Journal of Computational Physics 250 (2013) 1–12

Contents lists available at SciVerse ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Computing the partition function, ensemble averages, and density of states for lattice spin systems by sampling the mean

Dirk Gillespie *

Department of Molecular Biophysics and Physiology, Rush University Medical Center, Chicago, IL, United States

ARTICLE INFO

Article history:
Received 16 August 2012
Received in revised form 15 February 2013
Accepted 1 May 2013
Available online 20 May 2013

Keywords:
Lattice spin systems
Partition function
Ensemble averages
Random sampling

ABSTRACT

An algorithm to approximately calculate the partition function (and subsequently ensemble averages) and density of states of lattice spin systems through non-Monte-Carlo random sampling is developed. This algorithm (called the sampling-the-mean algorithm) can be applied to models where the up or down spins at lattice nodes interact to change the spin states of other lattice nodes, especially non-Ising-like models with long-range interactions such as the biological model considered here. Because it is based on the Central Limit Theorem of probability, the sampling-the-mean algorithm also gives estimates of the error in the partition function, ensemble averages, and density of states. Easily implemented parallelization strategies and error minimizing sampling strategies are discussed. The sampling-the-mean method works especially well for relatively small systems, systems with a density of energy states that contains sharp spikes or oscillations, or systems with little *a priori* knowledge of the density of states.

1. Introduction

Lattice spin models are classic models of many physical phenomena, dating back at least to Ising in 1925 [1] and Bethe in 1931 [2]. Since then, these models have been applied to magnetism [1,2], phase transitions [3,4], quantum mechanics [5], and solid state physics [6,7]. In this paper, a model from biology is introduced. In these models, N lattice nodes have spins (either up or down) that interact to change the spins of their neighbors.

Despite their age, lattice spin models continue to be useful tools in many fields, with thousands of recent papers using Ising models alone. Therefore, there continues to be interest in efficient numerical techniques to simulate traditional lattice spin models (e.g., [8,9]). In addition, there is a growing interest in lattice spin models from biology. In particular, there are a number of vital phenomena that involve flipping the state of proteins based on their neighbors' states which can be modeled as lattice spin systems. These include calcium-induced calcium release that creates Ca^{2+} sparks and Ca^{2+} puffs [10], the propagation of calcium waves in arrhythmic heart muscle [11,12], and action potential initiation and propagation both within a neuron [13] and across connexin-connected heart cells [14]. Many of these systems differ from traditional lattice spin system in that they involve a relative small number of lattice nodes (sometimes <100) and do not have a physically-defined Hamiltonian (although it is possible to define a pseudo-Hamiltonian). Moreover, the interactions are not only between nearest neighbors. In fact, because many biological systems involve diffusing ions as the interaction mechanism, the interactions can be very long-ranged. One of the things shown in this paper is that lattice spin systems with these properties can behave differently compared to traditional ones and that a new computational tool can be effective in studying them (and possibly some traditional lattice spin systems as well).

* Tel.: +1 3129423089.
E-mail address: dirk_gillespie@rush.edu

0021-9991/$ - see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.jcp.2013.05.001

The most common numerical technique is Monte Carlo (MC) simulation where the spins of randomly-chosen nodes are flipped to determine thermodynamic properties [15]. In this paper, random sampling is used in a different way to approximately calculate the partition function of the system, a quantity Metropolis–Hastings MC algorithms [16–18] explicitly avoid. This is done by rearranging the partition sum into $N+1$ terms by considering only states with exactly n ($n = 0, 1, \ldots, N$) up spins. Each term includes a mean (average) quantity that can be approximated using random sampling. It is therefore called the sampling-the-mean algorithm. After all the $N+1$ terms are computed separately, they are added together. With the same algorithm, it is also possible to approximately calculate ensemble averages and the density of energy states. Moreover, this highly parallelizable algorithm gives explicit estimates for the error of the approximation.

The sampling-the-mean algorithm is in many ways different from most MC methods. First, sampling-the-mean is based on the Central Limit Theorem of probability theory [19–21]; MC simulations are based on microscopic reversibility [18]. Secondly, there is no wasted computational effort; all the calculations are used, whereas in MC simulations many samples are rejected (with the Wang–Landau algorithm [22] being a notable exception). Lastly, all the random samples in the sampling-the-mean method are independent of each other; in MC simulations, the last accepted state is perturbed for the next attempted state, which sometimes can lead to the simulations being stuck in local energy minima.

However, the sampling-the-mean algorithm is also in many ways similar to some MC simulation techniques. For example, by holding the number of up spins constant for each set of random samples, the sampling-the-mean technique is similar to fixed-magnetization sampling [23] like that of Kawasaki [24] that interchanges the neighboring spins to keep the magnetization of the entire lattice constant [15].

The sampling-the-mean algorithm also shares some similarities with the Wang–Landau MC technique [22] in that it uses the density of states to compute thermodynamic properties. Specifically, Wang–Landau computes the number of states with energy in the interval $E + \Delta E_{\mathrm{WL}}$, the density of energy states denoted $g(E)$. The partition function can then be computed by

$$Z = \sum_{E} g(E) e^{-E/kT}, \tag{1}$$

where k is the Boltzmann constant and T the absolute temperature. Sampling-the-mean inverts this by picking sample states with a fixed number of up spins for which the density of up-spin states is known beforehand. Results are then combined in a sum similar to Eq. (1). Moreover, as will be shown, that the sampling-the-mean algorithm can also be used to compute the density of energy states.

In this paper, the sampling-the-mean algorithm is tested mainly on a new lattice spin model from biology. This model has several of the properties described above that make it different from many standard lattice spin models. First, it does not have a physically defined Hamiltonian so the transition from one state to another is not given by a Boltzmann factor. Second, the interaction potential between nodes is not additive. Lastly, the interaction of the nodes is not only between nearest neighbors, but is long-ranged across the entire lattice. The model was chosen partly because of these properties to show that the sampling-the-mean algorithm does not rely on any special properties like nearest neighbor interactions. This model also highlights differences in the density of energy states between lattice spin models with long-range interactions and traditional nearest-neighbor models. Sampling-the-means appears to be more effective for the former, as compared to Wang–Landau.

In the next section, it is shown how the Central Limit Theorem can be used to (approximately) compute the partition function and ensemble averages and how the error in the calculation can be reduced most efficiently for a given number of random samples. Convergence and error trends are then analyzed for the biological test system to show that they are consistent with the predictions of the Central Limit Theorem. Later, sampling-the-mean is used to compute the density of energy states. When compared head-to-head with the Wang–Landau method, it is shown that the sampling-the-mean method is more efficient for relatively small systems ($<\sim$400 nodes), especially for the test system because it produces a density of states $g(E)$ that is both discrete with large gaps between narrow spikes and highly oscillatory.

2. Theory and methods

2.1. General theory: partition function

For a general lattice spin system, we first consider the state of the system α, defined by the vector $\alpha = (\alpha_1, \ldots, \alpha_N)$ representing the spins of the N lattice nodes. The probability of the system being in state α is denoted P_α. Since the partition function Z is not known *a priori*, one can only calculate a quantity Q_α that is proportional to P_α. Then,

$$P_\alpha = Q_\alpha / Z, \tag{2}$$

where the partition function is

$$Z = \sum_{\alpha} Q_\alpha, \tag{3}$$

with the sum going over all possible states α of the system. For a system with a Hamiltonian $H(\alpha)$,

$$Q_\alpha = e^{-H(\alpha)/kT}, \tag{4}$$