ARTICLE IN PRESS

Computer Physics Communications ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Computer Physics Communications

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

Determining efficient temperature sets for the simulated tempering method

A. Valentim, M.G.E. da Luz, Carlos E. Fiore [*]

Departamento de Física, Universidade Federal do Paraná, CP 19044, 81531-980 Curitiba-PR, Brazil

ARTICLE INFO

Article history:
Received 4 February 2014
Received in revised form
10 April 2014
Accepted 14 April 2014
Available online xxxx

Keywords:
Strong first-order phase transitions
Simulated tempering
Monte Carlo methods
Replica temperatures optimal values

ABSTRACT

In statistical physics, the efficiency of tempering approaches strongly depends on ingredients such as the number of replicas R, reliable determination of weight factors and the set of used temperatures, $\mathcal{T}_R = \{T_1, T_2, \dots, T_R\}$. For the simulated tempering (ST) in particular – useful due to its generality and conceptual simplicity – the latter aspect (closely related to the actual R) may be a key issue in problems displaying metastability and trapping in certain regions of the phase space. To determine $\mathcal{T}_R$'s leading to accurate thermodynamics estimates and still trying to minimize the simulation computational time, here a fixed exchange frequency scheme is considered for the ST. From the temperature of interest T_1, successive T's are chosen so that the exchange frequency between any adjacent pair T_r and T_{r+1} has a same value f. By varying the f's and analyzing the $\mathcal{T}_R$'s through relatively inexpensive tests (e.g., time decay towards the steady regime), an optimal situation in which the simulations visit much faster and more uniformly the relevant portions of the phase space is determined. As illustrations, the proposal is applied to three lattice models, BEG, Bell–Lavis, and Potts, in the hard case of extreme first-order phase transitions, always giving very good results, even for $R = 3$. Also, comparisons with other protocols (constant entropy and arithmetic progression) to choose the set $\mathcal{T}_R$ are undertaken. The fixed exchange frequency method is found to be consistently superior, specially for small R's. Finally, distinct instances where the prescription could be helpful (in second-order transitions and for the parallel tempering approach) are briefly discussed.

1. Introduction

Keystone in the study of statistical physics problems, numerical methods are generally expected to fulfill two requirements: (i) first (and surely the most important), to yield precise estimates for the thermodynamical quantities analyzed; (ii) second, to be as simple and as fast as possible in their implementations.

Nevertheless, often the mentioned two requisites strike out in opposite directions. Indeed, consider, e.g., systems in the regime of phase transitions whose distinct regions of the phase space are separated by large free-energy barriers. This is a common situation not only for complex problems like spin glasses, protein folding, and biomolecules conformation [1–5], but also in lattice gas models displaying first-order phase transitions [6,7]. In all such examples there may be the occurrence of metastability [6,7]. Thus, when simulated, these systems can get trapped into local minima. Ways to circumvent this technical difficulty should demand more sophisticated evolution dynamics procedures and longer computational times.

Different proposals like (a) cluster [8], (b) multicanonical [9], (c) Wang–Landau [10], and (d) tempering [11,12], among others, are relevant algorithms trying to maintain a good balance between features (i) and (ii) above. In particular, (d) above relies on the straightforward idea of “heating up” the system to higher temperatures, so as to help it to cross the barriers at low temperatures. Moreover, tempering methods have attracted large interest due to their generality with a broad applicability [13].

There are two major formulations for the tempering approach, namely, parallel (PT) [11] and simulated (ST) [12], where always the start point is to choose a set of R distinct temperatures (with $T_r < T_{r+1}, r = 1, \dots, R-1$), $\mathcal{T}_R = \{T_1, T_2, \dots, T_R\}$, in which T_1 is the one of interest. In the PT, configurations from the distinct R replicas (running in parallel) at the different T's are exchanged. For the ST, a single realization undergoes many temperature changes (among the T's in $\mathcal{T}_R$). Thus, the temperature itself is a dynamical variable.

Each tempering implementation presents its own characteristics and advantages, as recently discussed in detail in Ref. [14] (see also the Refs. therein). In particular, although the ST has a higher

---

[*] Corresponding author.
E-mail addresses: luz@fisica.ufpr.br (M.G.E. da Luz), fiore@fisica.ufpr.br (C.E. Fiore).

http://dx.doi.org/10.1016/j.cpc.2014.04.010

Please cite this article in press as: A. Valentim, et al., Determining efficient temperature sets for the simulated tempering method, Computer Physics Communications (2014), http://dx.doi.org/10.1016/j.cpc.2014.04.010

probability than the PT to exchange temperature [15–18,14], it displays a less frequent tunneling between coexisting phases [14]. Hence, the ST requires large computational time for generating uncorrelated configurations, thus a slower convergence to the steady equilibrium. On the other hand, for proper estimates (at least at first-order phase transition regimes) the PT needs non-adjacent switch of temperatures, making the procedure a bit more involving—an implementation not necessary for the ST.

Furthermore, for the hard to treat case of strong discontinuous phase transitions, promising extensions for tempering methods have been proposed. In particular, the PT combined with modified ensembles (as multiple Gaussians [19,20]) comprise the so called generalized replica exchange approaches [21,22]. They have been applied with great success to problems like solid–liquid [23] and vapor–liquid transitions [24]. Also, enhancements for the usual ST are possible. Examples are (a) to consider for it modified distributions [25], leading to very good results for both lattice (e.g., Potts and Ising) and continuum (Lennard-Jones clusters) models; and (b) besides T to assume another dynamical variable, as the external field [26], quite helpful in dealing with crossovers in 2D Ising systems.

Thus, it would be desirable to improve the efficiency of the ST still preserving its positive aspects, notably the procedure simplicity. As a hint to do so, the previous comments indicate that a central point in the ST method is less the probability of a single attempt to exchange temperatures $T_r \rightarrow T_{r+n}$ (with $n = 1$) and more the overall frequency in which the different system phases are visited. Therefore, one should try to optimize the set $\mathcal{T}_R$ as a whole, investigating how the combination of the different transitions would speed up the convergence to the steady state (by a more uniform sampling of the microscopic configurations).

For the ST we then propose here a rather direct protocol to select $\mathcal{T}_R$ by means of a fixed exchange frequency (FEF) prescription. Given R, it consists in determining the T_r's such that the exchange frequency between any pair of adjacent temperatures is f. From simple preliminary tests we verify if the obtained set leads to an appropriate tunneling between coexisting phases. If this is not the case, another value of f is chosen, a new $\mathcal{T}_R$ is calculated, and the tests repeated. With relatively low computational effort (see the next section), we end up with a very efficient $\mathcal{T}_R$ for the full simulations. Through examples, we furthermore show that this optimal $\mathcal{T}_R$ works well for other values of the considered parameters and not only for the specific values employed in the set derivation. The same $\mathcal{T}_R$ can also be used in the vicinity of the original parameters' values as well as for other system sizes. Hence, in many applications $\mathcal{T}_R$ needs to be determined just once. We compare the FEF with other schemes to select the T_r's. We find that the present is not only superior to more simple recipes (like arithmetic progressions (AP)) but also to more physically oriented selection methods (like the constant entropy (CE) [27,28]). We finally confirm a somehow expected result (but not fully investigated in the literature) that the exact distribution of temperatures in $\mathcal{T}_R$ becomes less relevant as R increases.

As illustrations, we address three distinct systems, thus exploiting a relatively larger variety of first-order phase transition features. One is the Potts model, an ideal case test. For large q's, it presents strong discontinuous transitions (the regime we shall focus on), whose temperatures are exactly known. The others are the BEG and BL models, likewise interesting not only by displaying more complex phase diagrams than the Potts (e.g., having phases with distinct structural properties), but also for already being extensively analyzed through the PT and ST approaches [14,28,29]. Thus, all of them are nice examples to check for the reliability of the proposed protocol.

The work is organized as the following. In Section 2 we review the ST approach and how to characterize first-order phase transitions at low T's (the context we consider in this contribution). We also discuss in full detail the FEF protocol. In Section 3 we analyze the BEG, Bell–Lavis (BL), and Potts lattice models. For the BEG and BL we also compare the FEF results with those for two other schemes (AP and CE) and illustrate the methods' performance dependence on the number of replicas R. Lastly, we present final remarks and conclusion in Section 4.

2. The method details

In general, for systems displaying first-order transitions at low temperatures or with a large jump in the order parameter [30], the distinct coexisting phases are separated by large free-energy barriers, exhibiting trapping and metastable states. Hence, such cases are interesting instances to test the proposed scheme. So, next we first give a brief account of the ST method and discuss an appropriate way to analyze strong first-order phase transitions. Then, we pass to describe a FEF procedure for the ST method.

2.1. The simulated tempering (ST)

The ST follows a twofold procedure. First, at a certain T_r (and during a given number of MC steps, defined as the number of lattice sites), a standard Metropolis prescription evolves a system of Hamiltonian $\mathcal{H}$ throughout the phase space microstates $\{\sigma\}$. Second, an attempt for the change $T_{r'} \rightarrow T_{r''}$ (with $r', r'' = 1, 2, \ldots, R$; $\beta_r = 1/(k_B T_r)$; and σ the system state at the attempt time step) is drawn from

$$p_{r' \rightarrow r''} = \min\{1, \exp[(\beta_{r'} - \beta_{r''})\mathcal{H}(\sigma) + (g_{r''} - g_{r'})]\}. \qquad (1)$$

This scheme is repeated a large enough number of times.[1] Also, we consider only adjacent exchanges, i.e., $\Delta r = |r'' - r'| = 1$.

According to Eq. (1), the transition probability $p_{r' \rightarrow r''}$ strongly depends on the temperatures difference. Larger $\beta_{r'} - \beta_{r''}$ leads to lower acceptance probabilities, whereas lower $\beta_{r'} - \beta_{r''}$, although enhancing the exchanges, may not be efficient since the generated configurations at $T_{r''}$ in general will be similar to those at $T_{r'}$. Therefore, conceivably there is a compromise between opposite factors, implying the existence of a best set $\mathcal{T}_R$.

Finally, we comment that in some ST implementations, the correct weights $g_r = \beta_r f_r$ (with f_r the free energy) – whose role is to ensure a uniform visit to the distinct T's – are approximated [15,32]. For our examples, we obtain the g's exactly by means of the approach in Refs. [33,29]. In short (full details in Refs. [33,14,29,34]), suppose a lattice model composed of K layers of L sites each. The total number of sites (or the volume) is then $V = L \times K$. Also, assume the full Hamiltonian written in terms of these layers as

$$\mathcal{H} = \sum_{k=1}^{K} \mathcal{H}(S_k, S_{k+1}), \qquad (2)$$

where $S_k \equiv (\sigma_{1,k}, \sigma_{2,k}, \ldots, \sigma_{L,k})$ denotes the kth layer state configuration and $S_{K+1} = S_1$ (periodic boundary conditions). The transfer matrix $\mathbf{T}$ is defined in such a way that its elements are $\mathbf{T}(S_k, S_{k+1}) = \exp[-\beta \mathcal{H}(S_k, S_{k+1})]$. Thus, in the thermodynamic limit (achieved already at relatively small V's [29]) $f_r = -\ln[\lambda^{(r)}]/(\beta_r L)$, with

$$\lambda^{(r)} = \left. \frac{\langle \mathbf{T}(S_k, S_{k+1} = S_k) \rangle}{\langle \delta_{S_k, S_{k+1}} \rangle} \right|_{\beta = \beta_r}. \qquad (3)$$

---

[1] The appropriate number of MC steps, necessary to ensure a correct sampling, may depend on distinct features [31], such as the system temperature and size, the phase space "complexity", etc. Since there is no standard and direct procedure to find it, a rough estimation of such a time can be obtained by studying the time convergence τ of the order parameter towards the steady value (say, starting from a fully ordered configuration). Thus, throughout this work, from simple simulations we obtain τ and then take as our number of MC steps $10 \times \tau$ (detailed discussions in Section 3).

Please cite this article in press as: A. Valentim, et al., Determining efficient temperature sets for the simulated tempering method, Computer Physics Communications (2014), http://dx.doi.org/10.1016/j.cpc.2014.04.010