

Virtual-Move Parallel Tempering

Ivan Coluzza* and Daan Frenkel^[a]

The exponential increase in the speed of computers during the past few decades has made it possible to perform simulations that were utterly unfeasible one generation ago. But in many cases, the development of more efficient algorithms has been at least as important.

One of the most widely used schemes to simulate many-body systems is the Markov-chain Monte Carlo method (MCMC) that was introduced in 1953 by Metropolis et al.^[1] In this algorithm the average properties of a system are estimated by performing a random walk in the configurational space, where each state is sampled with a frequency proportional to its Boltzmann weight. In the Metropolis algorithm, this is achieved by attempting random moves from the current state of the system to a new state. Depending on the ratio of the Boltzmann weights of the new and old states, these trial moves may be either accepted or rejected. Metropolis et al.

[a] *Dr. I. Coluzza, Prof. Dr. D. Frenkel*
FOM Institute for Atomic and Molecular Physics
Kruislaan 407, 1098 SJ Amsterdam (The Netherlands)
Fax: (+31) 31-20-6684106
E-mail: coluzza@amolf.nl

showed that the acceptance probability of trial moves can be chosen such that Boltzmann sampling is achieved.

One important application of the MCMC method is the estimation of the Landau free energy F of the system given in Equation (1) as function of some order parameter:

$$F(Q) = -kT[\ln P(Q)] \quad (1)$$

There are many situations where the MCMC method does not yield an accurate estimate of F , because it fails to explore configuration space efficiently. This is, for instance, the case in "glassy" systems that tend to get trapped for long times in small pockets of configuration space. In the early 1990s the so-called parallel-tempering (PT) technique was introduced to speed up the sampling in such systems.^[3,6]

In a parallel-tempering Monte Carlo (MC) simulation, n simulations of a particular model system are carried out in parallel at different temperatures (or at different values of some other thermodynamic field, such as the chemical potential or a biasing potential). Each of these copies of the system is called a replica. In addition to the regular MC trial moves, one occasionally attempts to swap the temperatures of a pair of these systems (say i and j). The swapping move between temperatures i and j is accepted or rejected according to a criterion that guarantees detailed balance, for example, see Equation (2):

$$P_{\text{acc}}(ij) = \frac{e^{\Delta\beta_{ij}\Delta E_{ij}}}{1 + e^{\Delta\beta_{ij}\Delta E_{ij}}} \quad (2)$$

where $\Delta\beta_{ij}$ is the difference of the inverse of swapping temperatures, and ΔE_{ij} is the energy difference of the two configurations. Although there are other valid acceptance rules, we used the one in Equation (2) because it was easy to implement.

To facilitate the sampling of high free-energy states ("difficult" regions), we used adaptive umbrella sampling.^[8,9] In this (iterative) scheme, a biasing potential is constructed using the histogram of the states, sampled during an iteration as follows in Equation (3):

$$W_l(Q,T) = W_{l-1}(Q,T) - a \ln[P_l(Q)] \quad (3)$$

where W is the biasing potential function of an order parameter Q , l is the iteration number, a is a constant that controls the rate of convergence of W (a typical value for a is 0.05), and T is the temperature. After iteration, W converges to the Landau free energy. As a consequence, $P(Q) \sim \exp[-\beta F(Q)] \exp[-W(Q)]$ becomes essentially flat and the biased sampling explores a larger fraction of the configuration space. During the MC sampling we include the bias, and only at the end of the simulation do we compute the free energy $F(Q)$ from Equation (4):

$$F(Q) = -kT[\ln P(Q) + W(Q,T)] \quad (4)$$

where $P(Q)$ is the probability of observing a state characterized

by the order parameter Q , and $W(Q,T)$ is the biasing potential of the last iteration computed at temperature T . Combined with parallel tempering, the acceptance rule for the temperature swapping move is then described by Equations (5) and (6):

$$\text{acc}_{ij} = \frac{e^{\Delta\beta_{ij}\Delta E_{ij} + \Delta W_{ij}}}{1 + e^{\Delta\beta_{ij}\Delta E_{ij} + \Delta W_{ij}}} \quad (5)$$

$$\Delta W_{ij} = W_l(Q_i, T_j) - W_l(Q_j, T_j) + W_l(Q_j, T_i) - W_l(Q_i, T_i) \quad (6)$$

where i and j are replica indices, and l is the iteration number. We refer to this scheme as APT (adaptive parallel tempering).^[10,11]

In the conventional MCMC method all information about rejected trial moves is discarded. Recently one of us proposed a scheme that makes it possible to include the contributions of rejected configurations in the sampling of averages.^[2] Herein, we show how this approach can be used to increase the power of the parallel-tempering scheme.

In this scheme, we only retain information about PT moves that have been accepted. However, in the spirit of ref. [2], we can include the contribution of all PT trial moves, irrespective of whether they are accepted. The weight of the contribution of such a virtual move is directly related to its acceptance probability. For instance, if we use the symmetric acceptance rule for MC trial moves, then the weights of the original and new (trial) state in the sampling of virtual moves are given by Equations (7) and (8):

$$P_N = \frac{e^{\Delta\beta\Delta E_{0-N} + \Delta W_{0-N}}}{1 + e^{\Delta\beta\Delta E_{0-N} + \Delta W_{0-N}}} \quad (7)$$

$$P_0 = \frac{1}{1 + e^{\Delta\beta\Delta E_{0-N} + \Delta W_{0-N}}} \quad (8)$$

where ΔW_{0-N} is defined in Equation (6). We are not limited to a single trial swap of state i with a given state j . Rather, we can include all possible trial swaps between the temperature state i and all $N-1$ remaining temperatures. Our estimate for the contribution to the probability distribution P_i corresponding to temperature i is then given by the following sum in Equation (8)

$$P_i(Q) = \sum_{j=1}^{N-1} \left(\frac{1}{1 + e^{\Delta\beta_{ij}\Delta E_{ij} + \Delta W_{ij}}} \right) \delta(Q_i - Q) + \sum_{j=1}^{N-1} \left(\frac{e^{\Delta\beta_{ij}\Delta E_{ij} + \Delta W_{ij}}}{1 + e^{\Delta\beta_{ij}\Delta E_{ij} + \Delta W_{ij}}} \right) \delta(Q_j - Q) \quad (9)$$

where the delta functions select the configurations with order parameter Q . As we now combine the parallel-tempering algorithm with a set of parallel virtual moves, we refer to the present scheme as virtual-move parallel tempering (VMPT).

To measure the efficiency of VMPT, we computed the free-energy landscape of a simple lattice-protein model. In this model, interaction with a substrate can induce a conformation-

al change in the proteins. For the same system we had already explored the use of the conventional APT scheme.^[10]

Specifically, the model protein that we consider represents a heteropolymer containing 80 amino acids, while the substrate has a fixed-space arrangement and contains 40 residues (see Figure 1). The configurational energy of the system is defined as Equation (10):

$$E_C = E_{\text{intra}} + E_{\text{inter}} = \sum_i^{N_C} \left[\sum_{j \neq i}^{N_C} C_{ij} S_{ij} + \sum_{j' \neq i'}^{N_S} C_{ij'} S_{ij'} \right] \quad (10)$$

where the indices i and j run over the residues of the protein, while j' runs only over the elements of the substrate; C is the contact defined as Equation (11):

$$C_{ij} = \begin{cases} 1 & \text{if } i \text{ neighbor of } j \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

and S_{ij} is the interaction matrix. For S we use the 20×20 matrix fitted by Miyazawa and Jernigan^[7] on the basis of the frequency of contacts between each pair of amino acids in nature.

We change the identity of the amino acids along the chain by "point mutations" which, in this context, means changes of a single amino acid. In doing so we explore the sequence space of the protein and the substrate, and we minimize at the same time the configurational energy of the system in two distinct configurations, one bound (Figure 1a) and one unbound (Figure 1b). The design scheme is the same as that used in ref. [10]. In this scheme, trial mutations are accepted if the Monte Carlo acceptance criterion is satisfied for both configurations.

The result of the design process is a model protein that has the ability to change its conformation when bound to the substrate. The sampling of the configurations is performed with three basic moves: corner-flip, crankshaft, and branch rotation. The corner-flip involves a rotation of 180° of a given particle around the line joining its neighbors along the chain. The crankshaft move is a rotation by 90° of two consecutive particles. A




Figure 1. Spatial arrangement of the chain in the structures used to test the model (a,b), and intermediate structure (c) ($Q=25$).

Table 1. Simulation parameters used for comparing the VMPT algorithm with the old scheme. In simulation I we used the same parameters for both algorithms. The results in Figure 2 show that VMPT was much more efficient in sampling the free energy. In simulation II, we increased by two orders of magnitude the number of steps of the simulation with APT to obtain a sampling of comparable $F(Q)$ to the one computed using the new VMPT scheme (Figure 4). Execution times computed on an SGI Altix 3700 with Intel Itanium II, 1.3 GHz

Simulation	Temperatures [$k^{-1} T^{-1}$]	Number of iterations	Sampling steps	APT exec. time [s]	VMPT exec. time [s]
I	0.1, 0.125, 0.143, 0.167, 0.2, 0.222, 0.23, 0.25, 0.27, 0.29, 0.31, 0.33, 0.35, 0.37, 0.4, 0.444, 0.5	400	4×10^8	2600	3200
II	0.1, 0.125, 0.143, 0.167, 0.2, 0.222, 0.23, 0.25, 0.27, 0.29, 0.31, 0.33, 0.35, 0.37, 0.4, 0.444, 0.5	1000	2×10^{10}	150000	




Figure 2. Average free energy computed with five runs (10^8 MC steps, Table 1 I) of the old scheme, compared with the result of five VMPT simulations (10^8 MC steps, Table 1 I), at $T=0.1 k^{-1} T^{-1}$. The points with $F=0$ correspond to values of Q that have not been sampled.

branch rotation is a turn, around a randomly chosen pivot particle, of the whole section starting from the pivot particle and going to the end of the chain. For all these moves we use a symmetric acceptance rule [Eq. (12)] with the addition of the biasing potential calculated with the umbrella sampling scheme of Equation (3):

$$\text{acc}_{O \rightarrow N} = \frac{e^{\beta \Delta E_{O \rightarrow N} + \Delta W_{O \rightarrow N}}}{1 + e^{\beta \Delta E_{O \rightarrow N} + \Delta W_{O \rightarrow N}}} \quad (12)$$

where $\Delta E_{O \rightarrow N}$ is the energy difference between the new and the old state [Eq. (10)], and $\Delta W_{O \rightarrow N}$ is the difference in the bias potential from the same states [Eq. (3)]. We sample the free energy as a function of two order parameters, of which the first is the conformational energy defined in Equation (10), and the second is the difference between the number of contacts belonging to two reference structures (e.g., 1 and 2), that is [Eq. (13)]:

$$Q(C) = \sum_{i < j}^N [C_{ij}^{(1)} C_{ij} - C_{ij}^{(2)} C_{ij}] \quad (13)$$

where $C_{ij}^{(1)}$ and $C_{ij}^{(2)}$ are the contact maps of the reference structures, and C_{ij} is the contact map of the instantaneous configuration. The order parameter that measures the change in the number of native contacts is defined as follows: as we consider two distinct native states, we take these as the reference structures. Every contact that occurs to state 1 has a value +1 and every contact that belongs to structure 2 has a value -1. Contacts that appear in both 1 and 2, or do not appear in either, do not contribute to the order parameter.

The reason why we assign negative values to native contacts of structure 2 is that we compute the free-energy difference between the protein in configurations 1 and 2. If we had assigned 0 to the contacts of structure 2 then we would not have been able to distinguish it

from unfolded configurations that do not have any native contacts at all. For our specific case, $C_{ij}^{(1)}$ represents the structure in Figure 1 a, while $C_{ij}^{(2)}$ corresponds to the one shown in Figure 1 b, and Q has values between -15 and 30. Because the

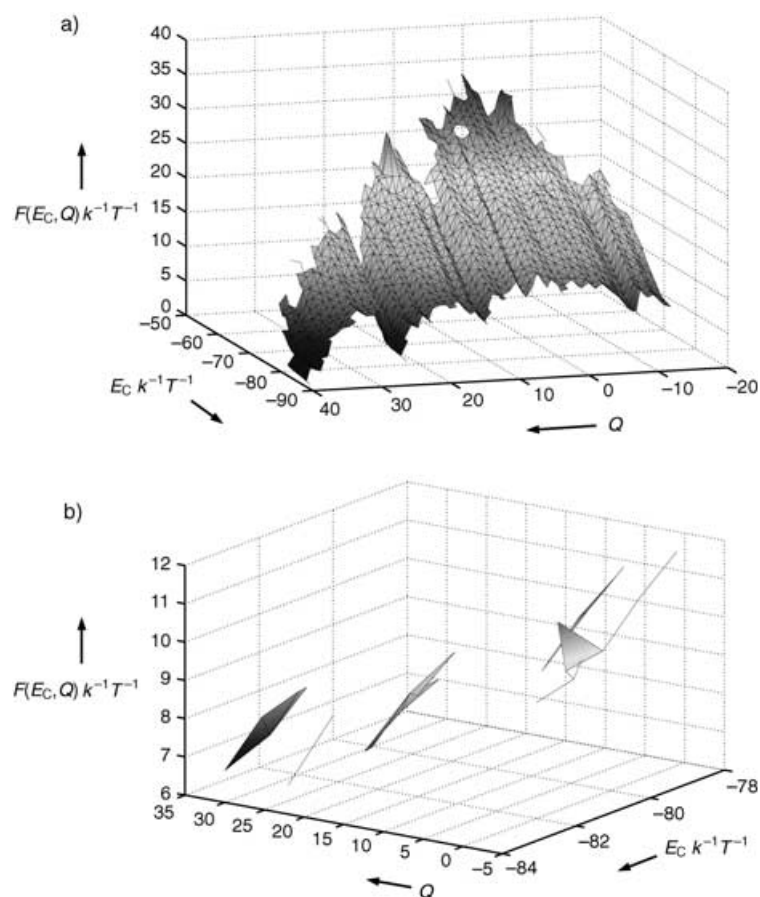


Figure 3. Plot of the free-energy landscapes computed with the VMPT algorithm (a) and the standard APT scheme (b). The free energies $F(E_C, Q)$ are a function of the conformational energy E_C [Eq. (10)] and of the number of native contacts Q [Eq. (13)]. It is important to notice the big difference in the sampling; in fact, the number of points sampled with VMPT is 30 times larger than that with APT.




Figure 4. Average free energy computed with five long runs (10^{10} MC steps, Table 1 II) of the old scheme, compared with the result of five shorter VMPT simulations (10^9 MC steps Table 1 I), at $T = 0.5 \text{ kT}^{-1}$.

number of native contacts includes the contacts with the substrate of the reference state, it can be used to compute the free-energy difference between the unbound state and the specifically bound one.

We performed 15 simulations, five of them with VMPT (using the parameters in Table 1 I) and the other ten with APT (five using the parameters in Table 1 I, and five with the parameters in Table 1 II). In Figure 2 we compare the average free energies at $T=0.1$ (with error bars). We only show those free energies that were sampled in all five simulations of each group. It is clear that the VMPT approach leads to a much better sampling of the free-energy landscape. The advantage of the VMPT approach becomes even more obvious if we plot the free-energy "landscape" as function of two order parameters (viz. the conformational energy [Eq. (10)] and the number of native contacts). In this case the APT method is almost useless, as only small fragments of the free-energy landscape can be reconstructed. The total number of points sampled with VMPT is 20 times larger than that with APT, and the energy range probed is one order of magnitude larger (see Figure 3).

To check the accuracy of the VMPT method, we compared the average free energy obtained by APT and VMPT at high temperatures where the APT scheme works reasonably well. As can be seen in Figure 4 the two methods agree well in this regime, although a much longer APT simulation was needed. Even though the APT runs required 20 times more MC cycles, the method still probes about 30% less of the free-energy landscape than the VMPT scheme.

As the implementation described above is not based on a particular feature of the system under study, the results obtained suggest that the VMPT method may be useful for the study of any system that is normally simulated using parallel tempering. Examples of the application of parallel tempering in fully atomistic simulations of protein folding can be found in refs. [12,13].

Acknowledgements

I.C. would like to thank Dr. Georgios Boulougouris for many enlightening discussions. This work is part of the research program of the "Stichting voor Fundamenteel Onderzoek der Materie (FOM)", which is financially supported by the "Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO)". An NCF grant of computer time on the TERAS supercomputer is gratefully acknowledged.

Keywords: computational chemistry · free-energy calculations · molecular modeling · protein folding

- [1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. N. Teller, E. Teller, *J. Chem. Phys.* **1953**, *21*, 1087.
- [2] D. Frenkel, *Proc. Natl. Acad. Sci. USA* **2004**, *101*, 17571.
- [3] D. Frenkel, B. Smit, *Understanding Molecular Simulations*, Academic Press, New York, **2002**, p. 389.
- [4] A. P. Lyubartsev, A. A. Martsinovski, S. V. Shevkunov, P. N. Vorontsov-Velyaminov, *J. Chem. Phys.* **1992**, *96*, 1776.
- [5] E. Marinari, G. Parisi, *Europhys. Lett.* **1992**, *19*, 451.
- [6] C. J. Geyer, E. A. Thompson, *J. Am. Stat. Assoc.* **1995**, *90*, 909.

- [7] S. Miyazawa, R. Jernigan, *Macromolecules* **1985**, *18*, 534, Table VI.
- [8] G. M. Torrie, J. P. Valleau, *J. Comput. Phys.* **1977**, *23*, 187.
- [9] B. A. Berg, T. Neuhaus, *Phys. Rev. Lett.* **1992**, *68*, 9.
- [10] I. Coluzza, H. G. Muller, D. Frenkel, *Phys. Rev. E* **2003**, *68*, 046703.
- [11] R. Faller, Q. Yan, J. J. de Pablo, *J. Chem. Phys.* **2002**, *116*, 5419.
- [12] C. Y. Lin, C. K. Hu, U. H. E. Hansmann, *Proteins Struct. Funct. Genet.* **2003**, *53*, 436.
- [13] A. Schug, W. Wenzel, *Europhys. Lett.* **2004**, *67*, 307.

Received: December 16, 2004

Revised: March 4, 2005