Multi-teraflops Spin Dynamics Studies of the Magnetic Structure of FeMn/Co Interfaces*

A. Canning,

NERSC, Lawrence Berkeley National Laboratory, Berkeley, CA 94720,

B. Ujfalussy,

Department of Physics and Astronomy, University of Tennessee, Knoxville TN 37996,

T. C. Schulthess, X. -G. Zhang, W. A. Shelton, D. M. C. Nicholson, and G. M. Stocks,

Oak Ridge National Laboratory, Oak Ridge, TN 37831,

Yang Wang,

Pittsburgh Supercomputer Center, 4400 Fifth Avenue, Pittsburgh, PA 15213,

T. Dirks,

IBM, Bethesda, Maryland.

Abstract

We have used the power of massively parallel computers to perform first principles spin dynamics (SD) simulations of the magnetic structure of Iron-Manganese/Cobalt (FeMn/Co) interfaces. These large scale quantum mechanical simulations, involving 2016-atom super-cell models, reveal details of the orientational configuration of the magnetic moments at the interface that are unobtainable by any other means. Exchange bias, which involves the use of an antiferromagnetic (AFM) layer such as FeMn to pin the orientation of the magnetic moment of a proximate ferromagnetic (FM) layer such as Co, is of fundamental importance in magnetic multilayer storage and read head devices. Here the equation of motion of first principles SD is used to perform relaxations of model magnetic structures to the true ground (equilibrium) state. Our code is intrinsically parallel and has achieved a maximum execution rate of 2.46 Teraflops on the IBM SP at the National Energy Research Scientific Computing Center (NERSC).

Keyword phrases: computation-intensive applications, parallel algorithms, material science

1   Introduction

We used the power of massively parallel processing (MPP) computers to implement first principles spin dynamics and to study the magnetic structure of complex inhomogeneous systems. The method we have used rests on three theoretical developments. The equation of motion of first principles Spin Dynamics (SD) [1], [2] which is used to relax model magnetic structures to the ground state or equilibrium T=0K configuration. Constrained density functional theory implemented within the Constrained Local Moment (CLM) model [3], [4] for calculating the magnetic structure, energy, and constraining field associated with the non-equilibrium, non-collinear, magnetic states that are required at each time step of a SD simulation. The Locally Selfconsistent Multiple Scattering (LSMS) method for performing first principles, non-collinear, constrained local moment, electronic structure calculations of a sufficiently large number of atoms for us to be able to realistically simulate the magnetic structure of such complex systems. This implementation of SD relaxation within the CLM-LSM code has allowed us to significantly extend the work that won the Gordon Bell prize at SC98 [6] where we studied CLM states of Iron above the Curie temperature.

We have used the new LSMS code to study the ground state magnetic structure of the technologically interesting interface between a $Fe_{0.5}Mn_{0.5}$ alloy and Co. These first principles quantum mechanical calculations have uncovered details of the interface magnetic structure that cannot be obtained by any other means - either experimentally or theoretically from spin model Hamiltonians. These calculations are being performed in order to shed light on the technologically important phenomenon of exchange bias which is an important ingredient in determining the utility of magnetic multi-layer Giant Magneto-Resistance (GMR) read heads for the latest generation of disc drives. Exchange bias refers to the asymmetric placement of the magnetic hysteresis loop that is induced when a ferromagnetic (FM) material such as Co is placed proximate to an antiferromagnetic (AFM) material such as equiatomic FeMn alloys. While the mechanism responsible for exchange bias is unknown, it is believed that the magnetic structure at the FM/AFM interface is a key ingredient. Of particular interest is the effect on the antiferromagnet of its interaction with a strong ferromagnet such as Cobalt.

Our first principles studies are based on the large cell model of a FeMn/Co interface shown in figure 1. The face centered cubic (fcc) solid solution alloy $Fe_{0.5}Mn_{0.5}$ has long been of scientific interest in that it is an example of a compositionally disordered fcc alloy that is, non-the-less, ordered antiferromagnetically. The AFM ordering in this material is believed to be the so called "3Q" state shown in figure 1. It is a rare example of a peculiar AFM state in which the magnetic moments are non-collinear. In our model interface structure the FeMn alloy is placed next to FM Co in which the magnetic moments are all aligned in the same direction. The microscopic description of the magnetic structure of this interface was the main goal of the current calculation.

More specifically, our large cell model consists of 2016-atoms on an fcc lattice constructed, in such a way that the z axis of the supercell was oriented along the (111) direction in the conventional fcc lattice coordinate system. The resulting structure has been divided into layers perpendicular to the z direction, out of which 6 layers were made up of Cobalt atoms and 15 layers were randomly populated with Fe or Mn to form a $Fe_{0.5}Mn_{0.5}$ alloy system. The orientation of the magnetization on the Co atoms was set parallel to the interface (and set to lay along the x-axis) and according to the 3Q state on the Fe and Mn atoms. Each layer consists of an 8x12 array of sites in the (111)-plane of the underlying fcc structure. The lattice parameter of the fcc lattice is taken to be 6.86 Bohr Radii (atomic units). The above chemical and orientational configuration is the initial state, from which the system is evolved to a different final state, according to the equations of *ab-initio* spin-dynamics([1], [2]). The bottom three layers of Cobalt were held fixed by a constraining field to better simulate the real interface. This means that only three layers of Cobalt were allowed to relax their moment orientations.
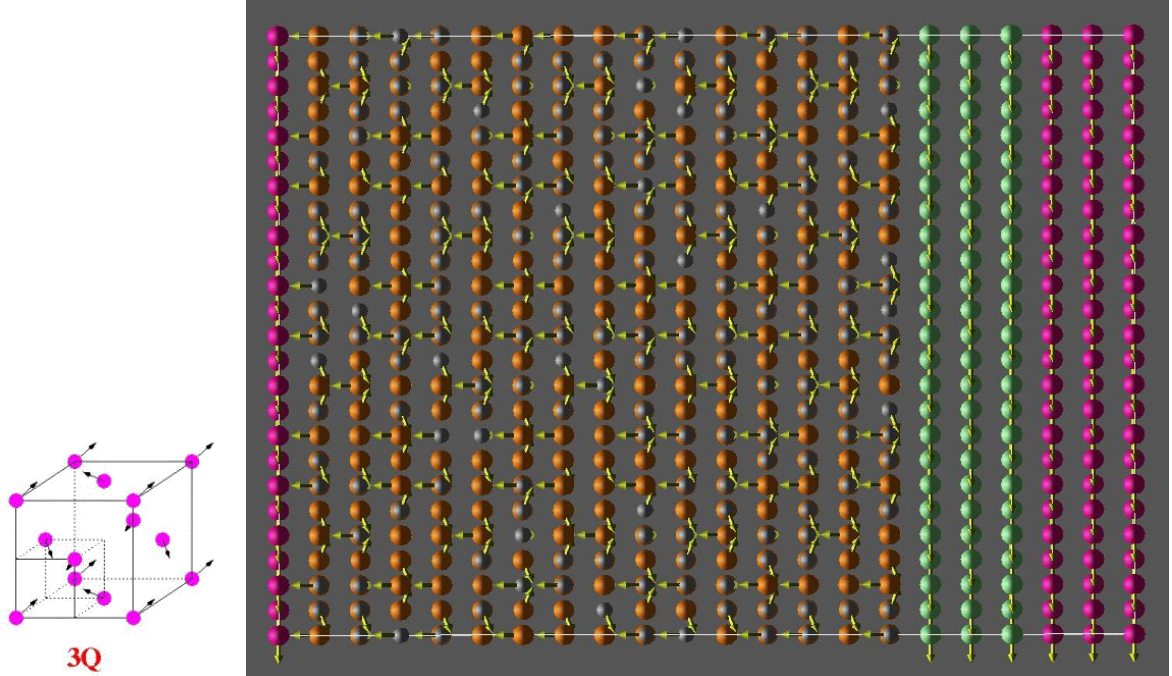
Figure 1: Visualization of the exchange-bias system initial state, viewed from the direction of the y-axis. The arrows represent the orientation of the magnetic moments. The gold balls represent the Fe atoms, the grey ones the Mn atoms, while the green and pink balls stand for the relaxed and unrelaxed Co atoms. On the left, the 3Q antiferromagnetic ordering is schematically displayed within the conventional fcc structure.

In this work the major computational problem, that of performing first principles local density approximation (LDA) density functional theory (DFT) electronic structure calculations on large systems, is solved using the LSMS method. The basic LSMS method is an $O(N)$ LDA method that was specifically designed for implementation on MPPs. In the present work the standard LDA task of self-consistently determining the charge and magnetization densities is made more difficult by the need to also self-consistently determine site dependent magnetic fields that serve to constrain the local magnetic moments to a prescribed set of directions. This is then further compounded by the necessity of doing this for many orientational configurations of the magnetic moments as the system is evolved towards its ground state through repeated application of the equation of motion of SD.

2 Methods and Algorithms

2.1 Methods

2.1.1 Local Density Approximation

Using DFT is now the accepted way of performing first principles electronic structure and total energy calculations for condensed matter [7]. DFT is an exact theory of the ground state, the LDA converts it into a practical computational scheme. DFT-LDA replaces the solution of the many particle Schrödinger equation by the solution of a set of single particle self-consistent field (SCF) equations. The complex many body exchange and correlation effects are included through an effective exchange correlation potential.

In DFT the central quantity of interest is the charge density $\rho(\mathbf{r})$ (and, additionally, in LS(pin)DA the magnetization density $\mathbf{m}(\mathbf{r})$). For a fixed external potential (e.g. that resulting from a set of nuclei) DFT-LDA returns the ground state $\rho(\mathbf{r})$ (and $\mathbf{m}(\mathbf{r})$). For non-spin-polarized systems the SCF equations take the form:

$$\left\{ \frac{\hbar}{2m}\nabla^2 - V_{eff} + \epsilon \right\} G(\mathbf{r},\mathbf{r}';\epsilon) = \delta(\mathbf{r} - \mathbf{r}') \tag{1}$$

$$\rho(\mathbf{r}) = -\frac{1}{\pi\Omega} \int_\Omega d\mathbf{r} \int_{-\infty}^{\infty} d\epsilon \, G(\mathbf{r},\mathbf{r};\epsilon) f(\epsilon - \mu) \tag{2}$$

$$V_{eff} = V_{ext} + e^2 \int d\mathbf{r}' \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} + \frac{\delta E_{xc}}{\delta\rho(\mathbf{r})} \tag{3}$$

The SCF algorithm proceeds as follows: an input guess $\rho^{in}(\mathbf{r})$ of the electronic charge is used in eq. 3 to construct an initial guess of the effective electron-ion potential $V_{eff}(\mathbf{r})$ appearing in the Schrödinger equation (eq. 1). The single particle Schrödinger equation is then solved to obtain the single particle wave functions or, as indicated in eq. 1 and as is appropriate to the multiple scattering Green's function methods used in this paper, the single particle Green's function, $G(\mathbf{r},\mathbf{r};\epsilon)$. Next, the output charge density $\rho^{out}(\mathbf{r})$ is obtained from eq. 2. Finally, a new $\rho^{in}(\mathbf{r})$ is constructed by mixing the original $\rho^{in}(\mathbf{r})$ with $\rho^{out}(\mathbf{r})$ according to some (linear or non-linear) mixing algorithm. This procedure is then iterated until $\rho^{in}(\mathbf{r})$ and $\rho^{out}(\mathbf{r})$ agree to a prescribed small tolerance. Once a SCF solution is found the total energy is obtained from:

$$E_{tot}[\rho] = E_{band} - \int d\mathbf{r}\rho(\mathbf{r})V_{eff}(\mathbf{r}) + E_{Coulomb}[\rho] + E_{xc}[\rho] \tag{4}$$

Clearly, LDA is a first principles, parameter free, theory of matter in that the only inputs are the atomic numbers of the constituent elements.

2.1.2 The LSMS Method for Non-Collinear Magnetism

Standard reciprocal space ($\mathbf{k}$-space) electronic structure methods for solving the LDA single particle Schrödinger equation (eq. 1) ultimately scale as the third power of the number of atoms, N, in the unit cell. This is due to the need to diagonalize, invert, or orthogonalize a matrix whose size is proportional to N. For the very large cells that are required for implementing SD this scaling is prohibitive. The LSMS method [5] avoids this $O(N^3)$ scaling by using real space multiple scattering theory to obtain the Green's function (eq. 1) and hence the electronic charge (eq. 2) and magnetization densities. In the LSMS method the atoms comprising the crystal are treated as disjoint electron scatterers. The multiple scattering equations are solved for each atom by ignoring multiple scattering processes outside a local interaction zone (LIZ) centered on each atom (see figure 2).

The potential throughout the whole space is reconstructed in the standard way by solving Poisson's equation (eq. 3) for a crystal electron density made up of a sum of the single site densities. For the LSMS method, the LIZ size is the central convergence control parameter. For face and body centered cubic (fcc and bcc) transition metals LIZ sizes of a few ($\sim 3 - 7$) neighboring atom shells are sufficient to obtain convergence of the charge and magnetization densities. Using a Cray T3E, quite precise $O(N)$ scaling of the basic LSMS method has been demonstrated previously for up to 1024 atoms (nodes) [6].

The details of the LSMS method have been described in the literature [5]. The extension to cover the case of non-collinear moment arrangements is discussed in references [3] and [4]. Here let us simply quote the two most salient formulae necessary to understand the basic algorithm. The first is the Green's
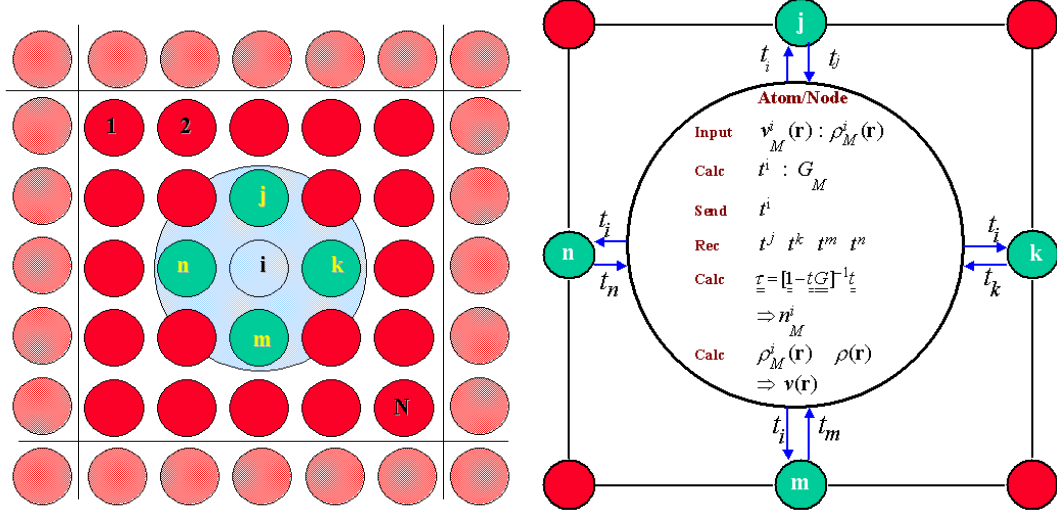
Figure 2: *Left* Schematic of a parallel SCF algorithm. We envision a basic unit cell consisting of N-atoms labeled $1, \ldots, N$ which is periodically reproduced as indicated. About each atom, for example the i^{th}, we define a local interaction zone consisting of M-atoms including itself, for example labeled i, j, k, l, m. *Right* Schematic of the implementation LSMS algorithm on a massively parallel computer.

function for the atom at the center of the LIZ, $\hat{\mathcal{G}}_M(\{\mathbf{e}\}_i; \mathbf{r}; \epsilon)$, which is given by

$$
\begin{aligned}
\hat{\mathcal{G}}_M(\{\mathbf{e}\}_i; \mathbf{r}; \epsilon) &= \sum_{LL'} \hat{Z}_L^i(\mathbf{r}; \epsilon) \left[\hat{\underline{\tau}}_M(\{\mathbf{e}\}_i; \epsilon)\right]_{LL'}^{ii} \hat{Z}_{L'}^i(\mathbf{r}; \epsilon) \\
&\quad - \sum_L \hat{Z}_L^i(\mathbf{r}; \epsilon) \hat{J}_L^i(\mathbf{r}; \epsilon),
\end{aligned}
\tag{5}
$$

The second is the scattering path matrix, $\hat{\underline{\tau}}_M(\{\mathbf{e}\}_i; \epsilon)$, which describes the multiple scattering effects within the LIZ. It is given by the inverse of the KKR matrix, $\mathbf{M}$, through

$$
\hat{\underline{\tau}} = \hat{\underline{\mathbf{M}}}^{-1}
\tag{6}
$$

$$
\hat{\underline{\mathbf{M}}} = \left[\hat{\underline{\mathbf{t}}}^{-1} - \hat{\underline{\mathbf{G}}} \right].
\tag{7}
$$

In the last equation $\hat{\underline{\mathbf{t}}}$ is the t-matrix which describes the scattering from a single site and $\underline{\mathbf{G}}$ is the structure constant matrix which describes the propagation of electrons from site to site. The above are matrices in site (i), angular momentum (L) and spin (denoted by ˆ) indices. This gives a linear dimension of $M \times L \times 2$, where M is the number of atoms in the LIZ (87 in the present work), L is a combined angular and azimuthal quantum number (16 for atoms near the center of the LIZ and 9 for the more distant ones), and the factor of 2 arises from the spin index. Not surprisingly, inversion of the KKR matrix (eq. 6) to obtain the scattering path matrix is the most time consuming task in the LSMS method.

2.1.3   Algorithm of *ab initio* spin dynamics.

A general orientational configuration of atomic magnetic moments $\{\mathbf{e}_i\}$ such as is encountered during an individual time step of spin-dynamics, is not a ground state of LSDA. In order to properly describe

the energetics of a general configuration it is necessary to introduce a constraint which maintains the orientational configuration. The purpose of the constraining fields is to force the local magnetization, $\mathbf{M}_i$, to point along the prescribed local direction $\mathbf{e}_i$. Note that at each site the constraining field $(\vec{B}_{cons})$ is transverse to the magnetization. The constraining field is not known *a priori*, and has to be calculated by some additional algorithm. The details of the theory and the constrained local moment (CLM) model are contained in reference [4]. The CLM state may be considered as a time freeze during the evolution of the system from a general initial state to a final state. This final state may be the ground state, as most often is, or another non-equilibrium state. We apply a simple model to describe the time-evolution process of the system between such CLM states: if the timestep is very small, the magnetic moments can be considered as classical moments of constant size. The relaxation of such moments to the ground state is then described by a Landau-Lifshitz like equation of motion of the form

$$
\begin{aligned}
\frac{d}{dt}\mathbf{e}_i &= \alpha\mathbf{e}_i \times (\mathbf{e}_i \times \mathbf{B}_{eff}) \\
\vec{B}_{eff} &= -\vec{B}_{cons}.
\end{aligned}
\tag{8}
$$

This simple first order differential equation needs to be solved for each moment. This gives the orientational configuration for the new timestep. A new constrained LDA calculation is then performed. This is repeated until a ground state is obtained in which $\mathbf{B}_{cons} = 0$.

2.2 Parallel Challenges

By making an atom (or atoms) to node (or CPU in the case of a machine that has multiple CPU's per node) equivalence the LSMS is highly scalable on MPPs since each node can be assigned the work involved in setting up and inverting the KKR matrix eq. (6) for the atom(s) to which it is assigned. Inversion of the KKR matrix is the most time consuming step in the LSMS method. It typically takes over 80% of the total run time and is implemented in terms of complex matrix-matrix multiplies using the highly optimized BLAS routine CGEMM.

Keeping angular momenta up to l_{max} (typically $l_{max} = 3$) the dimension of the matrix that is inverted is $2 \times (l_{max} + 1)^2 \times M$, where M is the number of atoms in the LIZ and the factor of 2 comes from the spin-index. Calculation of the charge and magnetization densities requires that each node invert this matrix for a set of energies along a contour in the complex energy plane of sufficient density to converge the necessary energy integrations. If M is smaller than the number of atoms in the system, N, this is clearly advantageous compared to conventional reciprocal space methods since these involve the inversion of a matrix of dimension $2 \times (l_{max} + 1)^2 \times N$ at a sufficient number of $\mathbf{k}$-points (a few for large systems to a few hundred for small systems) to converge necessary Brillouin zone integrals. The details of how the LSMS algorithm maps onto a parallel machine is shown in figure 2 (left) and explained in detail in the literature[5].

Necessary message passing involves passing of single site t-matrices for the sites/nodes in a given atom's LIZ, as well as global information regarding the total charge in the system, charge transfer and potential multipole moments. A typical run spends more than 80% of the time in BLAS3 routines and only a few percent in communications. This is a direct result of the complete rethinking of the electronic structure problem in terms of a parallel algorithm. This allows the algorithm to scale up to very large numbers of processors with the percentage of time spent in communications staying approximately constant when the number of processors increases with the number of atoms.

2.3 Special Optimization Issues: Fast Algorithms for Matrix Inversion

In the LSMS method each node calculates the scattering path matrix for its own atom(s). Hence, only the $L \times L$ ($L = 2(l_{max} + 1)^2$) upper left block of the inverse is required, which represents the local atom's contribution to the scattering path matrix. A significant enhancement in computational performance can be achieved by taking advantage of this fact. Furthermore, the design of the inversion algorithms can be largely based on BLAS level 3 matrix-matrix multiply routines to ensure optimal performance.

The size of the KKR matrix in the LSMS approach for an LIZ size of M atoms, is $Q = ML$, where $L = 2(l_{\max} + 1)^2$ for each atom. Here for simplicity we assume that all atoms in the LIZ are calculated to the same $l_{\max}$. The standard approach for the dense matrix inversion would be to apply the LU factorization of the whole matrix, which takes $\frac{1}{3}Q^3$ operations (one operation is one multiply and one add), or $\frac{2}{3}Q^3$ floating point operations, plus the forward and backward substitution for the first L columns, which take $Q^2 L$ operations or $2Q^2 L$ floating point operations, for a total of

$$Q^2(\frac{2}{3}Q + 2L) = \frac{2}{3}M^3 L^3 + 2M^2 L^3 \qquad (9)$$

floating point operations.

Recognizing that not the entire first L columns of the inverse are needed, we searched for a more efficient algorithm that could greatly reduce the total number of floating point operations for the calculation. The approach that we take is a variant of Strassen's [10] technique due to X.-G. Zhang [9]. If we partition a matrix of dimension Q into blocks of Q_1 and Q_2, $Q_1 + Q_2 = Q$, in the form,

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}, \qquad (10)$$

then the top $L \times L$ diagonal block of the inverse of this matrix is also the top $L \times L$ diagonal block of the inverse of the following matrix,

$$A' = A - BD^{-1}C. \qquad (11)$$

Thus we performed a transformation that reduced the problem from $Q \times Q$ to $Q_1 \times Q_1$. The number of operations (one multiply and one add) involved in this transformation is obviously $\frac{1}{3}Q_2^3 + Q_2^2 Q_1 + Q_2 Q_1^2$, which is smaller for smaller values of Q_2. However, Q_2 cannot be too small, otherwise the performance of the computer can degrade considerably due to cache misses. A reasonable choice is $Q_2 = L$ where L is typically 32 for our calculations. Thus we apply Eq. (11) recursively, until the size of the resulting A' is $L \times L$, at which time a simple inversion yields the desired upper $L \times L$ diagonal block of the inverse of the KKR matrix. The total number of floating point operations is,

$$2\sum_{m=1}^{M} \left\{ \frac{1}{3}L^3 + L^2[(M-m)L] + L[(M-m)L]^2 \right\} = \frac{2}{3}M^3 L^3 - 2M^2 L^3 + 2ML^3. \qquad (12)$$

The leading term of this result is the same as the standard LU algorithm, Eq. (9). However, the difference in the lower order terms of M can significantly reduce the number of operations for small M. Furthermore, since the new algorithm spends a very small fraction of the operations in inversion, ($\frac{2}{3}ML^3$ floating point operations), the majority of the operations are in matrix-matrix multiplies, it is much more easily optimized than the standard LU algorithm. For example, for a typical calculation where $L = 32$ and $M \approx 36$, we find that the new algorithm is about a factor of three faster than the standard LU on a wide range of platforms although the speedups are less significant for larger values of M.

3   Software Configuration

The Spin Dynamics (SD) version of the LSMS code is written in FORTRAN90 and is approximately 45,000 lines. The SD-LSMS code is derived from the CLM-LSMS code that won the 1998 Gordon Bell. The CLM-LSMS code for performing first principles constrained local moment (CLM) calculations is subsumed with the spin dynamics loop of the present code. The essential point being that the current generation of machines is sufficiently powerful to allow integration of the SD equation of motion over many time steps. In the present case SD time stepping is utilized to allow 'annealing' of an initial model configuration into the ground state. The CLM model is used to calculate the magnetic moments, energy, and constraining fields for the current orientational configuration. The SD equation of motion is then used to evolve the orientational configuration using the negative of the constraining field as the driving field.

In the current code message passing, which was previously distributed throughout the code and activated using a precompiler, according to the message passing interface being used - originally NX and PVM - is implemented through an interface of our own design. The specifics of message passing are relegated to a FORTRAN90 message-passing module and are implemented in both MPI and PVM. The consequence is that use of a precompiler is confined to the message-passing module and the rest of the source code does not depend on the message passing library used. Both the MPI and PVM versions of the code are used depending on the platform being used. Typically, PVM is used on work-station clusters and MPI is used on Cray T3E, IBM-SP, and Compaq MPP machines.

4   Results

4.1   Performance

In addition to performing the specific science calculations reported in the next sub-section, we performed a number of calculations on large cell models of disordered FeMn alloys design to study the performance of the SP3 within a real application context. These were part of a separate project to study the ground state magnetic structure and total energy of face centered cubic (fcc) disordered solid solution FeMn alloys.

The first calculations of this type were done on a 184-node IBM RS/6000 SP3 (Eagle) operated by the Center for Computational Sciences (CCS) at Oak Ridge National Laboratory (ORNL). Each "Winterhawk-II" node has four 375 MHz Power3-II processors and 2GB of memory. The peak computational power of this machine is over 1 TeraFLOPS in the compute partition of 176 nodes. We initially performed 32-atom per unit cell simulations of FeMn on 32 nodes using a 43-atom LIZ and using spin dynamics to relax the local moments in order to determine the true magnetic state. To perform the larger scale simulations the LSMS code was ported to Falcon which is a 64-node Compaq AlphaServer SC also located in CCS. Falcon has four 667 MHz Alpha EV67 processors and 2GB of memory per node and 2 TB of Fiber Channel disk attached, with an estimated computational power of 342 GFLOPS. Simulations were performed on 256-atom per unit cell configurations using an 87-atom LIZ and using spin dynamics to relax the local moments in order to determine the true magnetic state.

The specific performance runs documented here were for large cell models of FeMn alloys made up of NxNxN repeats of the underlying face centered cubic (fcc) lattice with the 3Q alignment of magnetic moments. Since the fcc cell contains 4-atoms, cells having N=1,2,3,4,5,6,7, and 8 correspond to system sizes of 32, 108, 256, 500, 864, 1372, and 2048-atoms respectively. In addition a calculation for a 4x8x17 system corresponding to 2176 atoms was performed to obtain performance figures for the full 136-node compute partition of the NERSC IBM-SP3. The NERSC IBM-SP has a total of 158 Nighthawk II SMP

nodes, each with 16, 375 MHz Power3+ processor and a peak speed of 3.26 Tflops for the 136-node compute partition. The performance runs comprised of 10 self-consistent field (SCF) iterations towards convergence of the CLM state of one SD-time step. This is sufficient to amortize the start-up and shut-down (or check pointing) cost of the code and is representative of how the code is used in a scientific simulation. The start-up and shut-down time increases slowly with system size but still only accounts for $\approx$ 60 seconds of a $\approx$ 4000 second run.

The LIZ was held fixed at 7 near-neighbor (nn) shells (87-sites) and $l_{max} = 3$ was used for the first two nn-shell and $l_{max} = 2$ for the rest which results in a KKR matrix dimension of 876$\times$2. This LIZ size is sufficient to reliably resolve the ground state configuration.

In order to establish the operations count, the N=2 (32-atom) system was run on 32 processors of the Cray T3E900 at NERSC, using the performance analysis tool PAT. The IBM-SP does not have a generally available parallel hardware performance monitor. The atom to processor mapping of the SD-LSMS code results in each node executing (approximately) the same number of operations independent of system size. Analysis of the PAT results showed a very small node to node variation in the instruction count (less than 0.1%). Thus, to estimate the SP3 performance, we used the T3E flop count averaged over the 32-processors as the per processor flop count for the IBM SP3 runs. For the 10-iteration run the average flop count on the T3E was 4.672×10^6 flop/processor(atom) including all start-up and shut-down operations. In the LSMS algorithm as the number of atoms in the supercell increases the number of operations per atom stays constant so it is simple to calculate the flop count for any system size. Internal timings embedded within the code were used to obtain execution times on the IBM SP3.
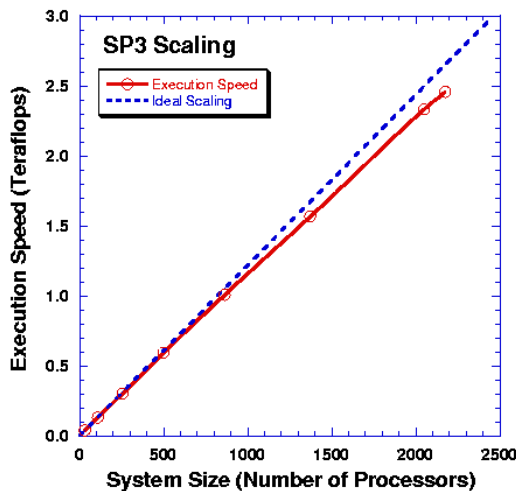

Figure 3: Execution rate of SD-LSMS code. The measured execution rates are shown as open circles and the solid line is a guide to the eye. The ideal linear scaling line (dashed line) is based on straight-line extrapolation of the 32-atom (processor) results. The performance runs were carried out during normal production cycles with an unknown mix of other jobs running of the machine

The measured execution rates are shown in figure 3 and are compared with ideal linear scale up from the 32-processor timings. For the largest system (2176-processors) the code ran at 2.46 Teraflops. For the 2016-atom production runs used to study the FeMn/Co interfaces the code was running in excess of 2.25 Teraflops. The high flop speeds are obtained in part by using the IBM native BLAS3 (matrix-matrix multiply) routines which run close to the peak performance of the Power3 chip for sufficiently

large matrices. For very large systems a slight deviation from ideal linear scaling is observed. In each SCF loop the KKR-matrix is inverted on a mesh of 31 energies in the complex energy plane needed to perform an energy integration. The time to perform an inverse is essentially constant, independent of system size. However, the time to perform a SCF loop gradually increases from ≈ 380 sec. for the 32-atom system to ≈ 430 sec. for the 2176-atom system. Thus, the additional time results from the time required to construct the KKR-matrix. Since the size of the matrix is again constant, independent of system size, the major difference between these runs is the total number of messages being passed during the set up process. Therefore the time increase can be ascribed to the increasing cost of message passing of the single site t-matrices during the set up of the KKR matrix as the number of atoms/processors increases (see figure 2).

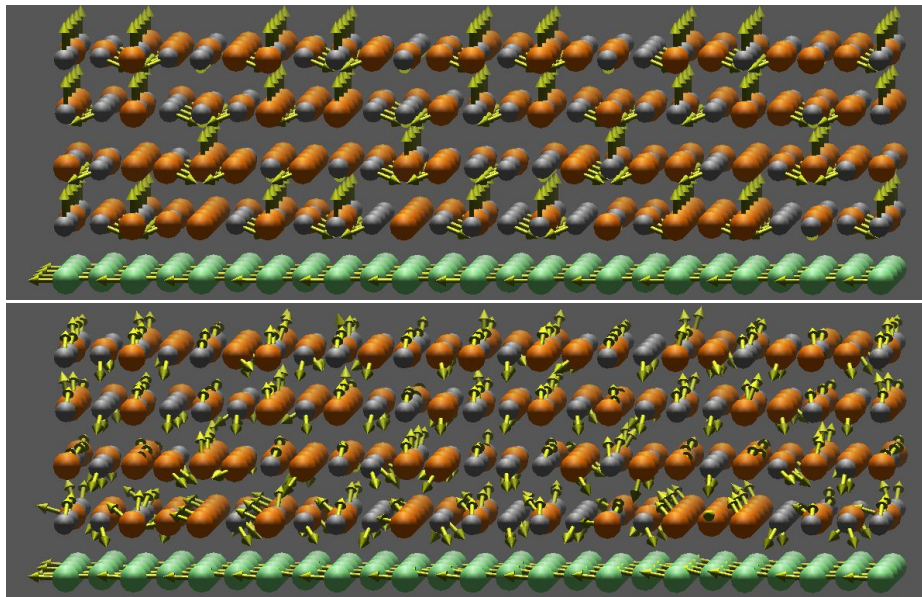## 4.2 Magnetic Structure of FeMn/Co Interfaces

Figure 4: (Upper) Section of the FeMn/Co interface showing the initial configuration of the magnetic moments. The visualization shows the Co layer and four FeMn layers closest to the interface. The Co layer is initialized in a Fm state and the FeMn layers are initialized in the 3Q structure. (Lower) Section of the FeMn/Co interface showing the final configuration of the magnetic moments for the same five layers as shown for the initial configuration.

The main science results of this paper are for the ground state magnetic structure of the FeMn/Co interface shown in figure 1 and discussed in the introduction.

In figure 4-lower we show a section of the final (ground state) interface magnetic structure. For comparison we show in figure 4-upper the same layers in the initial moment configuration (extracted from figure 1). Clearly, major changes have occurred in the FeMn-alloy layers. The magnetic moments have rotated and evolved from the 3Q-structure into a new magnetic structure. This evolution results from the interaction of the FeMn-layer with the Co. During the simulation the orientation of the magnetic moments of the lowest three layers of Co (see figure 1) was held fixed while all other Co and FeMn layers were free to rotate according to SD. Evolution of the structure involves a CLM calculation at each time step. For the initial configuration the local constraining fields, needed to hold the magnetic moments

in the prescribed configuration, are large. This is indicative of the fact that the initial state (shown in figure 1) is far from the equilibrium state. In the final state the constraining fields are reduced by several orders of magnitude and the total energy is lower in energy by some 2 mRy/atom, indicating evolution towards the ground state configuration.

Here, it is not our purpose to make a detailed analysis of the results of this calculation. We simply note that the FeMn layers are no longer in the 3Q structure (even allowing for a possible global rotation of the 3Q structure). In addition there is a 'random' component to the orientation of the Fe and Mn magnetic moments about some predominant directions. There has also been a small rotation of the moments in the uppermost Co layer relative to their original position. Analysis of the details of the magnetic structure, which appears to be very rich and could not have been predicted without the full machinery of first principles SD and the use of the massively parallel resources, is proceeding and when complete will be published elsewhere.

Acknowledgments

References

[1] V. P. Antropov et al. Phys. Rev. **54**, 1019 (1996).

[2] Q. Niu, Xindong Wang, L. Kleinman, Wu-Ming Liu, D. M. C. Nicholson, and G. M. Stocks, Phys. Rev. Letters, **83**, 207 (1999).

[3] G. M. Stocks, B. Ujfalussy, Xindong Wang, Xiaoguang Zhang, D. M. C. Nicholson, W. A. Shelton, Yang Wang, A. Canning, and B. L. Gyorffy, Phil. Mag. B **78**, 665, (1998).

[4] B. Ujfalussy, X. D Wang, D. M. C. Nicholson, W. A. Shelton, G. M. Stocks, Y. Wang, B. L. Gyorffy, J. Appl. Phys. 85, 4824 (1999)

[5] Y. Wang, G.M. Stocks, W.A. Shelton and D.M.C. Nicholson, Phy. Rev. Letters **75**, 2867 (1995).

[6] *High Performance First Principles Method for Complex Magnetic Properties*, Ujfalussy, B., Xindong Wang, Xiaoguang Zhang, D. M. C. Nicholson, W. A. Shelton, G. M. Stocks, A. Canning, Yang Wang, and B. L. Gyorffy, Proceeding of the ACM/IEEE SC98 Conference, Orlando, Florida, November 7-13, (1998). IEEE Computer Society, Los Alamitos, CA 90720-1264, CD-ROM.). First Place Gordon Bell Prize at SC98 Orlando. Proc. of SC98, Orlando, Florida, (1998).

[7] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964), W. Kohn and L. J. Sham, Phys. Rev. **140**, A1133 (1965).

[8] J. S. Faulkner and G. M. Stocks, Phys. Rev. **B21**, 3222 (1980).

[9] X.-G. Zhang, Ph.D. thesis, Department of Physics, Northwestern University (1989).

[10] V. Strassen, Numer. Math., 13, 354 (1969).