Hybrid Multicanonical Cluster Algorithm for Efficient Simulations of Long-Range Spin Models

Sylvain Reynal, Hung-The Diep

Hybrid Multicanonical Cluster Algorithm for Efficient Simulations of Long-Range Spin Models

S. Reynal* and H.T.Diep
Laboratoire de Physique Théorique et Modélisation, CNRS-Université de Cergy-Pontoise (UMR 8089), 2 avenue A. Chauvin, F-95302 Cergy-Pontoise Cedex, France
(Dated: October 7, 2005)

An efficient, flat histogram Monte Carlo algorithm is proposed that simulates long-range spin models in the multicanonical ensemble with very low dynamic exponents and drastically reduced computational effort. The method combines a random-walk in energy space with cluster updates, where bond weights depend continuously on the lattice energy. Application to q-state Potts chains with power-law decaying interactions is considered. Lattice sizes as high as 2^{16} spins, unattainable with conventional flat histogram algorithms, are investigated. Numerical results demonstrate the remarkable performance of the method over a wide spectrum of model parameters.

PACS numbers: 05.10.Ln, 64.60.Cn, 75.10.Hk

I. INTRODUCTION

Amidst powerful methods dedicated to the study of long-range spin models, Monte Carlo (MC) methods have now gained a prominent role [1, 2, 3, 4]. Whether they rely on single-spin or cluster updates, *canonical* MC simulations of long-range models still pose the same limitations as do numerical studies of their short-range counterparts: (i) the computation of free energies is cumbersome, rendering e.g., a precise determination of the order of the transition intractable; (ii) for models exhibiting complicated energy landscapes or first-order phase transitions, the dynamics of the Markovian chain often sticks to local free energy minima, making it necessary to carry out simulations over exceedingly long times [5]. Efficient algorithms that work out these limitations are flat histogram algorithms, which operate in the multicanonical ensemble [5, 6, 7, 8], and engender a random-walk in energy space. The Markovian chain is thus weighed by a *multicanonical* weight $w(E) \sim 1/n(E)$, where $n(E)$ is an estimate of the density of states. In the last decade, several efficient schemes have been devised that compute $n(E)$ in the course of the simulation itself, e.g., Wang-Landau's algorithm [6] or the transition matrix method [8].

Local-update implementations of these algorithms were shown to reduce tunneling times from an exponential- to a power-law of the lattice size $\tau_e \sim L^z$ [7], yet with dynamic exponents z being still substantially higher than the ideal value $z \sim D$ expected from a random-walk argument. For long-range spin models, the need to compute the lattice energy anew at every MC sweep brings about an additional hurdle, since of order L^{2D} operations are involved. This yields a total CPU load per independent measurement scaling as L^{2D+z}, and limits studies to modest lattice sizes [4], where strong finite-size effects constitute a serious hindrance.

Conversely, cluster algorithms [1, 9] are known to drastically reduce dynamic exponents, owing to their ability to minimize correlations by updating spins in a collective way. Because such algorithms hinge on particular symmetries of the model, and the multicanonical weight $w(E)$ does not keep track of them, embedding collective updates in a multicanonical algorithm is not straightforward. The purpose of the present work is to develop a hybrid method which successfully tackles this challenge in an efficient and versatile way. In the case of long-range interactions, the method also drastically reduces the simulational effort through an optimized computation of the lattice energy.

II. ALGORITHM

Our method is primarily based on the observation that, in local-update multicanonical algorithms, it is the microcanonical temperature which governs the dynamics of the Markovian chain [4]. Indeed, a single spin is flipped, and the move is accepted with the acceptance rate $W = \min[1, n(E)/n(E+\epsilon)]$, where E and ϵ denote the initial energy and the energy variation respectively. Carrying out a series expansion in ϵ yields $W \sim \min[1, \exp(-\beta(E)\epsilon)]$, where $\beta(E) = d \ln n(E)/dE$ is an estimate of the inverse microcanonical temperature; hence a multicanonical dynamics is *locally* equivalent to a canonical dynamics at inverse temperature $\beta(E)$. We therefore propose to build clusters of spins by placing bonds between spins with the same probability as given by a canonical cluster algorithm operating at $\beta(E)$. To be specific, we consider a ferromagnetic long-range Potts model whose Hamiltonian reads $E = -\sum_{i<j} J(|i-j|)\delta_{\sigma_i,\sigma_j}$. The σ_i variables take on integer values between 1 and q, and $J(|i-j|) > 0$. We write the multicanonical weight as $w(E) = \phi(E)\exp[-\beta(E)E]$ and, taking guidance from Swendsen-Wang's algorithm [9], expand the Boltzmann-like exponential term as a trace over the bonds of a random-bond cluster. This yields

$$w(E) = \phi(E) \sum_{[b]} \prod_{i<j} p_{|i-j|}(E)\delta_{\sigma_i,\sigma_j}\delta_{b_{ij},1} + \delta_{b_{ij},0},$$

where the trace is over all lattice bonds, a bond is active (inactive) whenever $b_{ij} = 1$ (0), and $p_{|i-j|}(E) = \exp[\beta(E)J(|i-j|)] - 1$ is interpreted as the statistical weight of a bond linking spins i and j. As opposed to Swendsen-Wang's algorithm,

*Electronic address: reynal@ensea.fr; URL: `http://www-reynal.ensea.fr`; Permanent address: E.N.S.E.A., 6 avenue du Ponceau, 95014 Cergy Cedex, France.

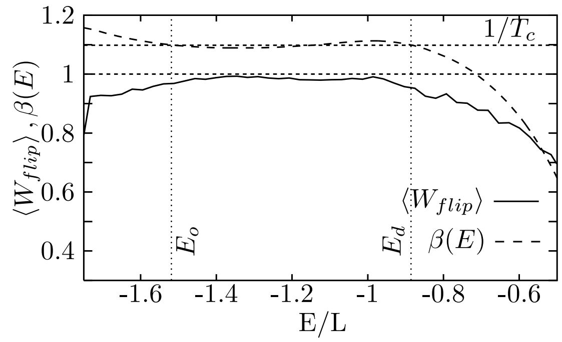

FIG. 1: Inverse microcanonical temperature $\beta(E)$ and mean acceptance rate $\langle W_{flip}\rangle$ as a function of the energy per spin, for the long-range Potts chain with $q = 6$, $\sigma = 0.7$, and $L = 1024$ spins. E_o and E_d denote the energy of the histogram peaks corresponding to the ordered and disordered phase, respectively. The finite size transition temperature T_c and the 100% line are shown for convenience.

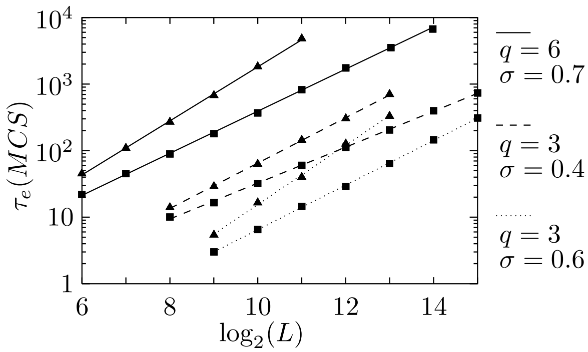

FIG. 2: Tunneling times for the long-range Potts chain. Triangles and squares refer to the local-update and to our algorithm, respectively.

bond weights thus vary continuously with the lattice energy. A simple way of constructing clusters may then consider each pair of spins in turn for bond activation, and activate a bond with a probability $1 - \exp[-\beta(E)J(|i-j|)]$. For interactions decaying with the interparticle distance, however, a large number of bonds have a negligible activation probability; a more efficient approach [1] consists, for each spin, in drawing at random the index of the next spin to be *provisionally* added to it using a cumulative probability, and then to place a bond if both spins match. This method was shown to reduce the number of operations per cluster construction from L^{2D} to roughly L^D. Finally, each cluster is assigned a random spin value, and the new configuration at energy $E + \epsilon$ is accepted with the following acceptance rate,

$$W_{flip} = \min\left(1, \frac{\phi(E+\epsilon)}{\phi(E)} \prod_{l>0} \left[\frac{p_l(E+\epsilon)}{p_l(E)}\right]^{B(l)}\right), \quad (1)$$

where $B(l)$ refers to the number of active bonds of length l. The computation of the lattice energy is obviously a crucial part of the update procedure, yet its complexity scales as $O(L^{2D})$. This can be efficiently cut down to $O(L^D \ln L^D)$ by relying on an FFT implementation of the convolution theorem [10]. Noteworthy enough, such a reduction is of benefit only because a whole lattice update is carried out at a time.

III. NUMERICAL RESULTS

In order to investigate the efficiency of our algorithm, we have conducted simulations on q-state Potts chains with $q = 3$, 6 and 12, and power-law decaying interactions, i.e., $1/|i-j|^{1+\sigma}$. The density of states $n(E)$ was estimated using Wang-Landau's method [6], and the microcanonical temperature $\beta(E)$ was periodically updated from it using a spline interpolation. It also proved advantageous to rely on a transition matrix scheme [8, 11] to bootstrap the computation of $\beta(E)$ during the very first iterations of Wang-Landau's algorithm. We briefly discuss some of our estimates of transition temperatures, which we computed in two distinct ways: T_{eqh}, corresponding to peaks of the reweighted histogram having equal height; and T_W, defined from the intersections of the ratio W_o/W_d at successive lattice sizes, where W_o and W_d denote the histogram weights of the ordered and disordered phases respectively [12]. Infinite-size temperatures were determined from a fit to $T(L) = T(\infty) + a/L^b$. For $q = 3$ and $\sigma = 0.5$, for instance, we found $T_{eqh} = 1.6877(2)$ and $T_W = 1.68744(2)$, with a perfect fit from $L = 2^9$ to $L = 2^{16}$. These estimates are far more precise than the value of $1.686(4)$ obtained recently with a local-update approach [4]. For the sake of completeness, we performed simulations of the two-dimensional nearest-neighbor Potts model for sizes up to 256×256 and $q = 7$ and 10, and obtained transition temperatures matching the exact results with a five-digits precision [11].

How fast does our algorithm explore the configuration space? A relevant indicator here is the mean acceptance rate, since by construction the acceptance rate in Eq. 1 is not equal to unity. Interestingly, a series expansion shows that $1 - \langle W_{flip}\rangle$ is proportional to $|\beta'(E)\epsilon|$. As sketched in Fig. 1, $\beta(E)$ varies smoothly between the energy peaks of the disordered and the ordered phase, and W_{flip} remains well above 90% inside this energy range, incidentally the energy range of primary interest where analysis of first order transitions is concerned. Other values of q and σ corresponding to the first-order regime yielded an acceptance rate which never fell below 80%. Let alone the benefit of the cluster dynamics itself, which we consider below, this already represents an improvement of a factor 3 over a single-spin-update implementation [4]. In terms of correlations between successive configurations, a convenient indicator is the so-called *tunneling time* [13, 14] defined as one quarter of the number of MC steps needed to travel from one histogram peak to the other, and back. Shown in Fig. 2 are fits to the power law $\tau_e \sim L^z$, giving $z = 0.89(1)$ and $z = 1.11(1)$ for $q = 3$, $\sigma = 0.4$ and $\sigma = 0.6$, and $z = 1.05(1)$ for $q = 6$, $\sigma = 0.7$. This represents a strong reduction with respect to the local-update implementation where we obtained, respectively, $z = 1.13(2)$, $z = 1.48(2)$ and $z = 1.35(3)$; moreover, prefactors clearly favor our method even at small lattice

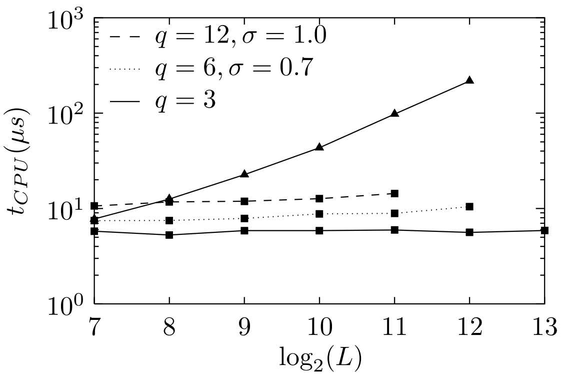

FIG. 3: CPU time per MC step and per spin for the long-range Potts chain. Triangles indicate typical CPU times for the local-update algorithm, irrespective of q and σ. Filled squares refer to our algorithm, where for $q = 3$ estimates were obtained by averaging over $\sigma = 0.4$, 0.5 and 0.6.

sizes. For the nearest-neighbor model, we found $z = 1.82(2)$ for $q = 7$ and $z = 2.23(1)$ for $q = 10$. These exponents are comparable with those obtained with the multibond approach [14]. Our exponents lie thus systematically closer to — and in some cases even below — the ideal value $z \sim D$ expected from a random-walk argument than those obtained with a local-update algorithm.

Restraining the rest of the discussion to the long-range case exclusively, we now consider the gain in CPU time introduced by the FFT accelerated computation of the energy. Figure 3 shows averages of the CPU (user) time per MC step and per spin as measured over a series of one-hour long simulations on various Intel-based CPU architectures. Clearly, the FFT acceleration cuts down the CPU load from a linear to a roughly constant one, with however small fluctuations owing to CPU caches differing in size. For higher q, a slight overhead can be witnessed; this is accounted for by the correspondingly higher number of FFT's to be computed, for the use of the convolution theorem for $q > 3$ requires first mapping the Potts Hamiltonian to an $O(q - 1)$ vector model, and then carrying out an FFT for each vector component separately. The local-update implementation is outperformed already at sizes of several hundreds spins; chains containing up to 2^{16} spins were simulated in a few days, whereas challenging such huge sizes with local updates would have demanded several months of intensive computation.

In conclusion, we have developed a new Monte Carlo method which combines in an efficient and straightforward way the benefits of flat histogram algorithms with the fast-decorrelating capabilities of cluster updates. We have shown that this method proves remarkably powerful when applied to long-range spin models, where the algorithm complexity reduces to that of a short-range model having the same number of spins. Our formulation is nonetheless versatile, and the method can be applied to any spin model for which a random-bond representation can be devised, and to a variety of density of states estimation schemes.

S. R. is greatly indebted to Prof. R. H. Swendsen for fruitful discussions regarding the estimation of the microcanonical temperature.

[1] E. Luijten and H. W. J. Blöte, Int. J. Mod. Phys. **6**, 359 (1995).
[2] K. Binder and E. Luijten, Comp. Phys. Comm. **127**, 126 (2000).
[3] K. Uzelac, Z. Glumac, and A. Anicic, Phys. Rev. E **63**, 037101 (2001).
[4] S. Reynal and H. T. Diep, Phys. Rev. E **69**, 026109 (2004).
[5] B. A. Berg and T. Neuhaus, Phys. Rev. Lett. **68**, 9 (1992).
[6] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).
[7] B. A. Berg, Comp. Phys. Comm. **147**, 52 (2002).
[8] J.-S. Wang and R. H. Swendsen, J. Stat. Phys. **106**, 245 (2002).
[9] R. H. Swendsen and J.-S. Wang, Phys. Rev. Lett. **58**, 86 (1987).
[10] M. Krech and E. Luijten, Phys. Rev. E **61**, 2058 (2000).
[11] S. Reynal and H. T. Diep, unpublished.
[12] C. Borgs and W. Janke, Phys. Rev. Lett. **68**, 1738 (1992).
[13] W. Janke, in *Quantum Simulations of Complex Many-Body Systems: From Theory to Algorithms*, edited by J. Grotendorst, D. Marx, and A. Muramatsu (John Von Neumann Institute for Computing, Jühlich, 2002), vol. 10 of *NIC Series*, pp. 423–445.
[14] W. Janke and S. Kappler, Phys. Rev. Lett. **74**, 212 (1995).