

Thermodynamics of Magnetic Systems from First Principles: WL-LSMS

M. Eisenbach¹, C.-G. Zhou^{1,*}, D. M. Nicholson¹, G. Brown², J. Larkin³, and T. C. Schulthess⁴

¹Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA

²Department of Physics, Florida State University, Tallahassee, Florida 32306, USA

³Cray Inc., Oak Ridge, Tennessee 37831, USA

⁴Institute for Theoretical Physics and Swiss National Supercomputer Center, ETH Zürich, 8903 Zürich, Switzerland

ABSTRACT: Density Functional calculations have proven to be a powerful tool to study the ground state of many materials. For finite temperatures the situation is less ideal and one is often forced to rely on models with parameters either fitted to zero temperature first principles calculations or experimental results. This approach is especially unsatisfactory in inhomogeneous systems, nano particles, or other systems where the model parameters could vary significantly from one site to another. Here we describe a possible solution to this problem by combining classical Monte Carlo calculations – the Wang-Landau method[2] in this case – with a first principles electronic structure calculation, specifically our locally self-consistent multiple scattering code (LSMS)[3]. The combined code shows superb scaling behavior on massively parallel computers. The code sustained 1.836 Petaflop/s on 223,232 cores of the Cray XT5 jaguar system at Oak Ridge.

KEYWORDS: Magnetism, Monte-Carlo, First Principles, Matrix Inversion

1 Introduction

Density Functional based first principles electronic structure calculations for condensed matter systems have reached a high level of maturity over the last few decades and are now a standard tool for the study of ground state material properties.[1] While these methods have evolved to provide greater accuracy and deal with wider classes of materials by developing new approximations to the exchange correlation functional, the important field of finite temperature behavior has received less attention. The phase

space of interesting systems is far too large to be dealt with directly. The usual methods of treating the thermodynamics of a physical system involve either the time evolution of an ensemble or the exploration of the most relevant parts of phase space by means of a Monte-Carlo method. Both these approaches require a large number of evaluations of the underlying Hamiltonian that describes the system ($> O(10^5)$), thus it is usually only feasible to treat severely simplified models that have to be designed to capture the essential physics, as opposed to a direct treatment of the Density-Functional Hamiltonian of the system.

To overcome this limitation we have developed the hybrid Wang-Landau/LSMS (WL-LSMS) code.

*Present address: Quantitative Research, J. P. Morgan Chase and Company, New York, New York 10017, USA

This code combines recent advances in computational statistical mechanics, namely the Wang-Landau method [2], that will allow the most efficient use of information obtained from individual evaluations of the system’s Hamiltonian, with the LSMS first principles that has already demonstrated superb scalability on massively parallel machines.

In the following sections we will first give a brief overview of the overall structure of the hybrid code, then we will describe the Wang-Landau and LSMS methods. Following this we will describe the ideas that enable the high fraction of theoretical peak performance that can be obtained with the block inversion kernel of LSMS that accounts for $\approx 95\%$ of the total execution time of the code. Finally we will present performance results for the hybrid WL/LSMS code and conclude.

2 Structure of WL-LSMS

The WL-LSMS code uses a hybrid parallelization scheme. At the top level, the code parallelizes over concurrent random walkers, where we use a master-slave scheme, with a master that accumulates the density of states of the system, and the slaves that execute the random walks, each running its own instance of the LSMS method. The second parallelization level is the LSMS portion of the code, where domain decomposition is used with one atom per processing core. In typical production runs, the WL method would use a hundred to a few thousand concurrent walkers, and the LSMS portion would be parallelized over up to a few thousand processing cores. The method hence will scale to hundred thousand or millions of processing cores. The schematics of the parallelization structure are shown in fig. 1.

3 The Wang-Landau Algorithm

All thermodynamic potentials can be derived from the partition function

$$Z(T) = \int e^{-E(\mathbf{X})/(k_B T)} d\mathbf{X} \quad (1)$$

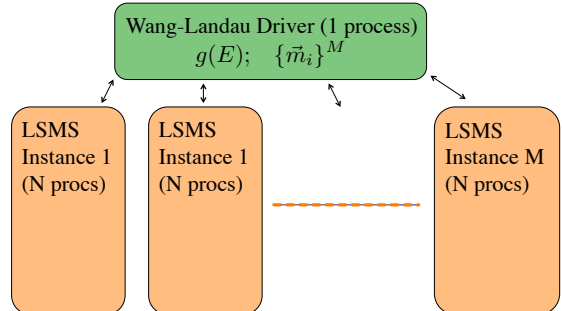


Figure 1: Parallelization strategy of the combined Wang-Landau/LSMS algorithm. The Wang-Landau process (Alg. 1) generates random spin configurations for M walkers and updates a single density of states $g(E)$. The energies for these N atom systems are calculated by independent LSMS processes (Fig. 2). This results in two levels of communication, between the Wang-Landau driver and the LSMS instances, and the internal communication inside the individual LSMS instances spanning N processes each.

where $E(\mathbf{X})$ is the internal energy of the system with the phase space described by the variable \mathbf{X} in some high dimensional space consisting of all the microscopic degrees of freedom of the system (atomic positions, velocities, and/or magnetic moments). In importance sampling Monte Carlo simulations one performs a random walk through phase space that is biased in such a way that the walker spends most of the time where the integrand in equation 1 is largest, that is, where the energy $E(\mathbf{X})$ is small. In the Metropolis algorithm, this is accomplished by accepting a proposed move from point \mathbf{X}_i to \mathbf{X}_{i+1} with the probability

$$\min[1, \exp(-k_B T [E(\mathbf{X}_{i+1}) - E(\mathbf{X}_i)])]. \quad (2)$$

This results in a very efficient computation of the partition function at a particular temperature when the energy function is not too complex. In many cases that are common in nano-science and biology, the energy function can have many local minima that are separated by large energy barriers. With the conven-

tional Metropolis algorithm, the random walk can be trapped for very long time around local minima, and sampling representative parts of phase space around other local minima of the energy function can become exceedingly difficult.

The partition function in equation 1 can be rewritten in the form

$$Z(T) = \int g(E) e^{-E/(k_B T)} dE \quad (3)$$

where the density of states is defined as

$$g(E) = \int \delta(E - E(\mathbf{X})) dX \quad (4)$$

and $\delta(E)$ is the Dirac δ -function.

Flat histogram methods, such as the Wang-Landau algorithm, use the density of states, $g(E)$, for importance sampling. The analog of the Metropolis formula in this case is to accept the new configuration with probability

$$\min[1, g(E_i)/g(E_{i+1})]. \quad (5)$$

The effect is to create an equal probability of visiting each energy level in the system. In other words, a histogram of where the random walk is at the end of each Monte Carlo move would be essentially flat (save for statistical $1/\sqrt{(N)}$ noise).

The main obstacle of flat-histogram methods is that $g(E)$ is not known. Instead, an estimate of the density of states $\tilde{g}(E)$ must be constructed self-consistently as the Monte Carlo estimate is generated. The Wang-Landau algorithm accomplishes this as follows. It begins with a prior estimate of the density of states, $\tilde{g}_0(E)$, which might be just a constant. Assuming that a Monte Carlo move to a new configuration with Energy E_{i+1} is accepted according to the criterion of equation 5, the density of states is updated with

$$\ln[\tilde{g}(E_{i+1})] \leftarrow \ln[\tilde{g}(E_{i+1})] + \ln f \quad (6)$$

where f is the modification factor that is initially set to $\ln f = 1$. For every accepted move, a histogram $H(E)$ is updated, which records where the random

walk has been. The estimate $\tilde{g}(E)$ is considered converged, when

$$\min[H(E)] \leq A \text{ mean}[H(E)], \quad (7)$$

where the flatness parameter $0 < A < 1$ controls the accuracy of the estimated $\tilde{g}(E)$, with increasing accuracy as A approaches unity. When this criterion is met, the modification factor is reduced such that $\ln(f) \leftarrow \ln(f)/2$ and the histogram reset. For each iteration in f , a new visit histogram $H(E)$ is calculated, and the process is repeated until $\ln f \leq 1 \times 10^{-6}$ or smaller, and the estimated $\tilde{g}(E)$ is considered the converged density of states.

Since the systems we set out to study here have continuous degrees of freedom, $g(E)$ is a function of a continuous variable as well. Hence, when the random walk arrives at a particular configuration with energy E_i , the estimated density of state $\tilde{g}(E)$ is updated by [8]

$$\tilde{g}(E) \rightarrow \tilde{g}(E) \times f^{k((E-E_i)/\delta)} \quad (8)$$

where $k(x) \geq 0$ is a continuous function with compact support. In particular we choose an Epanechnikov kernel, $k(x) = \max[0, 1 - x^2]$, and the width δ to be 2% of the difference between the minimal (ferromagnetic) and maximal (anti-ferromagnetic) energies of the system.

4 The LSMS Algorithm

For the energy evaluation, we employ the first principles framework of density functional theory (DFT) in the local spin density approximation (LSDA). To solve the Kohn-Sham equations arising in this context, we use a real space implementation of the multiple scattering formalism. The details of this method for calculating the Green function and the total ground state energy $E[n(\vec{r}), m(\vec{r})]$ are described elsewhere [3, 4]. For the present discussion it is important to note that the computationally most intensive part is the calculation of the scattering path matrix τ for each atom in the system by inverting the multiple scattering matrix.

$$\tau = [I - tG_0]^{-1} t$$

Algorithm 1 Wang-Landau/LSMS algorithm

- 1: initialize logarithmic density of states $\ln g(E) \leftarrow 0$, histogram $h(E) \leftarrow 0$, modification factor $\gamma \leftarrow 1$, and the set of magnetic moment directions for the M walkers $\{\hat{e}\}_{1\dots M}$
 - 2: **repeat**
 - 3: submit new random moment directions $\{\hat{e}^{\text{new}}\}$ to idle LSMS instances
 - 4: receive new energy E_n^{new} from walker n
 - 5: accept new set of directions $\{\hat{e}^{\text{new}}\}_n$ with probability $\min[1, g(E_n^{\text{old}})/g(E_n^{\text{new}})]$
 - 6: **if** Move accepted **then**
 - 7: $\{\hat{e}^{\text{old}}\}_n \leftarrow \{\hat{e}^{\text{new}}\}_n$
 - 8: **end if**
 - 9: update density of states $\ln g(E_n) \leftarrow \ln g(E_n) + \gamma$ and histogram $h(E_n) \leftarrow h(E_n) + 1$
 - 10: **if** $h(E)$ flat **then**
 - 11: $\gamma \leftarrow \gamma/2$, $h(E) \leftarrow 0$
 - 12: **end if**
 - 13: **until** $g(E)$ converged, *i.e.* $\gamma \approx 0$
-

The only part of τ that will be required in the subsequent calculation of site diagonal observables (*i.e.* magnetic moments, charge densities, and total energy) is a small (typically 32×32) diagonal block of this matrix whose rank is $O(4k)$. This will allow us to employ the algorithm described in the next section for maximum utilization of the on node floating point compute capabilities.

Most importantly for the application in the hybrid Wang-Landau LSMS method, our Locally Self-consistent Multiple Scattering (LSMS) method allows the possibility of non-collinear magnetism [5].

The orientation \hat{e}^i of the magnetic moment for each site is determined by

$$\hat{e}^i = \int_{\Omega^i} d\vec{r} \vec{m}^i(\vec{r}) / \left| \int_{\Omega^i} d\vec{r} \vec{m}^i(\vec{r}) \right|.$$

Since an arbitrary arrangement is not a DFT ground state we will have to deal with a constrained general state as presented by Stocks *et al* [6, 7]. In the constrained local moment (CLM) model the LSDA

equations are solved subject to a constraint

$$\int_{\Omega_i} \vec{m}_i(\vec{r}) \times \vec{e}_i d\vec{r} = 0 \quad (9)$$

that ensures that the local magnetizations lie along the directions prescribed by $\{\vec{e}_i\}$. The result is that, in order to maintain the specific orientational configuration, a local transverse constraining field must be applied at each site. The constraining field is obtained from the condition

$$\frac{\delta E^{\text{con}}[\{\vec{e}_i\}, \{\vec{B}_i^{\text{con}}\}]}{\delta \vec{e}_i} = 0 \quad (10)$$

applied to all sites and where E the generalized energy functional in the presence of the constraining field. Thus this method enables the calculation of the energies of arbitrary orientational states as generated by the Wang-Landau algorithm.

The Locally Self-consistent Multiple Scattering (LSMS) method calculates the electronic properties from first principles in real space, but introduces some approximations that make the treatment of infinite systems possible. Furthermore this method results in a code that scales linearly with the size of the system.

The LSMS method is based on the observation that good convergence can be obtained by solving the Kohn-Sham equation of density functional theory at a given atomic site by considering not the whole system, but only a sufficiently large neighborhood, the local interaction zone (LIZ), of each site.

5 Blockinversion in LSMS

The most time consuming part of the LSMS calculation is the inversion of the multiple scattering matrix. The amount of computational effort can be reduced by utilizing the fact that for each local interaction zone only the left upper block (τ_{00}) of the scattering path matrix τ is required. In this section we describe an algorithm that reduces the amount of work needed while providing excellent performance due to its reliance on dense matrix-matrix multiplications that are available in highly optimized form in vendor

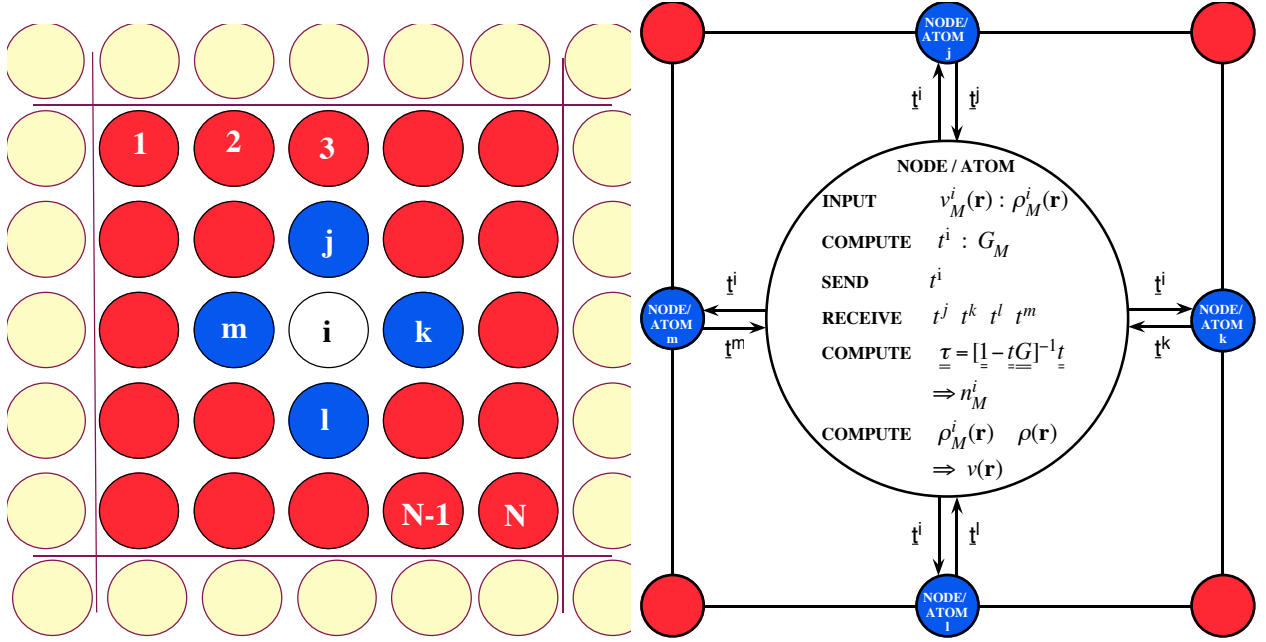


Figure 2: Schematic, Left: LIZ centered at processor/atom i ; Right: message passing and computation.

or third party provided implementations (*i.e.* ZGEMM in the BLAS library).

The method employed in LSMS to calculate the required block of the inverse relies on the well known expression for writing the invers of a matrix in term of inverses and products of subblocks:

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}^{-1} = \begin{pmatrix} U & V \\ W & Y \end{pmatrix}$$

where

$$\begin{aligned} U &= (A - BD^{-1}C)^{-1} \\ V &= -(A - BD^{-1}C)^{-1}BD^{-1} \\ W &= -D^{-1}C(A - BD^{-1}C)^{-1} \\ Y &= D^{-1} + D^{-1}C(A - BD^{-1}C)^{-1}BD^{-1} \end{aligned}$$

This this method can be applied multiple times to the subblock U until the desired block τ_{00} of the scattering path matrix is obtained.

To illustrate this algorithm we present the detailed steps of a calculation for a matrix that has been subdivided into 3×3 subblocks.

1. starting matrix

$$\begin{pmatrix} A & B & C \\ D & E & F \\ G & H & J \end{pmatrix}$$

2. LU factorization of J and calculation of product of its invers with the bottom row.

$$\begin{pmatrix} A & B & C \\ D & E & F \\ G' & H' & J' \end{pmatrix}$$

where

$$\begin{aligned} J' &= \text{LU factorization of } J \\ G' &= J^{-1}G \\ H' &= J^{-1}H \end{aligned}$$

3. multiplication of the rightmost column of the next row with the botom row.

$$\begin{pmatrix} A & B & C \\ D' & E' & F \\ G' & H' & J' \end{pmatrix}$$

where

$$\begin{aligned} D' &= D - FJ^{-1}G \\ E' &= E - FJ^{-1}H \end{aligned}$$

4. multiply top row of the rightmost column with the bottom row of the middle column.

$$\begin{pmatrix} A & B' & C \\ D' & E' & F \\ G' & H' & J' \end{pmatrix}$$

where

$$B' = B - CJ^{-1}H$$

5. LU factorization of E' and calculation of product of its invers with the rest of the middle row.

$$\begin{pmatrix} A & B' & C \\ D'' & E'' & F \\ G' & H' & J' \end{pmatrix}$$

where

$$\begin{aligned} E'' &= \text{LU factorization of } E' \\ D'' &= E'^{-1}D' = (E - FJ^{-1}H)^{-1} \\ &\quad (D - FJ^{-1}G) \end{aligned}$$

6. Multiply the top row and first column (final collection step)

$$\begin{pmatrix} A' & B' & C \\ D'' & E'' & F \\ G' & H' & J' \end{pmatrix}$$

where

$$\begin{aligned} A' &= A - B'D'' - CG' \\ &= A - CJ^{-1}G \\ &\quad -(B - CJ^{-1}H) \\ &\quad (E - FJ^{-1}H)^{-1}(D - FJ^{-1}G) \end{aligned}$$

7. Finally, the inversion of A' will yield the corresponding upper left block of the invers.

6 Performance

We analysed the performance of the code for systems consisting of both 250 and 1024 Fe atoms respectively. For these systems we study the scaling properties of the code as a function of number of walkers used in the WL simulation. Every individual LSMS calculation per walker can be distributed onto 250 or 1024 cores. On the Cray XT5 *jaguarpf* system at ORNL's National Center for Computational Sciences (NCCS), we can thus scale these calculations to up to 895 parallel WL walkers for 250 atoms on 223,752 cores and 218 walkers on 223,232 cores for 1024 Fe atoms. In these performance analysis runs, each walker executes 20 WL steps, which is far fewer than a real simulations. Since the setup time of the calculations remains the same if the runs were longer, the performance numbers and scaling properties we report in this section are conservative estimates of the real numbers one would measure during production runs. This becomes most evident in the increased startup time for large numbers of walkers in the 250 atom case.

In figure 3 we show how the time to solution scales if we increase the number of WL walkers, and thus the total number of samples taken in the WL-LSMS simulation. The result shown thus represent a weak scaling plot, and the scaling behavior of the WL-LSMS method looks close to optimal. We find a similarly optimal strong scaling behavior, if we fix the number of samples taken for every run and increase the number of walkers. With the available size of machines today, we are still far from saturating the method in terms of scaling behavior.

The sustained floating point performance of the runs that correspond to the results presented in figure 3 are shown in figure 4. In order to estimate the executed floating point operations of the benchmark runs, we have instrumented the WL-LSMS code with PAPI calls. The average number of floating point operations per second reported in figure 4 is computed from the total number of retired floating point operations as reported by PAPI_FP_OPS events divided by the measured execution time of the runs. The run time has been measured with PAPI calls to PAPI.get_real_usec as well as calls to the C time rou-

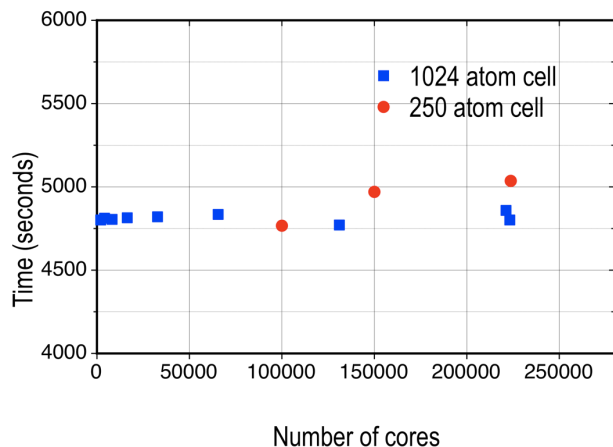


Figure 3: The weak scaling behavior of the WL-LSMS code for 250 and 1024 atom systems with a varying number of Wang-Landau walkers. The times shown represent the total runtime of the code and include the startup costs of the calculations. This accounts for the jump in the runtime for the 250 atom systems as each Wang-Landau walker reads its initial input file.

time. For the largest runs with 218 parallel Wang-Landau walkers of 1024 atoms each and 20 steps per walker, the measured sustained performances was 1.835 petaflop/s, which on the 223,232 AMD Opteron cores running at 2.6 GHz corresponds to a fraction of 79.0% of the theoretical peak performance.

7 Summary and Outlook

In the present paper we have demonstrated our approach for the first principles treatment of finite temperature behavior of magnetic systems. The combination of the most recent massively parallel supercomputing architectures and advances in both algorithmic developments and most importantly new computational methods have made this hybrid statistical mechanics/first principles method feasible.

The code presently is applicable to the evaluation of magnetic transition temperatures of transition metal alloys and has already reproduced the Curie

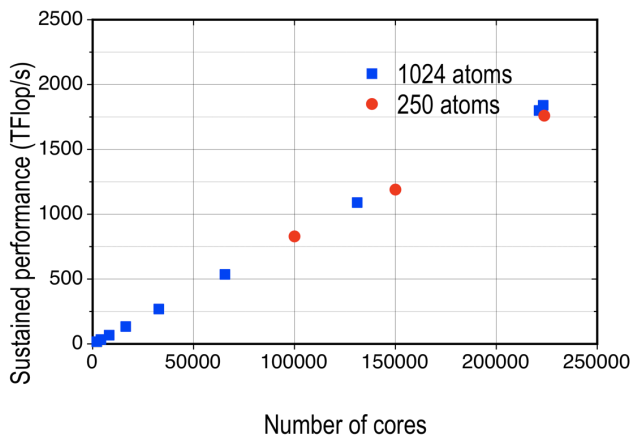


Figure 4: The sustained performance of the Wang-Landau LSMS code on `jaguar` for systems of 250 iron atoms (blue squares) and 1024 atoms (red circles). The code reaches a performance of 1.755 Petaflop/s on 223,752 cores for 250 atoms and 1.835 Petaflop/s on 223,232 cores for a 1024 atom system.

temperature of bulk Iron [9].

The code is being actively developed, both for extended functionality as well as being ported to new computer architectures. In particular the current code is limited by the rigid mapping of one MPI process to one atomic instance. To allow for an efficient deployment on highly multithreaded architectures a new version of LSMS is being developed that lifts this limitation. Also better sampling schemes are investigated that will minimize the number of first principles Hamiltonian evaluations that will be required to reach convergence.

Finally the general concept described here is not limited to magnetic systems or LSMS in particular and a combination of a Wang-Landau method with different constrained first principles calculations could be envisioned.

Acknowledgments

We acknowledge use of resources of the National Center for Computational Sciences and the Center for Nanophase Materials Sciences, which are sponsored

by the respective facilities divisions of the offices of Advanced Scientific Computing Research and Basic Energy Sciences (BES) of the U.S. Department of Energy. The present research was supported by the Laboratory Directed Research and Development fund of ORNL, as well as the division of scientific user facilities of BES. Computations were performed at the Oak Ridge Leadership Computing Facility under an INCITE allocation.

About the Authors

M. Eisenbach is a computational scientist at the National Center for Computational Sciences at Oak Ridge National Laboratory. He is one of the developers of the LSMS code as well as the main author of the WL-LSMS hybrid code.

C.-G. Zhou was a postdoctoral fellow with T. C. Schulthess at Oak Ridge while he developed improvements to the Wang-Landau method for systems of continuous variables, such as the Heisenberg like behavior studied in this work.

D. M. Nicholson is one of the original authors of the LSMS method. He conducts extensive computational research into the relation of magnetism with material properties and contributes to the advancement of density functional based methods.

G. Brown conducts research into the magnetism of model systems using various classical models. For the present work he provided Monte-Carlo reference calculations to aid in the evaluation of the WL-LSMS results. Together with M. Eisenbach and D. M. Nicholson he works on new methods based on the flat histogram sampling scheme that underlies the WL method, to provide methods for more efficient sampling of the Hamiltonian.

J. Larkin of Cray, Inc. provided invaluable help in understanding the scaling behavior of the code and guidance in the selection of the most appropriate numerical libraries.

T. C. Schulthess is the director of the Swiss National Supercomputing Center and a professor at the ETH Zürich. Prior to these appointments he was based at Oak Ridge and was one of the original initiators of the hybrid scheme described in the present

paper.

References

- [1] Martin, R. M., *Electronic Structure: Basic Theory and Practical Methods*, Cambridge, 2004.
- [2] F. Wang, D. P. Landau, *Phys. Rev. Lett.* **86**, 2050 (2001).
- [3] Yang Wang, G. M. Stocks, W. A. Shelton, D. M. C. Nicholson, Z. Szotek, and W. M. Temmerman, *Phys. Rev. Lett.* **75** 2867 (1995).
- [4] M. Eisenbach, B. L. Györfly, G. M. Stocks, and B. Újfalussy, *Phys. Rev. B* **65**, 144424 (2002).
- [5] G. M. Stocks, Y. Wang, D. M. C. Nicholson, W. A. Shelton, Z. Szotek, W. M. Temmerman, B. N. Harmon, V. P. Antropov, *Mater. Res. Soc. Symp. Proc.* **408**, 157 (1996).
- [6] G. M. Stocks, B. Újfalussy, X. Wang, D.M.C. Nicholson, W. A. Shelton, Y. Wang, A. Canning, and B. L. Györfly, *Philos. Mag. B* **78**, 665 (1998).
- [7] B. Újfalussy, X. Wang, D.M.C. Nicholson, W. A. Shelton, G. M. Stocks, Y. Wang, and B. L. Györfly, *J. Appl. Phys.* **85**, 4824 (1999).
- [8] C.-G. Zhou, T. C. Schulthess, S. Torbrügge, D. P. Landau, *Phys. Rev. Lett.* **96**, 120201 (2006).
- [9] M. Eisenbach, C.-G. Zhou, D. M. Nicholson, G. Brown, J. Larkin, T. C. Schulthess, SC'09: Proceedings of the Conference on High Performance Computing, Networking, Storage and Analysis, ACM, 2009.