Available online at www.sciencedirect.com

ScienceDirect

Physics Procedia 68 (2015) 125 – 129

Physics Procedia

28th Annual CSP Workshop on "Recent Developments in Computer Simulation Studies in Condensed Matter Physics", CSP 2015

Accelerating the convergence of replica exchange simulations using Gibbs sampling and adaptive temperature sets

Thomas Vogel, Danny Perez

Theoretical Division (T-1), Los Alamos National Laboratory, Los Alamos, NM 87545, U.S.A.

Abstract

We recently introduced a novel replica-exchange scheme in which an individual replica can sample from states encountered by other replicas at any previous time by way of a global configuration database, enabling the fast propagation of relevant states through the whole ensemble of replicas. This mechanism depends on the knowledge of global thermodynamic functions which are *measured* during the simulation and not coupled to the heat bath temperatures driving the individual simulations. Therefore, this setup also allows for a continuous adaptation of the temperature set. In this paper, we will review the new scheme and demonstrate its capability. The method is particularly useful for the fast and reliable estimation of the microcanonical temperature $T(U)$ or, equivalently, of the density of states $g(U)$ over a wide range of energies.

Peer-review under responsibility of The Organizing Committee of CSP 2015 Conference

Keywords: Replica exchange; Molecular dynamics; Gibbs sampling; Adaptive temperature set; Multicanonical; Microcanonical temperature

1. Introduction

Simulating physical systems exhibiting multifaceted structural transitions (encountered, for example, in many protein folding problems) or containing extremely hard to find but thermodynamically important states (e.g., crystalline states in hard materials) is a notorious challenge. Replica-exchange (RE) and multicanonical (MUCA) sampling schemes have both been applied, in conjunction with Monte Carlo (MC) and molecular dynamics (MD) methods, to address this issue. They enable the investigation of systems at multiple or varying temperatures, such facilitating the exploration of the configurational space. One of the main practical challenges in the case of RE is the determination of good temperature sets for the heat baths. To this day, many sophisticated schemes have been proposed to solve this problem, see, e.g., Rathore et al. (2005); Katzgraber et al. (2006); Patriksson and van der Spoel (2008); Guidetti et al. (2012); Ballard and Wales (2014). Another limitation of generic RE schemes is that exchange of configurations obviously does not, on its own, change the ensemble of configurations present in the simulation. This is not optimal if the simulation is not converged yet, e.g., if some replicas have yet to locate all thermodynamically relevant states.

E-mail address: {tvogel;danny_perez}@lanl.gov

Peer-review under responsibility of The Organizing Committee of CSP 2015 Conference
doi:10.1016/j.phpro.2015.07.121

In this case, some sort of population control, where "good" configurations can be multiplied and "irrelevant" ones deleted, can help (see Hsu and Grassberger (2011) for an example).

MUCA simulations offer a different strategy to efficiently sample configuration space. In the case of MD, MUCA sampling simply consists of rescaling the interatomic forces: $f_{\mathrm{muca}} = (T_0/T(U))\, f_{\mathrm{can}}$ (Hansmann et al., 1996; Junghans et al., 2014), where $f_{\mathrm{can}} = -\nabla U$ are the conventional forces and $T(U)$ is the microcanonical temperature. The difficulty lies in the estimation of this *a priori* unknown function. Metadynamics (Laio and Parrinello, 2002; Dama et al., 2014) and statistical temperature molecular dynamics (STMD) (Kim et al., 2006) were successfully applied to this problem. However, these methods can show dynamical artifacts like hysteresis when the system is driven over first-order like transitions, for example. Further, even with the knowledge of $T(U)$, the diffusion of a MUCA simulation in energy space can be very slow. However, a very powerful side-effect of MUCA simulations is that, once obtained, $T(U)$ can be used to infer the density of states, hence providing a wealth of thermodynamic information.

In the following, we review a generalized RE method, the Gibbs-sampling enhanced RE, which addresses some of the issues with generic RE and MUCA simulations. By uncoupling the measured microcanonical temperature from the heat-bath temperatures which drive the simulations we can introduce powerful conformational mixing steps and continuously adapt the heat-bath temperatures for a complete coverage of the whole energy range. We demonstrate this ability by presenting a particular adaptive simulation which is set up so that only a tiny part of the phase space in covered the beginning. Note that $T(U)$ is also available at the end of such a simulation, providing an alternative to conventional MUCA approaches.

2. Method review

We consider a RE scheme (Geyer, 1991) in an expanded-canonical ensemble (Lyubartsev et al., 1992) with reference temperatures T_0^i. In contrast to other sampling methods where thermodynamic functions are continuously adapted (Wang and Landau, 2001; Laio and Parrinello, 2002; Junghans et al., 2014), the basic idea in our method (Vogel and Perez, 2015) is to *measure* the microcanonical temperature while running canonical simulations, hence decoupling the measurements from the simulation-driving heat bath temperatures. We use an estimator for the microcanonical temperature based the microcanonical averages of time-derivatives of the product between the normalized energy gradient $\eta = \nabla U/(\nabla U \nabla U)$ and the particle momenta p:

$$g = \frac{\mathrm{d}F(U)}{\mathrm{d}U} = -\left\langle \frac{\mathrm{d}}{\mathrm{d}t}(\eta \cdot p) \right\rangle , \tag{1}$$

where F is the free energy, a relation which was derived in a more general form in the context of the adaptive biasing force (ABF) method (Darve et al., 2008). The time derivatives are estimated in an additional, constant-energy MD time step. As defined, g relates to the microcanonical temperature as:

$$T_{\mathrm{ABF}}(U) = \frac{T_0}{1-g} . \tag{2}$$

During the simulation, all replicas contribute measurements to the global calculation of $T_{\mathrm{ABF}}(U)$, while they also collectively fill a global conformational database.[1] An instantaneous estimator for the (unnormalized) entropy $S(U)$ can then be obtained:

$$S(U) = \int_{U_0}^{U} T_{\mathrm{ABF}}^{-1}(U')\,\mathrm{d}U' + C . \tag{3}$$

With $g(U) = \exp[k_B^{-1} S(U)]$ being the density of states, canonical energy probability distributions $P_i(U) \propto g(U)\, \mathrm{e}^{-\beta_i U}$ can be calculated. The key is that these distributions calculated from global data are different from those that would be inferred from the data gathered by individual walkers alone. For example, $P_i(U)$ can account for the presence of some low energy state, even though walker i might never have sampled it himself. This allows for the introduction of a rejection-free, global Monte Carlo move facilitating the propagation of important states through the simulation: at

[1] The management of all global data is done by an additional master process which does not perform a MD run himself.