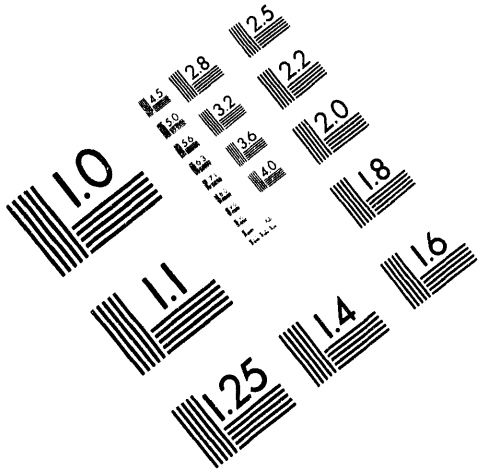

Association for Information and Image Management
1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910
301/587-8202

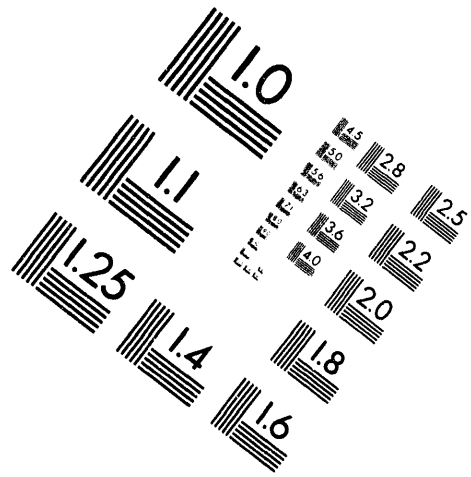

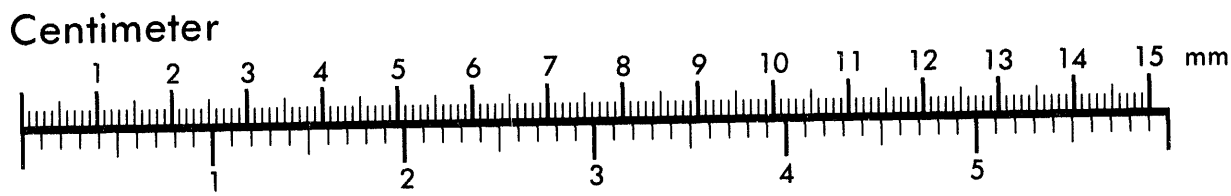

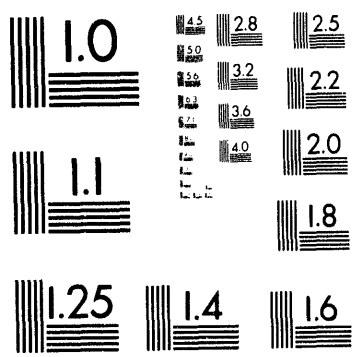

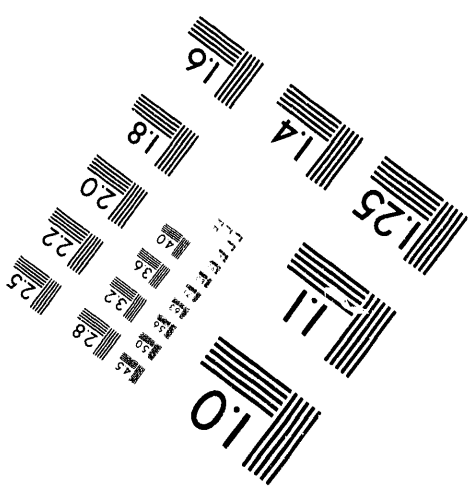

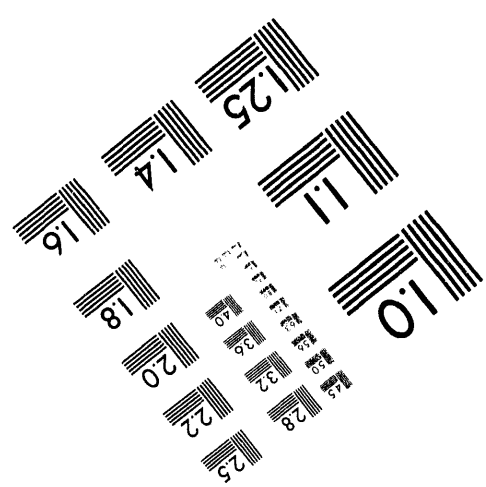

1 of 1

TOWARDS SCALABLE ELECTRONIC STRUCTURE CALCULATIONS FOR ALLOYS

G. M. Stocks, D. M. C. Nicholson, Y. Wang, W. A. Shelton
Oak Ridge National Laboratory, Oak Ridge, TN 37831-6114.

Z. Szotek, W. M. Temmermann *
Max-Plank-Institut Für Festkörperforschung, Stuttgart, Germany

Abstract

A new approach to calculating the properties of large systems within the local density approximation (LDA) that offers the promise of scalability on the massively parallel supercomputers is outlined. The electronic structure problem is formulated in real space using multiple scattering theory. The standard LDA algorithm is divided into two parts. Firstly, finding the self-consistent field (SCF) electron density, Secondly, calculating the energy corresponding to the SCF density. We show, at least for metals and alloys, that the former problem is easily solved using real space methods. For the second we take advantage of the variational properties of a generalized Harris-Foulkes free energy functional, a new conduction band Fermi function, and a fictitious finite electron temperature that again allow us to use real-space methods. Using a *compute-node* $\Rightarrow$ *atom* equivalence the new method is naturally highly parallel and leads to $O(N)$ scaling where N is the number of atoms making up the system.. We show scaling data gathered on the Intel XP/S 35 Paragon for systems up to 512-atoms/simulation cell. To demonstrate that we can achieve *metallurgical*-precision, we apply the new method to the calculation the energies of disordered $Cu_{0.5}Zn_{0.5}$ alloys using a large *random* sample.

INTRODUCTION

Understanding the metallurgical properties of alloys at the quantum mechanical level has been a long sought after goal of condensed matter physics (Hume-Rothery 1969). The underlying assumption being that such a understanding is not only intellectually satisfying but will also aid the development of new or improved materials. In recent years significant progress towards this goal has been made, particularly in the area of understanding the electronic driving mechanisms that underlay alloy phase stability (Stocks *et al.* 1994). These advances have been built on three major pillars. Firstly, the local density approximation (LDA) to density functional theory provides a basis for treating the energetics of condensed many-electron systems of sufficient *relative* precision to answer questions of metallurgical interest. Secondly, significant algorithmic developments have made problem areas once thought to be beyond the reach of LDA methods matters of routine (Stocks *et al.* 1994). Thirdly, the vast increase in computational power afforded by the 1980's generation of vector supercomputers not only allowed more complex calculations but also contributed to a constructive interplay of algorithmic and theoretical developments with available computational power.

Despite these major advances the surface of metallurgical problems has hardly been scratched. There are important classes of problems, involving interactions between large numbers of atoms, that are far beyond current methods and computational technology. The general area of the mechanical behavior of alloys, which involves such *macroscopic* defects as dislocations and grain boundaries, being one such area that is of particular importance to the design of new materials. Given the availability of a new generation of massively parallel supercomputers (MPS) it seems opportune to investigate their potential for allowing us to extend LDA methods to this regime. Applying LDA methods to basic simulation regions or unit cells that contain numbers of atoms, N, in the range of

*Permanent address: SERC, Daresbury Laboratory, Daresbury, Warrington WA4 4AD, UK

DISCLAIMER

100's-1000's is non-trivial; requiring that we address the issue of developing truly scalable algorithms [algorithms for which the computational effort grows linearly with the system size ($O(N)$-algorithms)].

The central problem in applying LDA methods to large systems involves a $O(N^3)$ divergence (the so called N^3-problem) in the amount of work associated with the linear algebra operations of inverting (in multiple scattering theory methods) or diagonalizing (in conventional basis set methods) a matrix whose size is proportional to N. The poor scaling of conventional LDA electronic structure methods generally limits them to $N < 50$ (typically $N < 10$). In addition to the N^3-problem, the straightforward adaptation of conventional methods to MPS by parallelization over reciprocal-space $\vec{k}$-points actually exhibits inverse scaling. This is because the number of $\vec{k}$-points required to obtain good convergence of Brillouin integrations *decreases* with *increasing* system size.

The exception to the above general remarks about system size is the, pseudo-potential based, Car-Parrinello method (Car and Parrinello 1985). For simple systems, *e.g.* Si and Al, and moderate system sizes, $N \sim 100$'s the scaling is $NlogN$ (Singh *et al.* 1992) and is limited by a fast Fourier transform required to judiciously switch between real and reciprocal space. However, as remarked above, the scaling reverts to N^3 for larger N due to an orthogonalization step. In addition, for systems containing elements for which the pseudo-potential is not truly weak, *e.g.* transition metal elements, the number of atoms that can be treated is much reduced because of the need to use more plane waves/atom.

In the remainder of this paper, we will briefly describe a real space multiple scattering theory approach to large systems that is specifically designed for application to massively parallel computer architectures such as the Intel Paragon (Nicholson *et al.* 1994). We call the new approach the Large System Multiple Scattering (LSMS) method. To date, we have demonstrated $O(N)$ scaling on system sizes that can be run on the available machines and project continued $O(N)$ scaling for systems containing thousands of atoms.

OUTLINE OF THE LSMS METHOD

For the purposes of discussion of the LSMS method it is useful to divide the LDA-algorithm (Kohn and Sham 1965) for calculating the total energy of a system into two parts. Firstly, the solution of the Euler-Lagrange equations to obtain the self-consistent electron density and potential corresponding to a given set of atomic positions. Secondly, the calculation of the total energy (and forces on the atoms) corresponding to that self-consistent-field (SCF) electron density. Conventionally, the same electronic structure method is used in to solve the Schrödinger equation in both of these stages. However, in recent years, it has become clear that, by taking advantage of the stationary properties of modified energy functionals of the Harris-Foulkes type (Harris 1985, Foulkes and Haydock 1989), it is possible to obtain accurate energies from approximate electron densities, provided that the kinetic energy is evaluated accurately.

At the heart of the LSMS method is the observation that we can use spatially local real space multiple scattering methods to obtain a solution of the SCF-problem. This approximation yields an approximate electron density of sufficient quality that when used in a Harris-Foulkes energy functional yields an essentially exact energy. In addition, by using an extended finite temperature free energy functional, a newly developed *conduction band* Fermi function, and by performing our calculations at a fictitious finite electron temperature we are able also to obtain accurate total energies using entirely real space methods.

Self-consistency

In the SCF-part of the LSMS method the electron density at the central site of a cluster of scatterers is used as a approximation to the true electron density at that site. That is to say, for the purpose of solving the Schrödinger equation, we replace all of the atoms in the crystal, except those within a few neighboring shells [we refer to this as the local interaction zone (LIZ)], by a constant background potential (in the case of the muffin-tin approximation by the muffin-tin zero, in a full potential calculation a suitable average over the boundary). In order to obtain an approximation to the local electron density on each site in the crystal, we consider every atom in the solid to be at the center of its own local cluster. We use multiple scattering theory to calculate the single particle Green function and hence the electron density on that site. We then reconstruct the potential throughout space

by solving Poisson's equation for a crystal electron density made up of a sum of the single site densities and use it in a mini self-consistency loop.

By utilizing the above device, the electronic structure problem has been reduced to that of calculating the electron density on the central atom of a finite cluster of sites. This simple procedure is particularly well suited to mapping onto a parallel computing paradigm by associating each atom in the system under consideration with a node on a parallel machine (for sufficiently powerful nodes more than one atom could be associated with a single node). This atom/node equivalence is illustrated schematically in fig. 1.

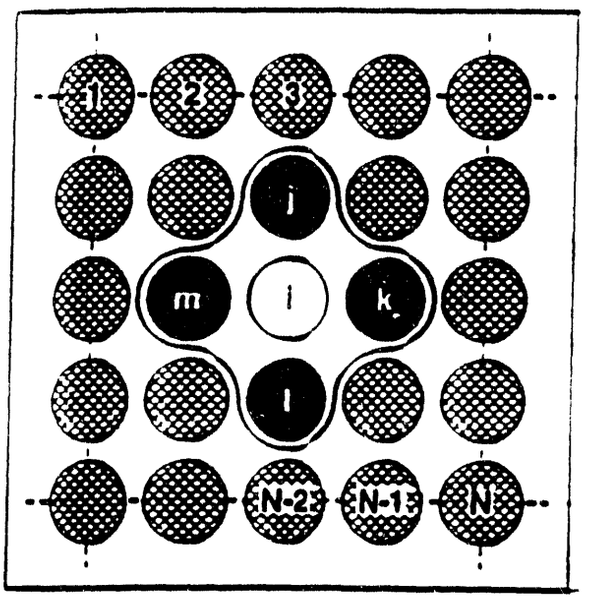

Figure 1: Schematic of a parallel SCF algorithm: we envision a basic unit cell consisting of N-atoms labeled $1, \ldots, N$ and about each atom, for example the i^{th}, we define a local interaction zone consisting of M-atoms including itself, for example atoms labeled i, j, k, l, m.

Given the above breakdown of the SCF problem, the central computational difficulty now becomes calculating the electron density on the central site of an M-site cluster. For this, multiple scattering theory is a particularly useful tool. If we restrict ourselves to muffin-tin potentials, although this is not necessary, the electron density $\rho^i_M(\vec{\mathbf{r}})$ on the central (i^{th}) site of the M-atom cluster is given by (Faulkner and Stocks 1980)

$$\rho^i_M(\vec{\mathbf{r}}) = \frac{-2}{\pi} \Im \int_{-\infty}^{\infty} d\epsilon f(\epsilon - \mu) \left[\sum_{LL'} \left[Z^i_L(\vec{\mathbf{r}};\epsilon) \tau^{ii}_{LL'} Z^i_{L'}(\vec{\mathbf{r}};\epsilon) - Z^i_L(\vec{\mathbf{r}};\epsilon) J^i_L(\vec{\mathbf{r}};\epsilon) \delta_{LL'} \right] \right] . \tag{1}$$

where $\Im$ denotes the imaginary part, $f(\epsilon - \mu)$ is the Fermi-Dirac distribution function, μ is the chemical potential, $Z^i_L(\vec{\mathbf{r}};\epsilon)$ and $J^i_L(\vec{\mathbf{r}};\epsilon)$ are regular and irregular solutions of the Schrödinger equation for a single scatterer, and, $\tau^{ii}_{LL'}$ is the scattering path matrix for the i^{th}-site (Faulkner and Stocks 1980). The scattering path matrix is found from

$$\underline{\tau}^{-1} = \left[\underline{t}^{-1} - \underline{g} \right] . \tag{2}$$

Equation 2 is a matrix equation both in site index i, which labels the position vectors $\vec{\mathbf{R}}_i$ of the atoms in the cluster, and angular momentum index L. The t-matrix $\underline{t}$ and real space structure constant matrix $\underline{g}$ have elements $t_{LL'}(\epsilon)\delta_{ij}$ and $g_{LL'}(\vec{\mathbf{R}}_i - \vec{\mathbf{R}}_j)$ respectively.

Clearly this real space method is highly scalable on MPS's since each node can be assigned the work involved in setting up and inverting the *rhs* of eq. 2 for the atom to which it is assigned. Keeping angular momenta up to $l_{max} = 3$ the dimension of the matrix is $(l_{max}+1)^2 M$. Thus, calculation of the electron density at each site requires

that each node invert a matrix of dimension $(l_{max}+1)^2 M$. If $M \ll N$ this is clearly advantageous compared to conventional reciprocal space methods since these involve the inversion of a matrix of dimension $(l_{max}+1)^2 N$ at a sufficient number of $\vec{\mathbf{k}}$-points to converge necessary Brillouin zone integrations. This advantage is, of course, provided that convergence of $\rho_M(\vec{\mathbf{r}}) = \sum_i \rho_M^i(\vec{\mathbf{r}})$ to the true SCF density $\rho_{scf}(\vec{\mathbf{r}}) \equiv \rho_\infty(\vec{\mathbf{r}})$ is sufficiently rapid.

How the LSMS algorithm maps onto a parallel machine is shown schematically in fig. 2. Consider the i^{th}-

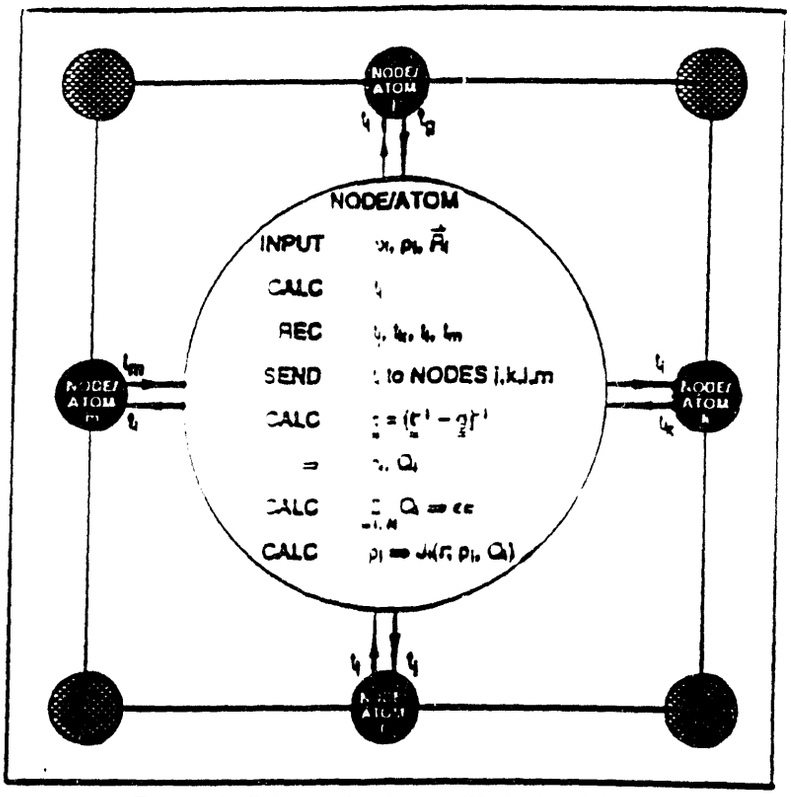

Figure 2: Schematic of the LSMS algorithm.

atom(node) of the N-atoms(nodes) that comprise the total (periodic) system and assume that the LIZ about each atom is confined to its nearest neighbor (NN) shell (atoms j, k, l, m in the illustration). The LSMS algorithm proceeds as follows. An initial guess of the potential $v_M^i(\vec{\mathbf{r}})$, and electron density $\rho_M^i(\vec{\mathbf{r}})$, for the i^{th} site, and the positions of all atoms in the system, are loaded onto the i^{th}-node. The node determines which other nodes are within its LIZ. It then calculates the t-matrix corresponding to its own potential. The i^{th}-node then requests and receives the t-matrices of other atoms/nodes within its interaction zone, in addition it sends its own t-matrix to other nodes for which it is one of the atoms in the remote nodes' LIZ. The i^{th}-node now has sufficient information to set up and invert the real space KKR-matrix, eq. 2, for the atom centered on itself, from which the local electronic charge and charge density can be calculated by repeating this procedure for all the energies on the energy integration contour required by eq. 1. The new $\rho_M^i(\vec{\mathbf{r}})$ can then be used in some appropriate mixing scheme and a new $v_M^i(\vec{\mathbf{r}})$ constructed. Apart from the calculation of the Madelung potential, this is an entirely local process. The calculation of the Madelung potential requires knowledge of the multipole moments on the remote sites which have to be interchanged amongst the nodes. The complete process can then be repeated until charge self-consistency is obtained.

In order to check the convergence of this method we have performed test calculations on a number of systems that are sufficiently small that they can also be treated by conventional methods, in this case the KKR method. For fcc based metals (e.g. Cu, Zn, L_{10}-structure CuZn) we find a local interaction zone consisting of the first shell of neighbors (13-atoms) sufficient to obtain energies that are accurate to 0.01 mRy provided that we use the LSMS-potentials in a full-Brillouin zone integration calculation of the Harris-Foulkes energy (see below). For bcc

based systems (e.g., Fe, Mo, B2-structure NiAl) two shells (15-atoms) are required.

Total Energy

So far we have only described how to obtain high quality electron densities in a scalable manner. As we have mentioned previously, calculation of the band structure contribution to the total energy requires that the Schrödinger equation be solved to a much higher precision than that required for obtaining the electron density. Clearly, this can be done using standard reciprocal space methods. However, this involves one last electronic structure calculation that scales as N^3. It is important that this be avoided.

To this end, we use a newly developed finite temperature *Harris-Foulkes* free-energy functional (Nicholson *et al.* 1994)

$$\begin{aligned} F_H[\rho] &= \int_{-\infty}^{\infty} d\epsilon \epsilon f(\epsilon-\mu) n(\epsilon, v(\rho,\vec{\mathbf{r}})) - \int d\vec{\mathbf{r}} \rho v(\rho,\vec{\mathbf{r}}) + U(\rho) + E_{xc}(\rho) \\ &\quad + \mu[N - \int_{-\infty}^{\infty} d\epsilon f(\epsilon-\mu) n(\epsilon, v(\rho,\vec{\mathbf{r}}))] \\ &\quad - k_B T \int_{-\infty}^{\infty} d\epsilon n(\epsilon, v(\rho,\vec{\mathbf{r}}))[f(\epsilon-\mu)\ln f(\epsilon-\mu) + (1-f(\epsilon-\mu))\ln((1-f(\epsilon-\mu))] \end{aligned} \tag{3}$$

where the effective potential appearing in this equation is the so called *output potential* obtained from

$$v(\rho,\vec{\mathbf{r}}) = \frac{\delta U + E_{xc}}{\delta \rho(\vec{\mathbf{r}})}. \tag{4}$$

This free energy functional has a number of appealing properties that can be effectively deployed in calculating accurate total energies. In particular, $F_H[\rho]$, defined by eqs. 3 and 4, is stationary with respect to the variations in the electron density ρ, effective potential $v(\rho,\vec{\mathbf{r}})$, Fermi function $f(\epsilon-\mu)$, chemical potential μ, and temperature T. These stationarity properties open up the possibility of using the approximate, LSMS, electron density $\rho_M(\vec{\mathbf{r}})$, potential $v_M(\vec{\mathbf{r}})$ and chemical potential μ_M as *stand-ins* for the fully SCF quantities without making a large error in the energy. Equally important, however, is the fact that eq. 3 allows us to perform energy evaluations at fictitious finite electron temperatures.

From the computational point of view, there are a number of advantages to performing the calculations at finite temperature. Firstly, energy integrations to obtain the charge density and band-structure energy are reduced to sums over the *Matsubara*-poles of the Fermi function, which can be made few in number by taking advantage of a new ***conduction band*** Fermi function (Nicholson *et al.* 1994). This greatly reduces the number of matrix inversions that are required to perform a total energy calculation. Secondly, since most of the Matsubara-poles are far off in the complex energy plane, the real space multiple scattering theory is more rapidly convergent, even for the troublesome band structure energy (the first term on the *rhs* of eq. 3). In fact, the Matsubara-pole closest to the real axis is approximately $\pi k_B T$ off in the complex plane, which for electron temperatures of 2500K (4000K) is already 50mRy (80mRy). Thirdly, a calculation at finite temperature provides the electron-hole entropy (last term on the *rhs* of eq. 3) as well as the energy (sum of the remaining terms). This information combined with the fact that the entropy is zero at $T = 0$ allows accurate projection of the energy to $T = 0$.

Again, using simple systems as tests we have demonstrated that we are able to obtain total energies of metallurgical precision. Typically, we use an interaction zone of about 55- (59)-atoms for fcc- (bcc)-structures respectively. For Cu we found that we could use a fictitious electron temperature as high as 4000K and still recover the ground state energy to an accuracy of better than 0.2mRy which is sufficient for many metallurgical applications.

SCALABILITY

The above algorithms have been implemented in a massively parallel computer code. The code assigns each atom in the system to a node of the MPS. Thus, we are able to treat as many atoms as we have nodes on the target

parallel machine. For simplicity, the current version of the code assumes periodic boundary conditions, however, this is not essential.

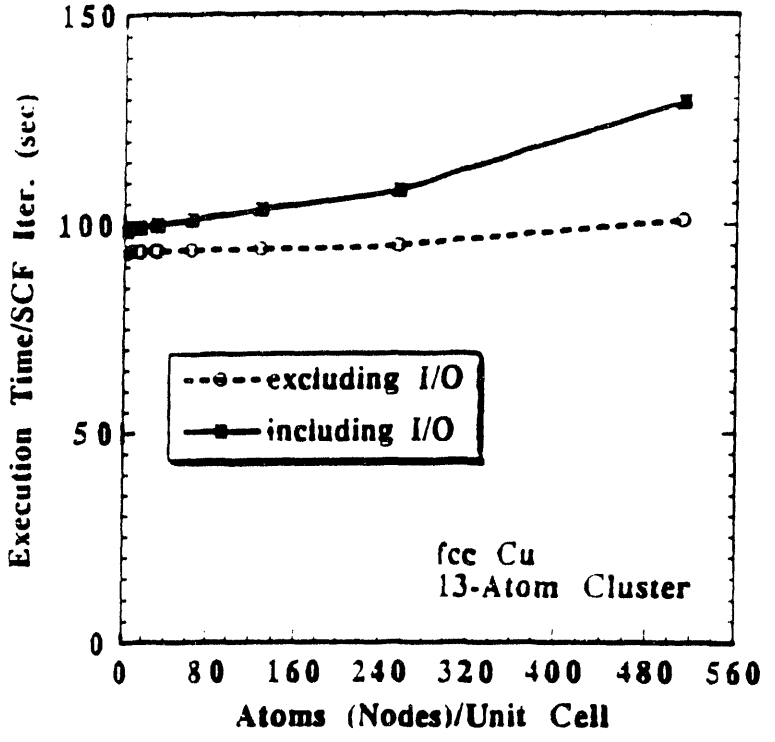

Figure 3: Scaling behavior of the LSMS algorithm showing the execution time per interation for a 10 iteration run for the SCF step. The calculation is for a large cell model of face centered cubic copper for a LIZ consisting of 13-atoms. Solid line includes I/O time, dashed line excludes I/O time.

In the fig. 3, we show how the LSMS algorithm scales with system size. To demonstrate scalability, we use face centered cubic (fcc) Cu because the SCF electron density and energy can be calculated by conventional (KKR) methods. In the LSMS-code, we first represent a single fcc unit cube as four interpenetrating simple cubic sub-lattices and use four nodes of the MPS. We then represent two unit cubes as eight interpenetrating simple tetragonal lattices and distribute the calculation across eight nodes. This procedure is repeated until the node limit of the machine in question is reached, which for the Intel Paragon XP/S-35 on which these calculations were performed is 512. In fig. 3 the solid line shows the total execution time per SCF iteration for a 10 SCF iteration run. The dashed line shows the time per SCF iteration alone, having subtracted out the time required to input the initial guesses of $\rho^i_M(\vec{r})$ and $v^i_M(\vec{r})$ and to output the final *converged* $\rho^i_M(\vec{r})$ and $v^i_M(\vec{r})$.

I/O apart, the LSMS algorithm exhibits essentially ideal $O(N)$-scaling. The time per SCF iteration per node is almost independent of the system size. Since there is little node to node communication in the LSMS algorithm and, at least in the SCF stage, the whole code and associated arrays fit into the available node memory, it is not surprising that the code scales essentially linearly. The departure from ideal $O(N)$-scaling seen when the I/O time is included results from the fact that the Parallel File System (PFS), where the input and output data sets are stored, does not scale ideally. This is because, on the Paragon XP/S-35, PFS is served by only 16 I/O-nodes. Clearly, when all 512 compute nodes are trying either to read or write, all at the same time, PFS becomes swamped.

When we use an interaction zone of 55-atoms to calculate the band structure energy the code has to *page out* since the real space KKR matrix is sufficiently large that it no longer fits onto a node. This difficulty results from the fact that the current XP/S-35 has only 16 megabytes of memory per node and the current operating system kernel is very large ($\approx$ 7-megabytes) resulting in only 9-megabytes of memory being available for applications. The resulting paging causes a reduction in the efficiency of the code, however, the total energy step still scales well.

APPLICATIONS

We have applied the LSMS-code to the direct calculation of the total energy of a large 256-atom/unit cell model of a random solid solution of $Cu_{0.5}Zn_{0.5}$. The sites of a $4 \times 4 \times 4$ repeat of an underlying fcc lattice were randomly occupied by 128 Cu and 128 Zn atoms. Calculation of the short range order parameter revealed that the *sample* had a near ideally random distribution of sites within the 256-atom unit cell. The LSMS-calculation was performed using $T = 0K$, LIZ's of both 13 and 55 atoms in the SCF step, and a LIZ of 55-atoms in total energy step. As expected, the final results were found to be independent of the size of the LIZ in the SCF step.

A quantity of considerable interest is the heat of mixing, $E_{mix} = E_{alloy} - cE_{Cu} - (1-c)E_{Zn}$ where E_{alloy} is the energy of the *random* alloy, E_{Cu} and E_{Zn} are the energies of fcc Cu and Zn respectively and c is the concentration of Cu. We obtain a calculated value of 4.5mRy/atom, this being close to the value calculated on the basis of the KKR-CPA method (Johnson *et al.* 1990). This value is somewhat smaller than the 7.0mRy/atom found on the basis of the most recent charge correlated (CC) KKR-CPA calculations of Johnson and Pinski (Johnson Pinski 1994) and a likely experimental value that is also in the range of 7mRy/atom. Use of the muffin-tin approximation in the present LSMS calculation and the atomic sphere approximation in the CC-KKR-CPA calculations may account for this discrepancy. This aside, the discrepancy is somewhat puzzling in view of the fact that the present calculation, by definition, includes precisely the charge correlations that the work of Johnson and Pinski was supposed to be correcting for and that were presumed absent from the original mean field KKR-CPA calculations.

In fig. 4, we show the calculated charges on Cu and on Zn sites obtained in the 55-atom run plotted as a function of the number of unlike neighbors in the first neighbor shell. Clearly, in a random alloy individual

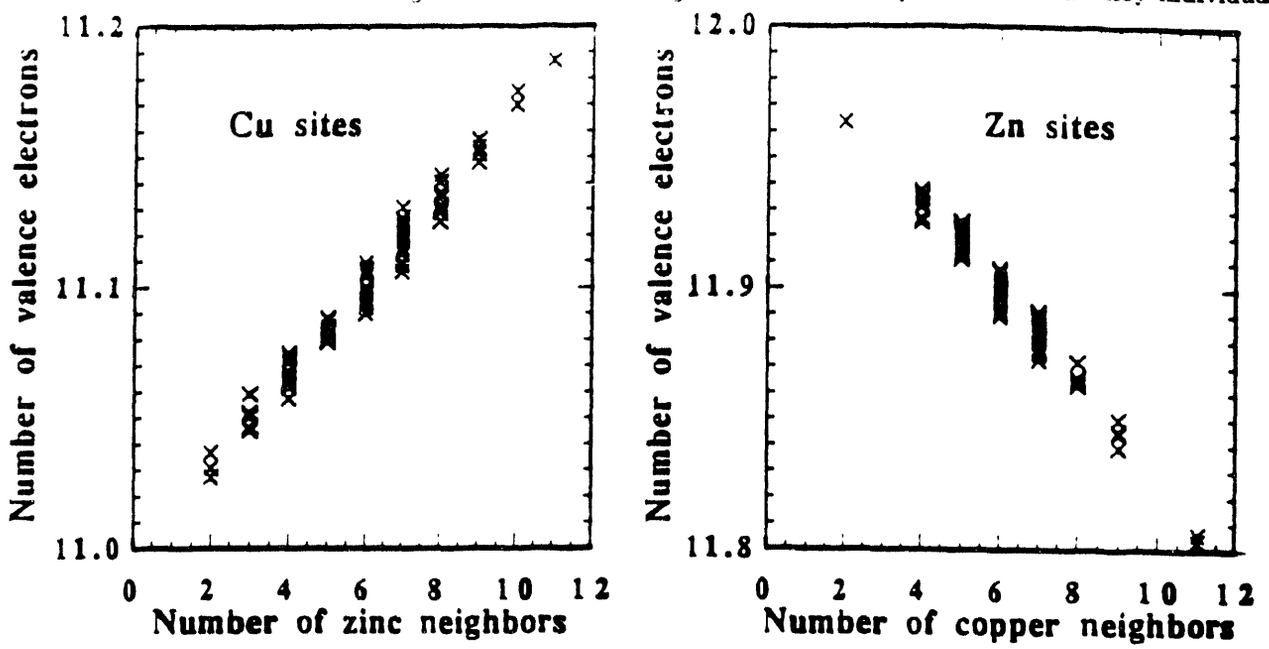

Figure 4: Charges on individual Cu- (left frame) and Zn-sites (right frame) plotted as a function of the number of unlike neighbors in the first neighbor shell taken from a 128-atom simulation of disordered $Cu_{0.5}Zn_{0.5}$.

sites will have different numbers of like and unlike near-neighbors. In fig. 4 are shown the charges for each of the 128-Cu and 128-Zn sites in the 256-site sample. Interestingly, the charges on individual sites are essentially proportional to the number of unlike neighbors. Scatter about the direct proportionality, resulting from the different possible arrangements of a given number of unlike and like neighbors, is clearly a secondary effect. This proportionality was pointed out by Margi *et al.* (Magri *et al.* 1990) in criticizing the mean-field theory SCF prescription conventionally used in the KKR-CPA method (Stocks *et al.* 1994) and underlies the improved energies obtained in the CC-KKR-CPA. In fact, the behavior seen in fig. 4 is almost identical to that found by

Johnso.. and Pinski for the same alloy system.

CONCLUSIONS

We have outlined a new method for performing first principles LDA calculations on large metallic systems. We have demonstrated linear scaling on the Intel XP/S-35 massively parallel supercomputer. We believe this to be the first LDA method for which true $O(N)$ scaling has been attained. In its current manifestation, the LSMS code is restricted to systems that can be reasonably treated using the muffin-tin approximation. However, the code is currently being extended to treat general shape potentials and to allow calculation of the forces on individual atoms. These capabilities will then allow a full relaxation of the internal degrees of freedom, greatly extending the range of applicability of the method.

ACKNOWLEDGEMENTS

Work supported by Division of Materials Science (GMS) and by the Applied Mathematical Sciences Research Program (GMS, DMCN, YW, WAS), Office of Energy Research, US. DOE under subcontract DE-AC05-84OR21400 with Martin-Marietta Energy Systems, Inc..

REFERENCES

Car, R. and M. Parrinello. 1985. "Unified Approach for Molecular Dynamics and Density-Functional Theory." *Physical Review Letters* **55**, no. 22: 2471-2474.

Faulkner, J. S. and G. M. Stocks. 1980. "Calculating Properties with the Coherent-Potential Approximation." *Physical Review* **B21**, no. 8: 3222-3244.

Foulkes, W. M. C. and R. Haydock. 1989. "Tight-Binding Models and Density-Functional Theory." *Physical Review* **B39**, no. 17: 12520-12536.

Harris, J. 1985. "Simplified Method for Calculating the Energy of Weakly Interacting Fragments." *Physical Review* **B31**, no. 4: 1770-1779.

Hume-Rothery, W. and B. R. Coles. 1969. *Atomic Theory for Students of Metallurgy.* Institute of Metals, 17 Belgrave Square, London, S.W.1.

Johnson, D. D. *et al.* 1990. "Total-Energy and Pressure Calculations for Random Substitutional Alloys." *Physical Review* **B41**, no. 14: 9701-9716.

Johnson, D. D. and F. J. Pinski. 1994. "Including Charge Correlations in Calculations of the Energetics and Electronic Structure for Random Substitutional Alloys." *Physical Review* **B**, in press.

Kohn, W. and L. J. Sham. 1965. "Self-Consistent Equations Including Exchange and Correlation Effects." *Physical Review* **A140**, no. 4A : 1133-1128.

Magri, R. *et al.* 1990. "Ground-State Structure and the Random-State Energy of the Madelung Lattice." *Physical Review* **B42**, no. 17: 11388-11391.

Nicholson, D. M. C. *et al.* 1994. "Stationarity of the Density Functional Free Energy: Application to Accelerated Multiple Scattering Calculations." *Physical Review* **B**, submitted.

Singh, D. J. *et al.* 1992. "Projector-Basis Techniques and Car-Parrinello Scaling in Mixing-Basis, Linearized-Augmented-Plane-Wave, and Extended Linearized-Augmented-Plane-Wave Electronic-Structure Methods." *Physical Review* **B46**, no. 20: 13065-13072.

Stocks, G. M. *et al.* 1994. "Electronic Theory of Ordering Phenomena in Metallic Alloys." In *Proceeding of NATO-ASI on Statics and Dynamics of Alloys Phase Transformation*, P. E. A. Turchi and A. Gonis, eds. Plenum Publ..ing Corporation, New York, in press.

DATE
FILMED
7/18/94

END