An efficient Wolff algorithm in an external magnetic field

Jaron Kent-Dobias and James P. Sethna

Laboratory of Atomic and Solid State Physics, Cornell University, Ithaca, NY, USA

(Dated: November 9, 2017)

We introduce an extension of the Wolff algorithm that preforms efficiently in an external magnetic field. Near the Ising critical point, the correlation time of our algorithm has a conventional scaling form that reduces to that of the Wolff algorithm at zero field and becomes more efficient at any nonzero field. As an application, we directly measure scaling functions of observables in the metastable state of the 2D Ising model.

The Ising model is a simple model of a magnet comprised of discrete locally interacting one-component spins. Like most systems in statistical mechanics, calculation of its ensemble properties cannot be done explicitly and is often performed using Monte Carlo methods. Near its continuous phase transition, divergent correlation length leads to divergent correlation time in any locally-updating algorithm, hampering computation. With no external field, this critical slowing-down is largely alleviated by cluster algorithms—the most efficient of which is the Wolff algorithm—whose dynamics are nonlocal since each step flips a cluster of spins whose average size diverges with the correlation length. While less efficient cluster algorithms, like Swendsen–Wang, have been modified to perform in nonzero field, Wolff only works at zero field. We describe an extension of the Wolff algorithm that works in arbitrary external field while preserving Wolff's efficiency throughout the entire temperature–field parameter space.

The Wolff algorithm works by first choosing a random spin and adding it to an empty cluster. Every neighbor of that spin that is pointed in the same direction as the spin is also added to the cluster with probability $1 - e^{-2\beta J}$, where $\beta = 1/T$ is inverse temperature and J is the coupling between sites. This process is repeated for the neighbors of every spin added to the cluster. When all sites surrounding the cluster have been exhausted, the cluster is flipped. Our algorithm is a simple extension of this. An extra spin—often referred to as the "ghost spin"—is introduced and made a nearest neighbor of all others with coupling $|H|$, the magnitude of the external field. The traditional Wolff algorithm is then preformed on this new extended lattice exactly as described above, with the extra spin treated no differently from any others, i.e., allowed to be added to clusters and subsequently flipped. Observables in the original system can be exactly estimated using the new one by a simple correspondence.

This paper is divided into three sections. First, the Ising model (and our notation for it) is introduced, along with extant Monte Carlo algorithms. Second, we introduce our algorithm in detail and compare its efficiency with the existing ones. Finally, we use this new algorithm to directly measure critical scaling functions for observables of the 2D Ising model in its metastable state.

INTRODUCTION

Consider an undirected graph $G = (V, E)$ describing a network of interacting spins. The set of vertices $V = \{1, \dots, n\}$ enumerates the sites of the network, and the set of edges E describes connections between neighboring sites. On each site is a spin that can take any value from the set $S = \{-1, 1\}$. The state of the system is described by a function $s : V \to S$, leading to a configuration space of all possible states $S^n = S \times \cdots \times S$. The Hamiltonian $\mathcal{H} : S^n \to \mathbb{R}$ gives the energy of a particular state $s \in S^n$ and is defined by

$$\mathcal{H}(s) = - \sum_{\{i,j\} \in E} J_{ij} s_i s_j - HM(s), \qquad (1)$$

where H is the external magnetic field, $J : E \to \mathbb{R}$ gives the coupling between spins on neighboring sites and $M : S^n \to \mathbb{R}$ is the magnetization of the system defined for a state $s \in S^n$ by

$$M(s) = \sum_{i \in V} s_i. \qquad (2)$$

For the purpose of this study, we will only be considering ferromagnetic systems where the function J is nonnegative.

An observable of the system is a function $A : S^n \to \mathbb{R}$ depending on the system's state. Both the Hamiltonian and magnetization defined above are observables. Assuming the ergodic hypothesis holds, the expected value $\langle A \rangle$ of any observable A is defined by its average over every state s in the configuration space S^n weighted by the Boltzmann factor $e^{-\beta \mathcal{H}(s)}$, or

$$\langle A \rangle = \frac{1}{Z} \sum_{s \in S^n} e^{-\beta \mathcal{H}(s)} A(s), \qquad (3)$$

where the partition function Z is defined by

$$Z = \sum_{s \in S^n} e^{-\beta \mathcal{H}(s)} \qquad (4)$$

and gives the correct normalization for the weighted sum. Unfortunately, the sum over configurations in (3) are intractable for all but for very small systems. Therefore

expectation values are usually approximated by some means. Monte Carlo methods are a common way of accomplishing this. These methods sample states s from the configuration space S^n according to the Boltzmann distribution $e^{-\beta \mathcal{H}(s)}$ so that averages of observables made using their incomplete samples asymptotically approach the true expected value.

The Metropolis–Hastings algorithm [1, 2] is very popular for systems in statistical physics. A starting state $s \in S^n$ is randomly perturbed to the state s', usually by flipping one spin. The change in energy $\Delta \mathcal{H} = \mathcal{H}(s') - \mathcal{H}(s)$ due to the perturbation is then computed. If the change is negative the perturbed state s' is accepted as the new state. Otherwise the perturbed state is accepted with probability $e^{-\beta \Delta \mathcal{H}}$. This process is repeated indefinitely and a sample of states is made by sampling the state s between iterations sufficiently separated that the successively sampled states are uncorrelated, e.g., at separations larger than the correlation time τ. Metropolis–Hastings is very general, but unless the perturbations are very carefully chosen the algorithm suffers in regimes where large correlations are present in the system, for instance near continuous phase transitions. Here the algorithm suffers from what is known as critical slowing-down, where likely states consist of large correlated clusters that take many perturbations to move between in configuration space.

The Wolff algorithm [3] solves many of these problems, but only at zero external field, $H = 0$. This algorithm solves the problem of critical slowing-down by flipping carefully-constructed clusters of spins at once in a way that samples high-correlated states quickly while also always accepting prospective states. A random site i is selected from the graph and its spin is flipped. Each of the site's neighbors j is also flipped with probability $1 - e^{-2\beta J_{ij}}$ if doing so would lower the energy of the bond i–j. The process is repeated with every neighbor that was flipped. While this algorithm successfully addresses the problems of critical slowing down at zero field, it doesn't directly apply at nonzero field. A common extension is the hybrid Wolff–Metropolis, where a Wolff cluster is formed using the ordinary formula, then accepted or rejected Metropolis-style based on the change in energy *relative to the external field*, $H\Delta M$. Because each spin has the same coupling to the external field, a strictly more efficient but exactly equivalent version of this hybrid is made by distributing the cluster acceptance or rejection across the spins in the cluster as they are added. In this version, the cluster is abandoned with probability $1 - e^{-2\beta H}$ every time a spin is added to it.

$z = 0.29(1)$ [4, 5] $z = 0.35(1)$ for Swendsen–Wang [6]

CLUSTER-FLIP IN A FIELD

Consider the new graph $\tilde{G} = (\tilde{V}, \tilde{E})$ defined from G by

$$\tilde{V} = 0 \cup V \tag{5}$$

$$\tilde{E} = E \cup \{\{0, i\} \mid i \in V\}, \tag{6}$$

or by adding a zeroth vertex and edges from every other vertex to the new one. Thee spin at site zero is often known as a "ghost spin." The network of spins described by this graph now have states that occupy a configuration space S^{n+1}. Extend the coupling between spins $\tilde{J} : \tilde{E} \to \mathbb{R}$ by

$$\tilde{J}(e) = \begin{cases} J(e) & e \in E \\ |H| & \text{otherwise,} \end{cases} \tag{7}$$

so that each spin is coupled to every other spin the same way they were before, and coupled to the new spin with strength $|H|$, the magnitude of the external field. Now define a Hamiltonian function $\tilde{\mathcal{H}} : S^{n+1} \to \mathbb{R}$ on this new, larger configuration space defined by

$$\tilde{\mathcal{H}}(s) = - \sum_{\{i,j\} \in \tilde{E}} \tilde{J}_{ij} s_i s_j. \tag{8}$$

This new Hamiltonian resembles the old one, but is comprised only of spin–spin interactions with no external field. However, by changing the terms considered in the sum we may equivalently write

$$\tilde{\mathcal{H}}(s) = - \sum_{\{i,j\} \in E} J_{ij} s_i s_j - H\tilde{M}(s) \tag{9}$$

where the new magnetization $\tilde{M} : S^{n+1} \to \mathbb{R}$ is defined for $s \in S^{n+1}$ by

$$\begin{aligned} \tilde{M}(s) &= \text{sgn}(H) s_0 \sum_{i \in V} s_i \\ &= M(s_0 \, \text{sgn}(H)(s_1, \ldots, s_n)). \end{aligned} \tag{10}$$

In fact, any observable A of the original system can be written as an observable $\tilde{A}$ of the new system by defining

$$\tilde{A}(s) = A(s_0 \, \text{sgn}(H)(s_1, \ldots, s_n)) \tag{11}$$

such that $\langle \tilde{A} \rangle = \langle A \rangle$. This can be seen readily by using the symmetry $\tilde{\mathcal{H}}(-s) = \tilde{\mathcal{H}}(s)$ of the Hamiltonian (8), the invariance of configuration space sums under negation of their summand, and the fact that $\tilde{\mathcal{H}}(1, s) = \mathcal{H}(s)$ for any $s \in S^n$. Notice in particular that this is true for the Hamiltonian $\tilde{\mathcal{H}}$ as well.

Our new spin system with an additional field is simply a ferromagnetic spin system in the absence of an external field. Therefore, the Wolff algorithm can be applied to it with absolutely no modifications. Since there is an exact correspondence between the expectation values of

our ordinary spin system in a field and their appropriately defined counterparts in our new system, sampling Monte Carlo simulations of our new system allows us to estimate the expectation values for the old. This "new" algorithm, if you can call it that, is shown in Algorithm **??**.

At sufficiently small $|H|$ both our algorithm and the hybrid Wolff–Metropolis reduce to the ordinary Wolff algorithm, and have its runtime properties. At very large $|H|$, the hybrid Wolff–Metropolis behaves exactly like Metropolis, where almost only one spin is ever flipped at a time with probability $\sim e^{-2\beta H}$, since the energy is dominated by contributions from the field.

We measured the autocorrelation time τ of the internal energy $\mathcal{H}$ of a square-lattice Ising model ($J = 1$ and E is the set of nearest neighbor pairs) resulting from using each of these three algorithms at various fields and temperatures. This was done using an initial convex sequence estimator, as described in [7]. Time was measured as "the number of spins that the algorithm has attempted to flip." For example, every Metropolis–Hastings step takes unit time, every Wolff step takes time equal to the number of spins in the flipping cluster, and every hybrid Wolff/Metropolis–Hastings step takes time equal to the number of spins in the cluster the step attempted to flip, whether it succeeds or not. The autocorrelation time as a function of reduced field βH is shown in Fig. 1. By this measure of time, the modified Wolff algorithm thermalizes faster than both the hybrid Wolff/Metropolis and ordinary Metropolis–Hastings in all regimes. At high field ($\beta H \sim 1$) the new algorithm has an autocorrelation time of roughly n, while both Metropolis–Hastings and the hybrid algorithm have an autocorrelation time of $2n$. At low temperature, neither the new algorithm nor Metropolis-Hastings have an autocorrelation time that depends on field, but the hybrid method performs very poorly at intermediate values of field in a way that becomes worse with system size, scaling worse than the hybrid algorithm does even it the critical temperature. At high temperature, all three algorithms perform similarly and in a way that is independent of system size. At the critical temperature, both the Metropolis–Hastings and hybrid algorithms experience critical slowing down in a similar way, whereas the new algorithm slows down in a far reduced way, always performing more efficiently at nonzero field than the zero-field Wolff algorithm.

Our algorithm for the Ising model in a field can be generalized to run on the q-spin Potts or $O(n)$ models in exactly the same way as the conventional Wolff algorithm.

MAGNETIZATION ESTIMATOR

At any size, the ensemble average magnetization $\langle M \rangle$ is identically zero at zero field. However, this is not what
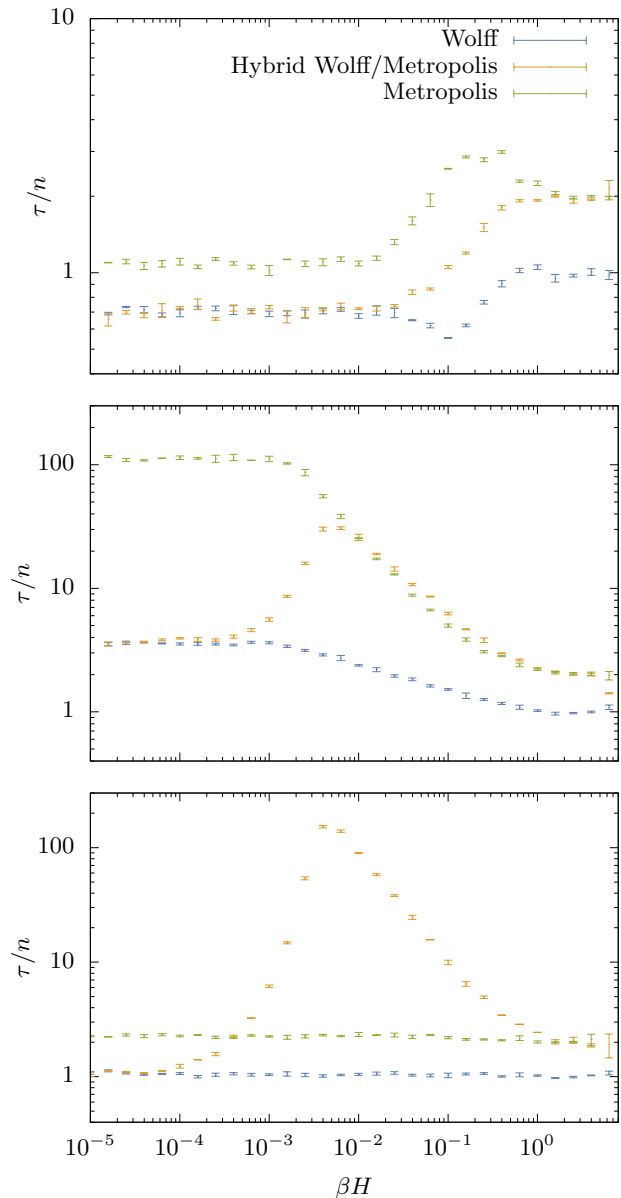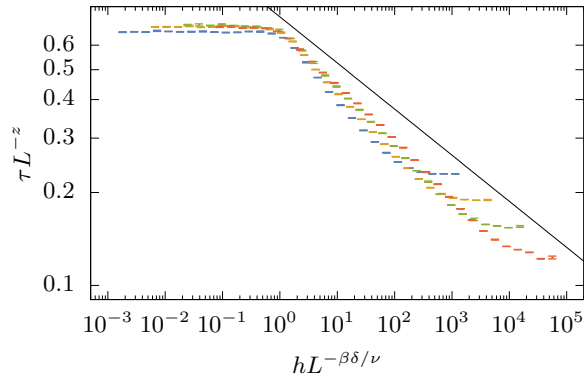

FIG. 1. The autocorrelation time τ of the internal energy $\mathcal{H}$ for a $n = 32 \times 32$ square-lattice Ising model simulated by the three nonzero field methods detailed in this paper. The top plot is for simulations at high temperature, $T = 5.80225$, the middle plot is for simulations at the critical temperature, $T = 2.26919$, and the bottom plot is for simulations at

is observed at low temperature. This is because in the low temperature phase the ergodic hypothesis—that the time-average value of observables is equal to their ensemble average—is violated. As the system size grows the likelihood of a fluctuation in any reasonable dynamics that flips the magnetization from one direction to the other becomes vanishingly small, and therefore it is inappropriate to estimate expected values in the low temperature phase by averaging over the whole configuration

space. Instead, values must be estimated by averaging over the portion of configuration space that is accessible to the dynamics.

For finite size systems, like any we would simulate, dynamics at zero field or even small nonzero field do allow the whole configuration space to be explored. However, people usually want to use the results from finite size systems to estimate the expected values in the thermodynamic limit, where this is no longer true. At zero field, for instance, it is common practice to use $\langle |M| \rangle$ to estimate the expected value for the magnetization instead of $\langle M \rangle$. But what to do at finite field? Is this approach justified?

The provenance of the restricted configuration space in the thermodynamic limit is the fact that, at criticality or for nonzero H under the critical temperature, an infinite cluster forms and can never be flipped by the dynamics. In a finite system there is no infinite cluster, and even the largest cluster has a nonzero probability of being flipped. However, we still want to be able to use finite systems to estimate the quantities of infinite ones. We approach this by thinking of a finite system as a piece of an infinite one. Clearly, time-averages of quantities in the chunk are equal to time-averages of the whole system, simply with much slower convergence. However, simulation of finite systems present a problem: in each sample, we are given a snapshot of a chunk of the system but do not know which direction the infinite cluster in the surrounding infinite system is pointing. Therefore, we must make a guess such that, in the limit of larger and larger system size, the probability that our guess is wrong and that the infinite cluster is pointed in a direction opposite the one we expected approaches zero. To estimate the equilibrium values of quantities in the infinite system, we only admit values resulting from states whose corresponding infinite cluster is expected to point in the direction of our external field. If the external field is zero, we choose one direction at random, say positive.

Snapshots of the chunk what we suspect are likely to be a piece of an infinite cluster facing opposite the external field are discarded for the purpose of computing equilib-

rium values of the system, but still represent something physical. When the model's state contains an infinite cluster oriented against the external field, it is known as *metastable*. Therefore in a finite simulation one can make convergent estimates of quantities in both the equilibrium and metastable states of the infinite system by dividing sampled states into two sets: those that are likely snapshots of a finite block in an infinite system whose infinite cluster is oriented with the field, and those whose infinite cluster is oriented against the field. All we must do now is specify how to distinguish between likely oriented with and likely oriented against. Others have used the direction of the largest cluster in the finite system to make this distinction. Far from the critical point in the low-temperature phase this does a clear and consistent job. However, we would like to study systems very near the critical point. We will take the likely direction of the infinite cluster to be given by *the direction in which the system is magnetized*. Clearly, as our finite chunk grows larger and larger, the probability that its magnetization is different from the direction of the infinite cluster goes to zero. In the thermodynamic limite

Since, in the thermodynamic limit expected values are given by an average over a restricted configuration space, we can estimate those expected values at finite size by making the same restriction. Defining the reduced configuration spaces

$$S_e^n = \{s \in S^n \mid \text{sgn}(H)M(s) > 0\} \qquad (12)$$
$$S_m^n = \{s \in S^n \mid \text{sgn}(H)M(s) < 0\} \qquad (13)$$
$$S_0^n = \{s \in S^n \mid \text{sgn}(H)M(s) = 0\} \qquad (14)$$

where

$$\text{sgn}(H) = \begin{cases} 1 & H \geq 0 \\ -1 & H < 0. \end{cases} \qquad (15)$$

Clearly, $S^n = S_e^n \cup S_m^n \cup S_0^n$, which implies that $Z = Z_e + Z_m + Z_0$. Since $S_e^n = \{-s \mid s \in S_m^n\}$, $|S_e^n| = |S_m^n|$. At zero field, $\mathcal{H}(s) = \mathcal{H}(-s)$, and therefore $Z_e = Z_m$. This can be used to show the near-equivalence (at zero field) of taking expectation values $\langle |M| \rangle$ of the absolute value of the magnetization and taking expectation values $\langle M \rangle_e$ of the magnetization on a reduced configuration space, since

$$\langle |M| \rangle = \frac{1}{1 + Z_0/2Z_e} \langle M \rangle_e \qquad (16)$$

At infinite temperature, $Z_0/Z_e \simeq n^{-1/2} \sim L^{-1}$ for large L, n. At any finite temperature, especially in the ferromagnetic phase, $Z_0 \ll Z_e$ in a much more extreme way.

If the ensemble average over only positive magnetizations can be said to converge to the equilibrium state of the Ising model in the thermodynamic limit, what of the average over only *negative* magnetizations, or the space S_m^n, ordinarily unphysical? This is, in one sense, precisely

the definition of the metastable state of the Ising model. Expected values of observables in the metastable state can therefore be computed by considering this reduced ensemble.

How, in practice, are these samplings over reduced ensembles preformed? There may be efficient algorithms for doing Monte Carlo sampling inside a particular reduced ensemble, but for finite-size systems with any of the algorithms described here the whole configuration space is available. Any of these can be used, however, to simultaneously sample all reduced configuration spaces by allowing them to sample the whole space and then only add a contribution to a given average if the associated state is in the reduced space of interest.

Thanks!

---

[1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, The journal of chemical physics **21**, 1087 (1953).
[2] W. K. Hastings, Biometrika **57**, 97 (1970).
[3] U. Wolff, Physical Review Letters **62**, 361 (1989).
[4] U. Wolff, Physics Letters B **228**, 379 (1989).
[5] C.-W. Liu, A. Polkovnikov, and A. W. Sandvik, Physical Review B **89**, 054307 (2014).
[6] R. H. Swendsen and J.-S. Wang, Physical review letters **58**, 86 (1987).
[7] C. J. Geyer, Statistical science , 473 (1992).