A New Wavelet-Based Paradigm for Hierarchical Coarse-Graining applied to Materials Modeling*

Ahmed E. Ismail, Gregory C. Rutledge, and George Stephanopoulos
Department of Chemical Engineering, MIT, Cambridge, MA 02139

Abstract— **We outline how the wavelet transform, a hierarchical averaging scheme, can be used to perform both structural and topological coarse-graining in systems with multiscale physical behavior, such as Ising lattices and polymer models. We illustrate how to create sampling mechanisms, which we call wavelet-accelerated Monte Carlo (WAMC), to study these systems and obtain qualitative and quantitatively accurate results in orders of magnitude less time than using atomistic simulations.**

I. Introduction

While there have been impressive computational gains afforded in recent years through advances both in computer hardware and in the expected gains predicted from Moore's law, it is also clear that relying on improvements in computational performance will ultimately be insufficient for advancing the state-of-the-art in molecular simulation. In response, many researchers have turned their attention toward the development of a variety of algorithms for advancing the time and length scales accessible to molecular simulations. One limitation common to many of these algorithms is their limited focus—they are generally designed to study only a specific set of molecular chemistries. However, it is well known that many physical systems possess common structural properties, including self-similarity. As a result, we have created a new simulation paradigm for studying systems with structured behavior over various length scales which exploits these links.

Our method is based upon the wavelet transform [1–3, among others]. Although most commonly used for signal processing and image analysis, we use it for its data compression properties. The basic principle in our application of the wavelet transform is that we use a characteristic property of our objects—such as the spin of a magnetic particle, or the position of an atom along the backbone of a chain molecule—as the basis for our sampling. The wavelet transform is then used to develop coarse-grained representations of "effective" or "block" variables, as well as potentials, describing the behavior of the system over successively larger length scales. This approach has several advantages, including generality with respect to the range of systems to which it can be applied; extensibility, as it can be used as the basis for a hierarchical simulation scheme; and efficiency, as the resulting algorithm is capable of yielding qualitatively accurate predictions of behavior over a wide range of parameters orders of magnitude faster than is capable with traditional atomistic simulations. We discuss our development of this paradigm, in the form of a new simulation technique, *wavelet-accelerated Monte Carlo*, through its application to lattice systems and to polymer random walks.

To date, the wavelet framework transform has not been extensively applied to models in statistical mechanics. Huang uses wavelet analysis to observe the statistical distribution of multiplicity fluctuations in a lattice gas [4], while Gamero *et al.* employ wavelets to introduce their notion of multiresolution entropy, but for dynamic signal analysis rather than statistical mechanics simulations [5]. O'Carroll attempts to establish a theoretical foundation connecting wavelets to the block renormalization group [6], [7]. The most extensive discussion of the relationship between wavelets and renormalization group theory is a recent monograph by Battle [8].

II. Using the wavelet transform as a coarse-graining scheme

There are two principal means for carrying out coarse-graining of a physical system: we can call these approaches *structural* and *topological.* Structural approaches operate on the simulation space, dividing it into regions, and then combining the regions from one scale to another. Topological approaches operate on the particles inserted into the simulation space, so that the rules for coarse-graining the system do not depend upon any specific physical structure of the simulation space, such as a lattice, but instead upon the structure of the particles. Ising lattices and polymer chains typify each of these individual approaches. We show how to develop coarse-grained simulations for each of these models below.

A. The Ising model

The Ising model is the standard model for studying the thermodynamic behavior of lattice systems, such as spin magnets, lattice gases, and binary alloys [9]. The Hamiltonian for the Ising model, which contains both nearest-neighbor pairwise interactions as well as interactions between lattice sites and an external field, can be generally written in the form

$$-\beta H = \sum_i h_i \sigma_i + \sum_i \sum_j J_{ij} \sigma_i \sigma_j, \qquad (1)$$

where σ_i is either the occupation number or the spin of lattice site i, h_i is the strength of the external field in the

direction of the spins σ_i at site i, and J_{ij} is the strength of the interaction between the pair of sites i and j. The spin-$\frac{1}{2}$ Ising model, the most frequently encountered, has $\sigma_i = +1$ or -1 and $J_{ij} = 0$, unless i and j are nearest neighbors on the lattice, in which case J_{ij} is a constant. The inverse temperature $\beta = (k_B T)^{-1}$; for convenience we choose $k_B = 1$, so that temperature, external field, and nearest-neighbor interactions are all dimensionless quantities.

While (1) is a compact representation of the Hamiltonian of the system, the expansion of the lattice variables σ_i and σ_j as a sum of wavelet coefficients makes these equations impractical for applying the wavelet transformation. Since the system is described on a regular lattice, we want to use discrete wavelets (also known as filter banks), so a matrix formulation of the Hamiltonian would be convenient. Using graph theory to represent the connectivity of the lattice as a matrix [10], we define the vectors $u = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ and $h = (h_1, \ldots, h_N)$ as the values of the N lattice variables in the system and the set of external-field strengths. Furthermore, define the matrix J such that element J_{ij} is the strength of the interaction between site i and site j. If sites i and j do not interact, then $J_{ij} = 0$. Then, the Hamiltonian (1) can be written in the form of a matrix equation:

$$-\beta H = h^T u + u^T J u, \tag{2}$$

where the superscript T denotes the transpose of the vector (or matrix) which precedes it.

The matrix W which defines the wavelet transform [1–3, among others] satisfies by construction $W^T W = I$, where I is the identity matrix. Therefore, to apply the wavelet transform, we insert $W^T W$ between each pair of terms in (2), thereby obtaining

$$-\beta H = \left(h^T W^T\right)(Wu) + \left(u^T W^T\right)\left(WJW^T\right)(Wu). \tag{3}$$

Using the general matrix property that $B^T A^T = (AB)^T$, (3) can be rewritten in terms of the wavelet-transformed vectors $\tilde{h}^{(K)} = Wh$, $\tilde{u}^{(K)} = Wu$, and the wavelet-transformed matrix $\tilde{J}^{(K)} = WJW^T$:

$$-\beta \tilde{H} = \left(\tilde{h}^{(K)}\right)^T \tilde{u}^{(K)} + \left(\tilde{u}^{(K)}\right)^T \tilde{J}^{(K)} \tilde{u}^{(K)}. \tag{4}$$

Using (4) as the basis for a Monte Carlo simulation requires the calculation of the change of energy ΔE_{nm} from microstate $\mathbf{u}_m$ to microstate $\mathbf{u}_n$:

$$\Delta E_{nm} = h^T (u_n - u_m) + (u_n - u_m)^T J (u_n - u_m). \tag{5}$$

If moves are restricted to changes of single spin flips, then only a single entry of $u_n - u_m$ is nonzero, and therefore the calculation (5) reduces to a dot product, instead of a matrix multiplication.

The effect of applying the wavelet transform in (4) is to create a representation in which the $\tilde{u}^{(K)}$ represent "block spins" whose values are determined by wavelet averaging over a well-defined region of the original system. The Hamiltonians (1) and (4) have the same formal structure, so that Monte Carlo simulations of the two systems are essentially identical. The only modifications needed to simulate a coarse-grained Hamiltonian are the ability to select new microstates $\tilde{u}_i^{(K)}$ which are generated through wavelet transformations of the original microstates u_i, and the elimination of unwanted degrees of freedom from (4). It should be noted that in (4), the elements of $\tilde{u}^{(K)}$ are not restricted to the same values as in the original system, but are free to take on any value which is consistent with the wavelet transform applied to the system.

To draw a new microstate $\tilde{u}^{(K)}$, we need an estimate for the probability distribution $p\left(\tilde{u}_i^{(K)}\right)$ for the individual sites in the new, coarse-grained lattice. Determining the correct distribution for a given $\tilde{u}_i^{(K)}$ would require a detailed simulation of the original system. An alternative, ignoring the effect of neighboring block spins, would be to perform an exact enumeration of the spins within a block, which is possible only for the smallest of block spins. Since we would like to apply this method to systems of arbitrary size, we want to avoid both of these options. Therefore, we simulate a sublattice with the same dimensions as $\tilde{u}_i^{(K)}$, ignoring physical interactions with the rest of the system by using either free or periodic boundary conditions. Using the standard Metropolis acceptance criterion, we compute distributions for the properties of the small lattice, such as the magnetization. Then, since the wavelet transform defines a single block spin $\tilde{u}_i^{(K)}$ as a linear function of the individual spins at level $K-1$ which it replaces, we can use the linearity properties of probability distributions to convert the distribution of the properties directly into a distribution for the block spin $\tilde{u}_i^{(K)}$ [11]. Finally, using the distribution for $\tilde{u}_i^{(K)}$ as a starting point, we perform a Monte Carlo simulation on the system of block spins defined by the Hamiltonian (4).

Although (1) and (4) are structurally the same, we cannot impose a one-to-one correspondence between the states in the configuration space of (5) and the states in the configuration space of (4). Consequently, the thermodynamic information obtained from the two will not necessarily be identical: the inability to sample correctly the exact distribution $p\left(\tilde{u}_i^{(K)}\right)$ ensures that there will be a loss of entropy associated with the coarse-graining process. However, we can still ensure that the detailed balance condition for the simulation based on (4) is satisfied for the new simulation by requiring

$$\frac{\alpha\left(\tilde{u}_n^{(K)} \to \tilde{u}_m^{(K)}\right)}{\alpha\left(\tilde{u}_m^{(K)} \to \tilde{u}_n^{(K)}\right)} = \frac{p\left(\tilde{u}_n^{(K)}\right)}{p\left(\tilde{u}_m^{(K)}\right)} e^{-\beta\left(H\left(\tilde{u}_n^{(K)}\right) - H\left(\tilde{u}_m^{(K)}\right)\right)},$$

where $\alpha(m \to n)$ is the probability of accepting a move from microstate m to microstate n, and $p(m)$ is the probability of selecting microstate m as determined from simulations on finer-grained lattices at lower scales.

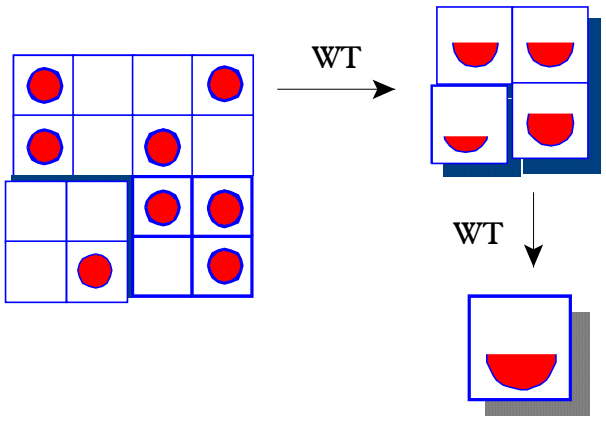

Fig. 1. Coarse-graining model for a lattice gas. Each 2×2 block is converted into a new block whose occupation is the average of the occupation numbers at the previous scale.

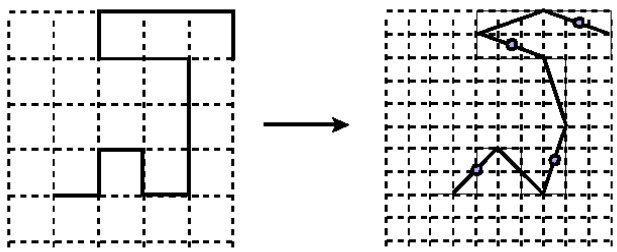

Fig. 2. Coarse-graining of a 16-site random walk in two dimensions using "center-of-mass aggregation" of adjacent points along the chain, as suggested by the wavelet transform method. The four sites which would be created after another iteration of this method are indicated on the graph; note that each is at a quarter-integer lattice point. [*N.B.* While shown in two dimensions for simplicity, all simulations were three-dimensional.]

The effect of this procedure is shown schematically in Figure 1, which represents a lattice gas. At the original scale, there are only two options for each site—occupied or unoccupied. Applying the wavelet transform to a 2×2 block on this lattice leads to a new "block site," which can have occupations ranging from fully empty to fully occupied, with one-quarter, half, and three-quarters occupied as the other options. If we choose, we can then re-block each 2×2 segment of the new model to produce an even coarser model, where each site can take one of 17 different values (corresponding to occupation fractions of $0, \frac{1}{16}, \frac{1}{8}, \ldots, 1$). The Hamiltonian which describes each of these systems is structurally equivalent, but the numeric constants h_i and J_{ij} in 4 will be different in each case.

B. Polymer chains

The fundamental concept behind the application of the wavelet transform to a polymer chain is that we apply the wavelet transform topologically along the chain backbone to the positions of the individual "atoms" (or functional groups or repeat units). Thus, we take as our input data the set of positions $R = \{r_1, \ldots r_{2N}\}$, and define the unnormalized Haar wavelet transform as the mapping

$$r_n^{(k)} = \frac{1}{2}\left[r_{2n-1}^{(k-1)} + r_{2n}^{(k-1)}\right], \tag{6}$$

$$w_n^{(k)} = \frac{1}{2}\left[r_{2n-1}^{(k-1)} - r_{2n}^{(k-1)}\right]. \tag{7}$$

The output is the set of averages $\left\{r_1^{(k)}, r_2^{(k)}, \ldots, r_N^{(k)}\right\}$ and the set of differences $\left\{w_1^{(k)}, w_2^{(k)}, \ldots, w_N^{(k)}\right\}$, where the superscript denotes the number of times we have repeated the process (on the set of scaling coefficients). The effect of the averaging operator in (6) is to create a new coarse-grained bead $r_n^{(k)}$ at the center of mass of the beads at $r_{2n-1}^{(k)}$ and $r_{2n}^{(k)}$; the differencing operator (7) returns the distance between the original particles' positions and their center of mass. A schematic representation is shown in Figure 2.

The methods most similar to the approach presented here are the "soft colloid" method of Louis *et al.* [12] and the "bond fluctuation" model of Carmesin and Kremer [13]. The specific hierarchical method discussed here is closely related in spirit to renormalization group procedures, and in particular to the blocking technique of Kadanoff [14], [15].

While we perform our simulations using only the averaged coordinates $\left\{r_1^{(k)}, r_2^{(k)}, \ldots, r_N^{(k)}\right\}$, if we were also able to keep track of the difference coordinates $\left\{w_1^{(k)}, w_2^{(k)}, \ldots, w_N^{(k)}\right\}$, we could reconstruct the chain at the next finest scale $\left\{r_1^{(k-1)}, r_2^{(k-1)}, \ldots, r_N^{(k-1)}\right\}$ as a consequence of perfect reconstruction [3]. The challenge of this approach is that the probabilistic models which would describe the difference coordinates is in general more complicated than those for the averaged coordinates, which makes working with them much more computationally expensive. Thus, if it were vital to develop a model for the differencing coordinates, it will almost always be to our advantage to simulate the system at scale $k-1$, and then use (6) and (7) to compute the difference coordinates.

The total length N of a chain can be written as $N = N_b \times N_e$, where N_b is the number of beads actually being simulated in the chain, and N_e is the "effective size" of a bead; that is, N_e represents the number of beads on the original representation of the chain that are combined to form a single coarse-grained bead. By definition, an atomistic simulation has $N_e = 1$ and therefore $N = N_b$. If we proceed through multiple stages of a hierarchical simulation, we write the vector $N_e = (N_{e,1}, N_{e,2}, \ldots, N_{e,m})$, where m is the number of stages in the simulation so far. The effective bead size $N_e^{(m)}$ is then given by $N_e^{(m)} = \prod_{i=1}^{m} N_{e,i}$.

Our hierarchical algorithm proceeds as follows. We start by performing an atomistic simulation ($O\left(10^6\right)$ attempted MC moves) on a chain containing $N_{b,1}$ beads, with $N_{e,1} = 1$. Then, every $N_{b,1}$ steps, we use (6) and (7) to determine the centers-of-mass $\mathbf{r}^{(K)}$ of the coarse-grained beads at a given length scale $N_{e,2} = 2^K$, and the wavelet coeffi-

cients $\mathbf{w}^{(1)}, \ldots, \mathbf{w}^{(K)}$. From this data, we can store the necessary information for the distribution of coarse-grained bond lengths, bond angles, and torsion angles at the next stage of the simulation. The next step is a coarse-grained simulation of a chain of length $N_{b,2}$ beads, where each coarse-grained bead represents $N_{e,2}$ beads of the original chain, where $N_{e,2} = 2^K \leq N_{b,1}/4$. If the chain length $N = N_{b,2} \times N_{e,2}$ is the desired length, we can terminate the simulation; otherwise, we can proceed to a third stage simulation of $N_{b,3}$ beads, each with effective size $N_{e,3} \leq (N_{b,2} \times N_{e,2})/4$. This process can be repeated as many times as is necessary to achieve the desired system length. The coarse-grained algorithm runs independently of the source of the input distributions: they can be derived from experimental data, from an atomistic simulation, or from a previous iteration of the coarse-grained algorithm.

A move on the coarse-grained lattice consists of replacing a selected bond with a new bond whose internal coordinates are drawn from the bond-length, bond-angle, and torsion-angle distributions derived in the previous stage of the simulation. Following the work of Dautenhahn and Hall [16], the remainder of one side of the chain is displaced correspondingly to preserve the connectivity of the chain. Once the new configuration is obtained, it is checked for self-avoidance.

III. Results

A. The Ising model

The wavelet transform method tends to produce overestimates for the critical point of the system; therefore, if we start with the high-temperature limit of our algorithm and slowly reduce the temperature in our simulation, we can observe the movement toward the critical point by watching various fluctuation parameters, such as the heat capacity $C_H = \left(\langle E^2 \rangle - \langle E \rangle^2\right)/k_B T^2$. Near the critical point, we expect to see a rapid increase in the value of C_H, consistent with the expected logarithmic divergence observed in the limit of finite-size systems [17], [18]. If we use the onset of this logarithmic divergence as an indicator, we can then "step down" and use a finer lattice with more degrees of freedom. This system will naturally better reflect the physics of our system, particularly in the vicinity of the critical point. We expect that very near the critical point, we will have to simulate the system at the original scale, since this will be effectively the only level which accurately represents the underlying behavior of the system. However, the region of parameter space where this is necessary is relatively small compared to the complete parameter space. This is especially true when we consider that as we proceed below the critical temperature of the system, the logarithmic divergence of C_H will also vanish. As a result, as we move increasingly far away from the critical point, we begin to approach the other fixed-point behaviors associated with the low-temperature limits of the system. Since these are reasonably well-preserved using the wavelet transformation, we can safely return to increasingly coarse-grained descriptions of our system as the simulation proceeds past the critical point.

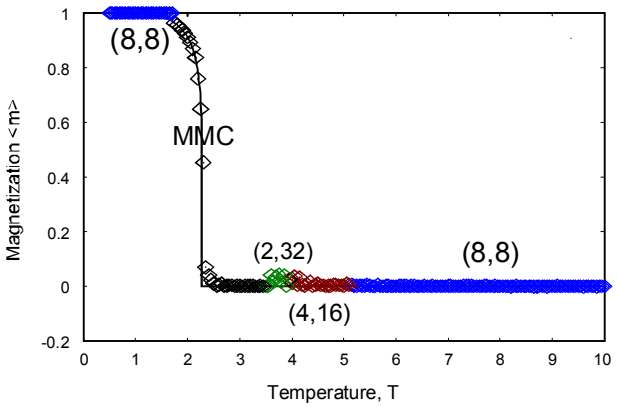

Fig. 3. Illustration of "adaptive" coarse-graining scheme for a 64×64 Ising model. The simulation sizes are as indicated; the model spends approximately 90 per cent of its computational time in the region marked MMC (Metropolis Monte Carlo).

As an example, we compute the spontaneous magnetization curve for a 64×64 Ising lattice in the temperature range $0.5 \leq T \leq 10.0$, with $\Delta T = -0.05$, and choosing as our refinement criterion $\Delta C_H/\Delta T \leq -0.5$, until we reach the finest scale, corresponding to the original problem. We begin by coarse-graining the system to an $(8,8)$-model, where we find that the criterion is triggered only at $T = 5.1$; we then continue with a $(4,16)$-model, down to $T = 4.0$, at which point the refinement criterion is exceeded. Refining once more, we proceed with a $(2,32)$-model until $T = 3.4$, at which point the threshold is again crossed. Since the next refinement is the original problem, we proceed at this level of resolution until we have passed the critical point, so that $\Delta C_H/\Delta T$ is positive. As a coarsening criterion, we select for simplicity the opposite of the refinement criterion, $\Delta C_H/\Delta T \leq 0.5$. Using this criterion, we find that we coarsen the model to the $(2,32)$-, $(4,16)$- and $(8,8)$-models at temperatures of $T = 1.75$, $T = 1.65$, and $T = 1.55$, respectively. The rapid coarsening of the model results from the higher estimates of the critical point in the coarsened models. Since we are well past the critical point, we expect changes in the heat capacity as a function of temperature to be relatively small, and thus it is possible to obtain accurate results from a relatively coarse model. Computationally, the time required to create this diagram was only 28 per cent that required to perform a standard Metropolis Monte Carlo simulation with the same number of steps. Moreover, in the regions that were not simulated using MMC, the computation time required was just 8 per cent of the time required for MMC. The resulting plot of magnetization versus temperature, shown as Figure 3, compares favorably to the analytical solution of Onsager, which is also shown [19].

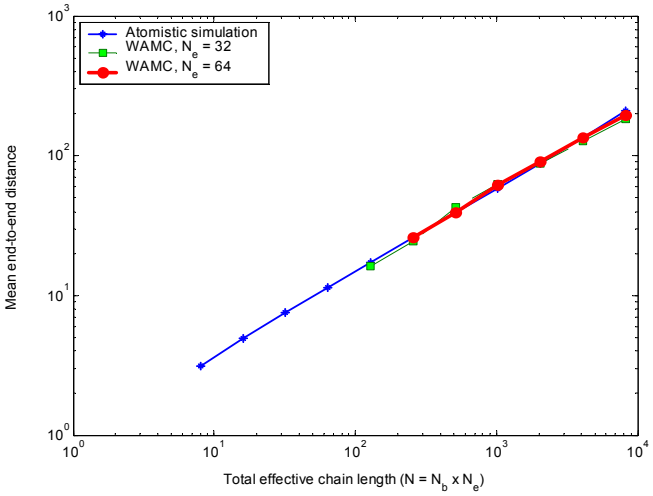

Fig. 4. Root mean square end-to-end distance $\langle R^2 \rangle^{1/2}$ of the pivot algorithm (diamonds) and of coarse-grained simulations based on bead sizes of 32 and 64 (squares and circles, respectively).

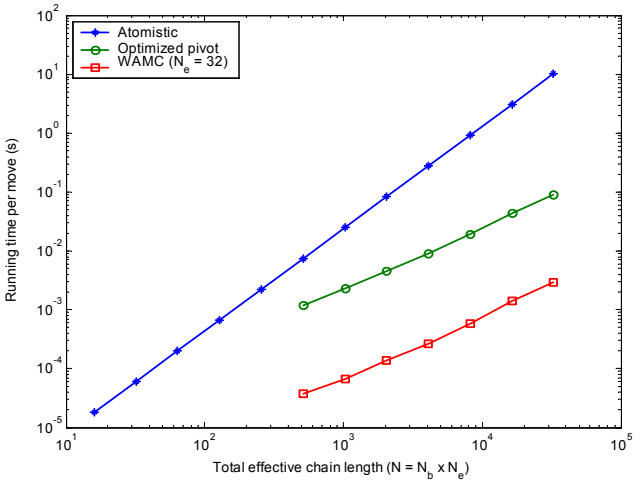

Fig. 5. A comparison of the running times of atomistic, optimized pivot, and WAMC algorithms for polymer chains as a function of the total effective chain length (top, middle, and bottom curves, respectively).

B. Polymer chains

A comparison of the results obtained for a detailed atomistic simulation and for two-scale coarse-grained simulations of chains up to length $2^{13} = 8192$ beads are shown in Figure 4. Two different coarse-grained simulations are plotted, showing $N_{e,2} = 32$ and 64, respectively; the mean end-to-end distances of the three simulations show excellent agreement with one another, indicating that for relatively large values of $N_{e,2}$, there is little effect on the end-to-end distance. Additionally, the Flory exponent for the scaling of the end-to-end distance, $\langle R \rangle \sim N^{\nu}$, is shown to be $\nu_{cg} \approx 0.578 \pm 0.006$ for $N_{e,2} = 32$ and $\nu = 0.581 \pm 0.006$ for $N_{e,2} = 64$, both of which are within one percent of the best available estimate $\nu \approx 0.577$ [20].

To show that this hierarchical simulation strategy is effective for studying large systems, we can measure its performance over a variety of chain sizes, as shown in Figure 5, which shows the running times for three different algorithms. The topmost line in the graph corresponds to a detailed atomistic simulation using an algorithm comparable to the WAMC algorithm; the running time of this algorithm is approximately $O\left(N^{7/4}\right)$. If we optimize the atomistic algorithm by using the pivot algorithm, along with some of the additional improvements suggested by Kennedy [21], we, obtain the "optimized atomistic" plot shown. The running time has been reduced to roughly $O\left(N^{6/5}\right)$; still further changes can reduce the running time to $O(N)$ or below; however, the resulting algorithm is impractical for use in models with interparticle potentials other than self-avoidance.

The lower line plotted in Figure 5 represents the running time of the WAMC algorithm with $N_e = 32$. We note that the algorithm has a running time of $O\left(N_b^{7/5}\right)$, which is in between the results for the atomistic and optimized atomistic algorithms. However, since by definition we have $N_b \ll N$, we expect that the WAMC algorithm will be faster on a per-move basis than either atomistic algorithm, unless the associated prefactor is much larger for the WAMC algorithm. However, since both methods are executing variations on the same basic process, this does not occur, and the WAMC algorithm is generally faster than either atomistic technique.

IV. Conclusions

The wavelet-accelerated Monte Carlo scheme is a robust technique for performing coarse-grained molecular systems with multiscale behavior. The simulations, although requiring more structural information than the initial atomistic simulation, have many fewer degrees of freedom, and therefore execute much more quickly. In addition, its adaptable nature means that we can adjust our coarse-graining strategy without knowing *a priori* any information about the system to be studied. Moreover, this gives us the ability to explore physics at any scale between those of the atomistic models and the model with the maximum allowable coarse-graining, a feature which other "handshaking" coarse-graining methodologies cannot offer.

References

[1] I. Daubechies, *Ten Lectures on Wavelets*, vol. 61 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Philadelphia: SIAM, 1992.

[2] W. Sweldens, "The lifting scheme: A new philosophy in biorthogonal wavelet constructions," in *Wavelet Applications in Signal and Image Processing III* (A. F. Laine and M. Unser, eds.), pp. 68–79, Proc. SPIE 2569, 1995.

[3] G. Strang and T. Nguyen, *Wavelets and Filter Banks*. Cambridge, MA: Wellesley-Cambridge, 1996.

[4] D.-W. Huang, "Wavelet analysis in multiplicity fluctuations," *Phys. Rev. D*, vol. 56, pp. 3961–3969, 1997.

[5] L. G. Gamero, A. Plastino, and M. E. Torres, "Wavelet analysis and nonlinear dynamics in a nonextensive setting," *Physica A*, vol. 246, pp. 487–509, 1997.
[6] M. O'Carroll, "Multiscale representation of generating and correlation functions for some models of statistical mechanics and quantum field theory," *J. Stat. Phys.*, vol. 73, pp. 945–958, 1993.
[7] M. O'Carroll, "Lattice and continuum wavelets and the block renormalization group," *J. Stat. Phys.*, vol. 71, pp. 415–423, 1993.
[8] G. Battle, *Wavelets and Renormalization*. Singapore: World Scientific, 1999.
[9] D. A. Lavis and G. M. Bell, *Statistical Mechanics of Lattice Systems 1: Closed-Form and Exact Solutions*. Berlin: Springer, 1999.
[10] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*. Cambridge, MA: McGraw Hill-MIT Press, 1990.
[11] P. Billingsley, *Probability and Measure*. New York: Wiley Interscience, 1995.
[12] A. A. Louis, P. G. Bolhuis, J.-P. Hansen, and E. J. Meijer, "Can polymer coils be modeled as "soft colloids"?," *Phys. Rev. Lett.*, vol. 85, pp. 2522–2525, 2000.
[13] I. Carmesin and K. Kremer, "The bond fluctuation method: A new effective algorithm for the dynamics of polymers in all spatial dimensions," *Macromolecules*, vol. 21, pp. 2819–2823, 1988.
[14] N. Goldenfeld, *Lectures on Phase Transitions and the Renormalization Group.* Reading, MA: Addison-Wesley, 1992.
[15] L. P. Kadanoff, *Statistical Physics: Statics, Dynamics, and Renormalization*. Singapore: World Scientific, 2000.
[16] J. Dautenhahn and C. K. Hall, "Monte Carlo simulation of off-lattice polymer chains: Effective pair potentials in dilute solution," *Macromolecules*, vol. 5399, pp. 5399–5412, 1994.
[17] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications*. San Diego: Academic Press, 1996.
[18] R. K. Pathria, *Statistical Mechanics*. Woburn, MA: Butterworth-Heinemann, 1996.
[19] L. Onsager, "Crystal statistics 1: A two-dimensional model with an order-disorder transition," *Phys. Rev.*, vol. 65, pp. 117–149, 1944.
[20] B. D. Hughes, *Random Walks and Random Environments*. New York: Oxford-Clarendon Press, 1995.
[21] T. Kennedy, "A faster implementation of the pivot algorithm for self-avoiding walks," *J. Stat. Phys.*, vol. 106, pp. 407–429, 2002.