The infinite swapping limit for parallel tempering

Large Deviation Methods for the Analysis and Design of Monte Carlo Schemes in Chemistry

Paul Dupuis
Brown University
Providence, RI 02912

Abstract

We consider the development of efficient simulation methods for the equilibrium distributions of high dimensional complex systems arising from chemistry and physics, and discuss an approach based on the use of large deviation rate functions to quantify performance.

Parallel tempering (also known as replica exchange sampling), is one of the most widely used methods for simulating complex systems. Simulations are conducted in parallel for a set of temperatures, and the key improvement over standard Monte Carlo is a swap mechanism that exchanges configurations between these parallel simulations at a given rate. The mechanism is designed to allow the low temperature system to escape from deep local energy minima where it might otherwise be trapped, via those swaps with the higher temperature components.

Using large deviation theory, we show that the rate of convergence of the empirical measure is a monotone increasing function of the swap rate. We also describe how one can construct a simulation scheme that is equivalent to the limit of the parallel tempering schemes in the sense of distributions, but which involves no swapping at all. With this scheme, which we call *infinite swapping* (INS), the effect of the swapping is captured by a collection of weights that influence both the dynamics and the empirical measure.

While the INS scheme optimizes the convergence rate, it has practical limitations in implementation due to the complexity of the weights when the number of temperatures is large (>7). We present an approximation to the full infinite swapping which is based on alternating between *partial infinite swapping* (PINS) algorithms, which can be shown to approximate (theoretically and practically) the INS scheme. Numerical studies on fairly complex Lennard-Jones systems are presented, and improvements of three orders of magnitude in performance over conventional parallel tempering are observed at an increased computational cost of 5-15%.