. The solid line shows an $O(N)$ scaling, and the dotted line shows an $O(N^3)$ scaling.

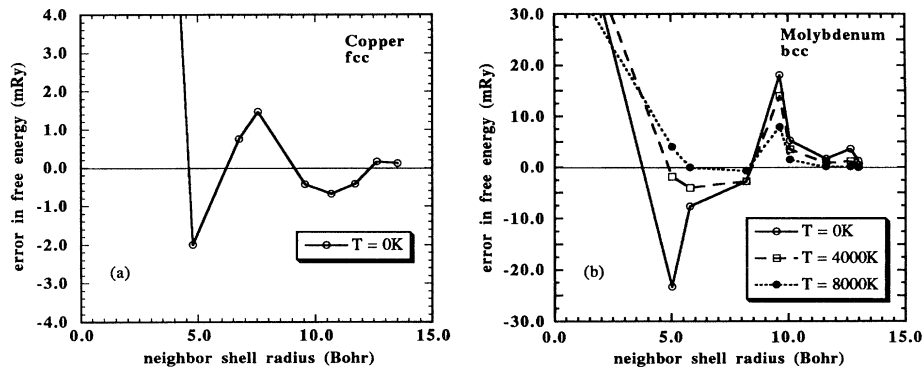


FIG. 3. Free energy error as a function of neighbor shell radius. (a) Cu, fcc, $a_0 = 6.76$ bohr. (b) Mo, bcc, $a_0 = 5.8$ bohr. The free energy errors for a LIZ of one atom ($R_M = 0.0$) are 197.9 mRy for Cu at $T = 0$ K; 69.3, 60.0, and 42.6 mRy for Mo at $T = 0, 4000,$ and 8000 K, respectively.

energy is 3.37 mRy per atom while for a LIZ containing the first seven neighboring shells (89 atom cluster) it is 3.30 mRy per atom. In comparison, the corresponding experimental values range from 1.0 [12] to 3.5 mRy per atom [13].

In conclusion, we have outlined the LSMS method for performing first-principles LDA calculations on systems comprised of a large number of atoms. The method takes advantage of massive parallelism. We have demonstrated its linear scaling on the Intel MP Paragon XP/S-150 MPP supercomputer. One might speculate that the convergence of the SCF procedure may become more difficult with increasing system size detracting from the $O(N)$ scaling. However, we have seen no evidence for this. It is also interesting that we find that quite small supercells, typically $N \sim 128$, are sufficient to produce energies of disordered phases that appear converged with respect to the system size [14]. In its present manifestation, the LSMS code is restricted to metallic alloy systems that can be reasonably treated using the muffin-tin or atomic sphere approximation. However, the code is currently being extended to treat potentials of general shape and to calculate the forces on individual atoms. These capabilities will allow investigation of the metallurgical problems that involve such macroscopic defects as dislocations and grain boundaries.

This work was supported by Mathematical, Information, and Computational Sciences Division of the Office of Computational and Technology Research, and by Division of Materials Sciences, Office of Energy Research, U.S. DOE under Contract No. DEAC05-84OR21400 with Lockheed Martin Energy Systems.

- [1] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964).
 [2] W. Kohn and L.J. Sham, Phys. Rev. **140**, A1133 (1965).
 [3] W. Yang, Phys. Rev. Lett. **66**, 1438 (1991); S. Baroni and P. Giannozzi, Europhys. Lett. **17**, 547 (1992); G. Galli and M. Parrinello, Phys. Rev. Lett. **69**, 3547

- (1992); F. Mauri, G. Galli, and R. Car, Phys. Rev. B **47**, 9973 (1993); X.-P. Li, R. W. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10891 (1993); M.S. Dow, Phys. Rev. B **47**, 10895 (1993); W. Kohn, Chem. Phys. Lett. **208**, 167 (1993); P. Ordejón, D.A. Drabold, M.P. Grumbach, and R.M. Martin, Phys. Rev. B **48**, 14646 (1993); J.R. Chelikowsky, N. Troullier, and Y. Saad, Phys. Rev. Lett. **72**, 1240 (1994); S. Goedecker and L. Colombo, Phys. Rev. Lett. **73**, 122 (1994); E.B. Stechel, A.R. Williams, and P.J. Feibelman, Phys. Rev. B **49**, 10088 (1994); F. Mauri and G. Galli, Phys. Rev. B **50**, 4316 (1994); S.-Y. Qiu, C.Z. Wang, K.M. Ho, and C.T. Chan, J. Phys. Condens. Matter **6**, 9153 (1994); M. Aoki, in "Hume Rothery Award Symposium for Professor David Pettifor: Structure and Phase Stability," edited by G.M. Stocks, C.T. Liu, and P.E.A. Turchi (TMS, Warrendale, PA, to be published).
 [4] J. Koringa, Physica (Utrecht) **13**, 392 (1947).
 [5] W. Kohn and N. Rostoker, Phys. Rev. **94**, 111 (1954).
 [6] Y. Wang, G.M. Stocks, and J.S. Faulkner, Phys. Rev. B **49**, 5028 (1994).
 [7] Lord Rayleigh, Philos. Mag. **34**, 481 (1892).
 [8] B.L. Györfy and M.J. Stott, in *Band Structure Spectroscopy of Metals and Alloys*, edited by D.J. Fabian and L.M. Watson (Academic, New York, 1973), p. 385.
 [9] J.S. Faulkner and G.M. Stocks, Phys. Rev. B **21**, 3222 (1980).
 [10] D.M.C. Nicholson, G.M. Stocks, Y. Wang, W.A. Shelton, Z. Szotek, and W.M. Temmerman, Phys. Rev. B **50**, 14686 (1994).
 [11] C.B. Walker and D.T. Keating, Phys. Rev. **130**, 1726 (1963).
 [12] P.J. Spencer, Calphad: Comput. Coupling Phase Diagrams Thermochem. **10**, 175 (1986).
 [13] M. Ahlers, Prog. Mater. Sci. **30**, 135 (1986).
 [14] Y. Wang, G.M. Stocks, D.M.C. Nicholson, W.A. Shelton, E.L. Hines, Z. Szotek, and W.M. Temmerman, in *High-Temperature Ordered Intermetallic Alloys VI*, edited by J. Horton, I. Baker, S. Hanada, R.D. Noebe, and D.S. Schwartz, MRS Symposia Proceedings No. 364 (Materials Research Society, Pittsburgh, 1995), p. 381.