Optimal allocation of replicas in parallel tempering simulations

Nitin Rathore, Manan Chopra, and Juan J. de Pablo[a]
Department of Chemical and Biological Engineering, University of Wisconsin-Madison, Madison, Wisconsin 53706

We have studied the efficiency of parallel tempering simulations for a variety of systems including a coarse-grained protein, an atomistic model polypeptide, and the Lennard-Jones fluid. A scheme is proposed for the optimal allocation of temperatures in these simulations. The method is compared to the existing empirical approaches used for this purpose. Accuracy associated with the computed thermodynamic quantities such as specific heat is also computed and their dependence on the trial-exchange acceptance rate is reported. © *2005 American Institute of Physics.*
[DOI: 10.1063/1.1831273]

I. INTRODUCTION

Parallel tempering or replica-exchange Monte Carlo[1–4] provides an efficient means for improving conformational sampling in systems with rugged energy landscapes, particularly at low temperatures. The basic premise behind the method's usefulness is that at high temperatures the system of interest can explore phase space relatively unhindered. This method relies on simultaneously simulating multiple replicas of the system. Exchanges or swaps of low- and high-temperature configurations are attempted periodically, thereby allowing low-temperature configurations to escape local free energy minima. Parallel tempering and its variants have been applied for the study of complex fluids and macromolecules, including proteins.[5–12]

In parallel tempering or replica exchange, N noninteracting copies or replicas of the protein molecule are simulated in N boxes, each at a different temperature. In addition to the standard Monte Carlo (parallel tempering) or molecular dynamics (replica exchange) moves in each box, the conformations in different replicas are swapped at regular intervals. Trial swaps are accepted with probability

$$p(E_1,\beta_1 \rightarrow E_2,\beta_2) = \min[1,\exp(\Delta\beta\Delta E)], \tag{1}$$

where ΔE is the difference in energy between conformations in two adjacent boxes, and $\Delta\beta$ is the difference between their inverse temperatures. To satisfy detailed balance, N replicas should be synchronized whenever a swap move is proposed. This is usually implemented by performing an equal number of Monte Carlo or molecular dynamics steps on each replica before attempting a swap move. The exchange between the structures in different replicas facilitates relaxation of structures that might otherwise be trapped in local energy minima.

For parallel tempering simulations to be most efficient, each replica should spend the same amount of time at each temperature. This can be achieved by a temperature distribution that yields the same acceptance probability for swap moves between all adjacent pairs of replicas. A protocol for

determining such an optimal distribution has not been identified yet. A simplistic approach assumes a geometric distribution of temperatures; this approach yields an optimal distribution if the specific heat of the system of interest stays constant with temperature. However, if the specific heat does change with temperature, the performance of a geometric-distribution approach is suboptimal, particularly near a phase transition.

An alternative approach for optimal replica allocation is proposed in this work. We derive a near-optimal distribution analytically by considering the particular case of Gaussian energy distributions of varying widths. The formalism is then extended to realistic model systems, including proteins and the Lennard-Jones fluid. A rigorous comparison with existing schemes for replica allocation is presented. We also discuss how the statistical errors in several computed thermodynamic quantities change with the number of replicas. The thermodynamic quantities are calculated for a given temperature range using histogram reweighting.[13] Results are presented in the form of deviations in estimated specific heat from its final value for a model protein. Using a larger number of replicas results in more accurate thermodynamic-property estimates, this increased accuracy, however, comes at the expense of higher computational demands.

II. MODEL

The generality of the proposed approach is explored by considering a range of model systems. Sample calculations are performed for a Go-type model of protein A, an atomistic β-hairpin system in the continuum, a coarse-grained lattice protein, and the Lennard-Jones fluid. These models are described only briefly; the reader is referred to the literature for additional details.

A. Protein A using Go-type model

The majority of our calculations are performed on a Go-type[14] representation of protein A, Fig. 1. Protein A is a three helix bundle protein consisting of 46 residues [residues 10–55 of the B domain of *Staphylococcus aureus* protein A (Protein Data Bank No. 1BDD)].[15] Very briefly, the protein

---

[a]Author to whom correspondence should be addressed. Electronic mail: depablo@engr.wisc.edu
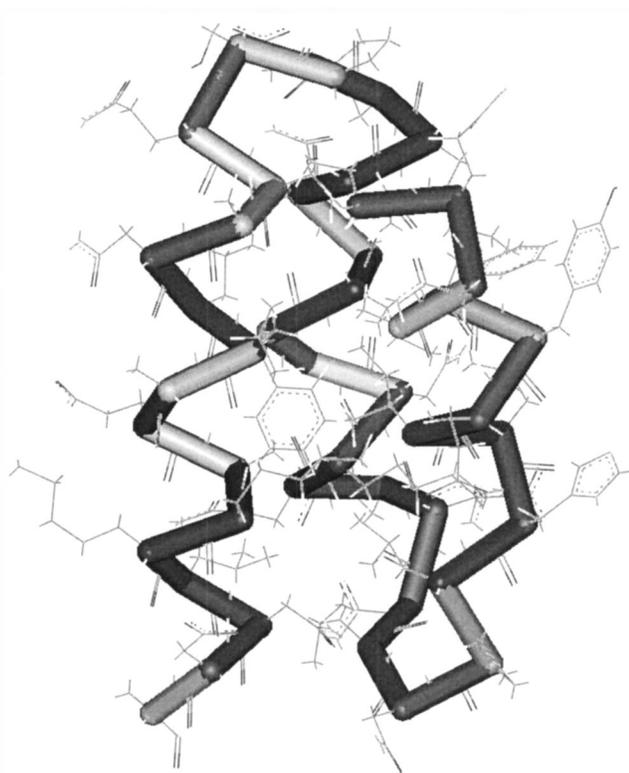
FIG. 1. Go-type model representation of the three helix bundle protein, protein A. Each amino acid is replaced by a single bead at the position of the alpha carbon atom.

FIG. 2. United atom representation of the native hairpin structure of the C-terminal fragment of protein G.

model is coarse grained so that each amino acid is represented by individual beads located at the C^{α} atoms. Consecutive beads along the chain interact through an anharmonic potential given by

$$V_{BB} = \sum_{i=1}^{N-1} U_{BB}(d_i), \tag{2}$$

where d_i is the distance between bead i and $i+1$ and

$$U_{BB} = k_1(d-d_o)^2 + k_2(d-d_o)^4, \tag{3}$$

where $d_o = 3.8$ Å is the equilibrium separation between the beads, $k_1 = 1\,\epsilon/\text{Å}^2$ and $k_2 = 100\epsilon/\text{Å}^4$. In addition to these bond potentials, the nonadjacent beads that form native contacts in the folded structure interact through a Lennard-Jones potential, whereas the non-native contacts interact solely through a repulsive potential. In our calculations we have used $\epsilon = 1$ kcal/mol.

B. β-hairpin using atomistic model

The C-terminal domain (GEWTYDDATKTFTVTE) of protein G (Protein Data Bank accession No. 1GB1) is also used for some of our calculations. Thermodynamic and structural studies have shown that this hairpin exhibits many of the basic features of protein folding, including the formation of a hydrophobic core and hydrogen bonds that stabilize the native conformation. Figure 2 shows a schematic representation of this peptide in its native hairpin configuration. The CHARMM19 (Ref. 16) force field is used with a united atom representation where the nonpolar hydrogen atoms are
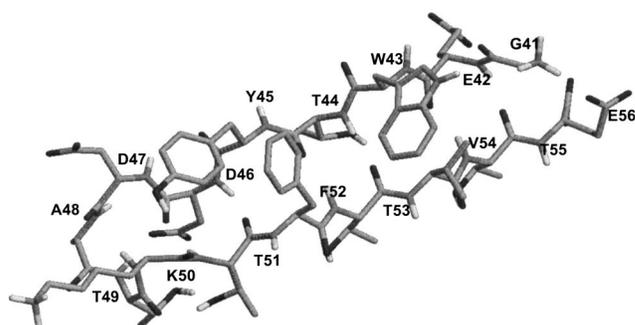
combined with the heavy atoms to which they are bonded. We use the EEF1 model parameters,[17] where the partial charges on the amino acids are modified to neutralize the side chains and the patched molecular termini. The density of states was computed for this system in our previous work,[18] readers are referred to that reference for further details.

C. Polyalanine on a lattice

To capture the behavior of a non-Gaussian system (see below) we have also considered the folding transitions of helical $[\text{-Ala-}]_{28}$ on a lattice. The lattice model employed here is based on the SICHO (side chain only) model by Kolinski *et al.*[19–22] Each amino acid is represented by the centroid of its side chain (Fig. 3); a protein is modeled as a chain connecting these virtual particles on a cubic lattice, with the lattice spacing corresponding to 1.45 Å in real proteins. The chain vectors representing virtual bonds between interaction centers are of variable length, ranging from $9^{1/2}$ to $30^{1/2}$ lattice units. The knowledge-based force field proposed by
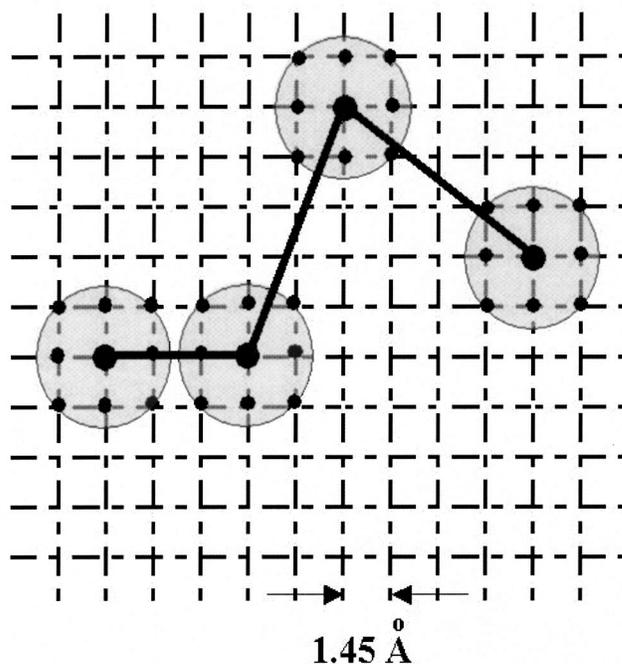

FIG. 3. Schematic representation of the side chain only lattice model.

Kolinski *et al.*[19–22] is employed. It includes a chain stiffness potential, a secondary structure bias, short-range interactions, hydrogen-bond interactions, and long-range interactions. Such a model shows high cooperativity in the folding behavior and a clear coexistence of two phases (folded and unfolded) near the folding transition.[23]

In addition to these peptide models, we also present results for the Lennard-Jones fluid.[24]

III. SIMULATION METHOD

Conventional molecular dynamics is employed to perform "control runs." To determine the optimal distribution of temperatures, replica-exchange simulations are conducted by performing swap moves between adjacent replicas. The configurations in individual boxes are propagated simultaneously and independently using molecular dynamics. The swap moves are performed at regular molecular dynamics step intervals by using the acceptance criteria given in Eq. (1). The momenta in the boxes are rescaled uniformly according to the new temperatures. This algorithm is discussed in detail in Ref. 5.

In a canonical Monte Carlo simulation, sample conformations with potential energy E are generated based on the distribution

$$P_\beta(E) \propto \exp(-\beta E).\tag{4}$$

In parallel tempering simulations, swap moves are accepted according to the acceptance criteria given in Eq. (1).

If the density of states of a system $\Omega(E)$ is known, this probability of acceptance for the swap moves between replicas at inverse temperatures β_1 and β_2 can be computed analytically based on the following expression:

$$P_{\mathrm{acc}}(\beta_1,\beta_2) = \int_1 \int_2 P_{\beta_1}(E_1) P_{\beta_2}(E_2)$$
$$\times \min[1,\exp(\Delta\beta\Delta E)]dE_1 dE_2,\tag{5}$$

where

$$P_\beta(E) = \frac{\exp(-\beta E)\Omega(E)}{Q(\beta)},\tag{6}$$

and $Q(\beta)$ is the canonical partition function, given by

$$Q(\beta) = \int \exp(-\beta E)\Omega(E)dE.\tag{7}$$

We have computed the density of states of the systems of interest by using a Wang-Landau approach.[18,23,25,26] Based on Eq. (5) we could then compute the probability of acceptance of swap moves for different overlapping ranges of energy distributions without performing any additional simulations. This information provide us with a means of arriving at optimal temperatures that will ensure equal acceptance of swaps.

IV. RESULTS AND DISCUSSION

Our goal is to determine the temperature distribution that ensures equal acceptance probability of swap moves between all the adjacent pairs of replicas. This approach would maxi-mize the flux of configurations from low- to high-temperature boxes. In the canonical ensemble, the probability of accepting a swap between two replicas is correlated with the area of overlap between the corresponding energy distributions. We therefore start by looking at the dependence of this area of overlap between two model Gaussians as a function of the separation between them.

A. Model Gaussians

A normalized Gaussian distribution in energy centered at $\bar{E}$ and having width σ is given by

$$P(E) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left[\frac{-(E-\bar{E})^2}{2\sigma^2}\right].\tag{8}$$

Now consider two such Gaussian distributions centered at $\bar{E}_1$ and $\bar{E}_2$ ($\bar{E}_2 > \bar{E}_1$) and having widths σ_1 and σ_2. The area of overlap is calculated by finding the energy value E_i at which the two Gaussians intersect [have the same $P(E)$]. The area is then computed by integrating the Gaussian on the right from $-\infty$ to E_i and the Gaussian on the left from E_i to $+\infty$.

$$A_{\mathrm{overlap}} = \int_{-\infty}^{E_i} \frac{1}{\sqrt{2\pi}\sigma_2}\exp\left[\frac{-(E-\bar{E}_2)^2}{2\sigma_2^2}\right]dE$$

$$+ \int_{E_i}^{+\infty} \frac{1}{\sqrt{2\pi}\sigma_1}\exp\left[\frac{-(E-\bar{E}_1)^2}{2\sigma_1^2}\right]dE$$

$$= \frac{1}{2}\mathrm{erfc}\left[\frac{E_i-\bar{E}_1}{\sqrt{2}\sigma_1}\right] + \frac{1}{2}\mathrm{erfc}\left[\frac{\bar{E}_2-E_i}{\sqrt{2}\sigma_2}\right],\tag{9}$$

where

$$\mathrm{erfc}(x) = \int_x^{+\infty} \frac{2}{\sqrt{\pi}}\exp(-t^2)dt.\tag{10}$$

For the case where the two Gaussians have equal widths, i.e., $\sigma_1 = \sigma_2 = \sigma$, we get

$$E_i = \frac{\bar{E}_1 + \bar{E}_2}{2}\tag{11}$$

and hence,

$$A_{\mathrm{overlap}} = \mathrm{erfc}\left[\frac{\bar{E}_2-\bar{E}_1}{2\sqrt{2}\sigma}\right] = \mathrm{erfc}\left[\frac{\Delta E}{2\sqrt{2}\sigma}\right],\tag{12}$$

where ΔE is the difference between the means of the two distributions. Figure 4 shows how the area of overlap changes with the spacing between the means of the two Gaussians of equal widths. However, as shown in inset of Fig. 4, these curves collapse onto a single curve if the energy spacing is divided by the width of the Gaussians.

In most systems of interest to our work the width of the Gaussian distributions (σ) changes with energy E. We therefore now consider the case of overlap between two Gaussians of different widths. The exact solution to the area of overlap is given by Eq. (9). The two complementary error functions can be merged and Eq. (9) can be rewritten as
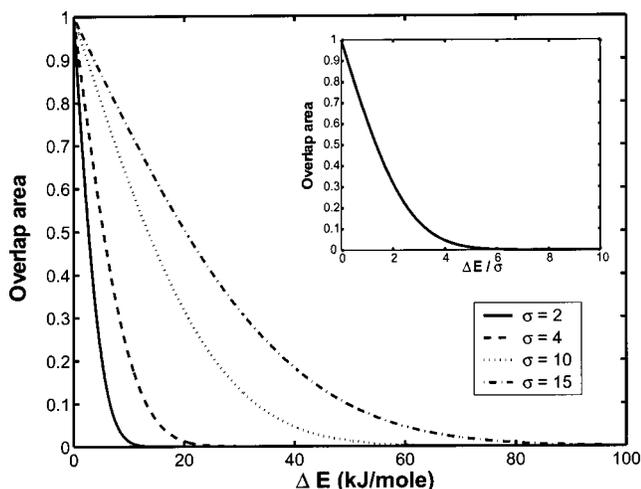
FIG. 4. Area of overlap between two Gaussians as a function of energy spacing ΔE between them. Inset: Area of overlap between Gaussians as a function of energy spacing normalized by the Gaussian width σ. Separate curves collapse on top of each other implying that the overlap area is a function of $\Delta E / \sigma$.
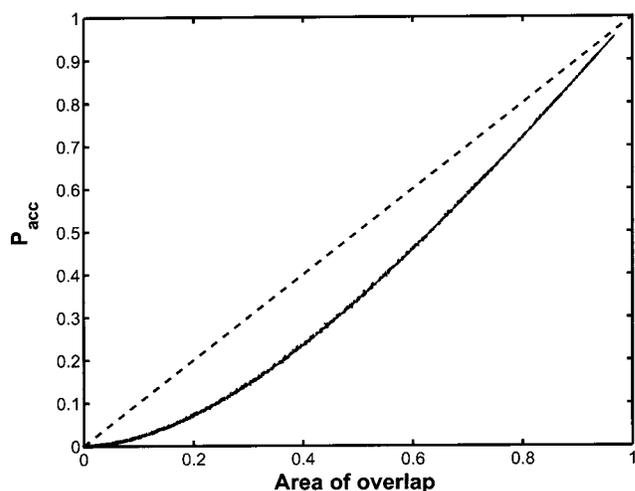

FIG. 5. The probability of accepting a swap move as a function of overlap area. The data is generated for Go-type model of protein A using DOS simulations. The dashed line corresponds to slope of one when P_{acc} is equal to area of overlap.

$$A_{\mathrm{overlap}} = \mathrm{erfc}\left[\frac{z}{2\sqrt{2}}\right] - \frac{e^{z^2/g}\left(\frac{\sigma_2}{\sigma_1} - 1\right)^2}{2(\sqrt{2}\pi z)} + O\left[\left(\frac{\sigma_2}{\sigma_1} - 1\right)^3\right],$$

(13)

where $z = \Delta E / \sigma_m$ and $\sigma_m = (\sigma_1 + \sigma_2)/2$. Therefore, to first order, we have

$$A_{\mathrm{overlap}} = \mathrm{erfc}\left[\frac{\Delta E}{2\sqrt{2}\sigma_m}\right].$$

(14)

This relation, which is exact if the two Gaussians have the same width, is true for the case of unequal widths to the order of the second term in Eq. (13).

This approximation turns out to provide an excellent estimate of the area of overlap between two Gaussians. For practical purposes, however, it is more relevant to look at the probability of acceptance of swaps rather than the area of overlap. The following section discusses how these two quantities are related to each other.

B. Complex systems

We begin by computing the density of states for Go-type model of protein A, an atomistic model of a β-hairpin, and the Lennard-Jones fluid. Once the density of states of a system is known, the energy distribution can be calculated at any temperature. The area of overlap is then determined from the area under the curves that is common to any two adjacent distributions. The probability of acceptance of swaps is computed by solving Eq. (5). Figure 5 shows how these two quantities are correlated with each other for the case of protein A. The curve is generated by considering a broad range of locations and spacings between the overlapping replicas. We can see from Fig. 5 that the relationship is not linear and that the proportionality changes considerably depending on the extent of overlap between the two replicas.

Throughout the remainder of this work, only the probability of accepting trial swap moves, P_{acc} is considered. Based on our findings for model Gaussians, we examine the dependence of P_{acc} with $\Delta E / \sigma_m$. Figure 6 presents data for three different systems. For all systems, P_{acc} has the same functional dependence on $\Delta E / \sigma_m$. To obtain a certain target acceptance of swaps, we need to allocate the temperatures so as to achieve the same value of $\Delta E / \sigma_m$ for each of the adjacent pairs. This target value of $\Delta E / \sigma_m$ is not system dependent. Such an agreement suggests that all the three systems considered here behave like Gaussians and follow the relationship given under Eq. (13).

Figures 7 and 8 show the energy distribution and the specific heat as a function of temperature for protein A. This Go-type model exhibits a sharp folding transition (as evidenced by the dominant peak of the specific heat plot), the energy distributions are nearly Gaussian, even close to the transition temperature.
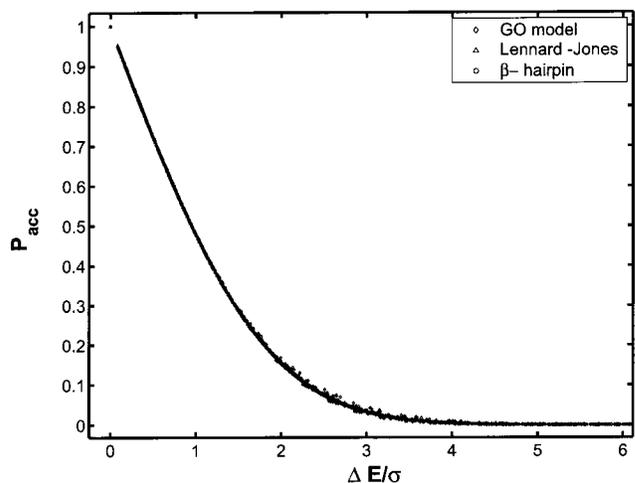

FIG. 6. The probability of acceptance of swap moves as a function of $\Delta E / \sigma$. It is observed that for three different systems the plots collapse on a single curve.
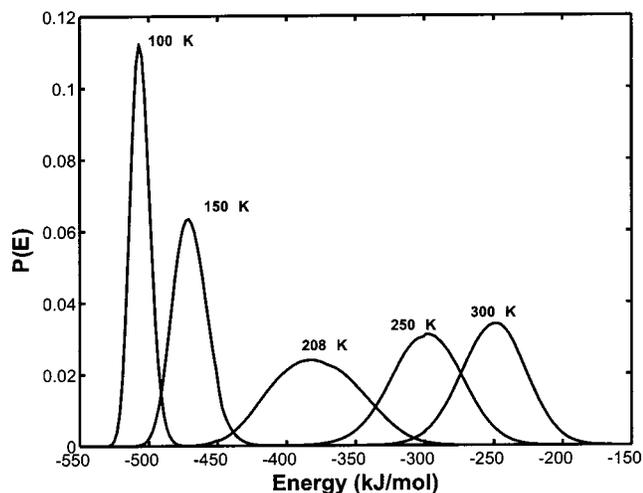
FIG. 7. Energy distribution for Gō model of protein A as obtained using DOS simulations. For the entire temperature range including the transition temperature ($T = 208$ K) the system is Gaussian-like.

FIG. 9. Energy distribution for lattice model of polyalanine for the entire temperature range as obtained using DOS simulations. The system is non-Gaussian close to the transition temperature.

For the particular case of our model protein on a lattice, as shown in Fig. 9, the distributions deviate from Gaussian behavior near and at the transition temperatures. In Fig. 10 we compare the swap acceptance probability for this system with that of a model Gaussian. One can see clear deviations in the middle range of $\Delta E/\sigma_m$. However, even with these deviations, if the target P_{acc} is more than 15%, one can anticipate with reasonable confidence the value of $\Delta E/\sigma_m$ required to achieve a given P_{acc}.

C. Replica-exchange simulations

More generally, one does not have *a priori* knowledge of the density of states. One can perform a short simulation run with a few replicas and obtain the average energy $E(T)$ as a function of temperature. Similarly, $\sigma(T)$ can be roughly estimated. For a target value of P_{acc}, the corresponding $\Delta E/\sigma$ value can be obtained from Fig. 6. We then set the temperature in the first box (T_1) to be equal to the starting point of
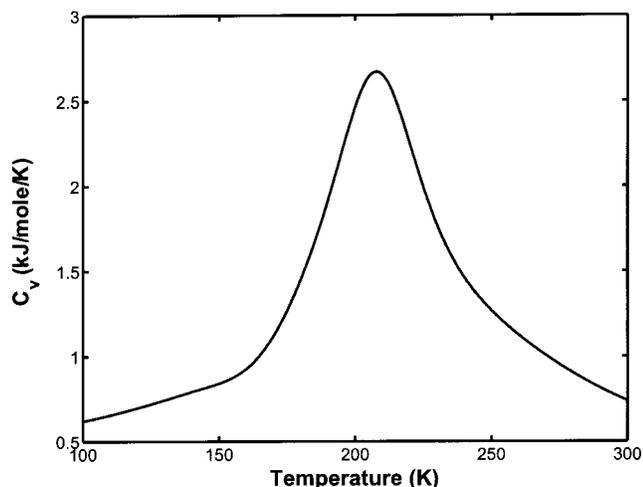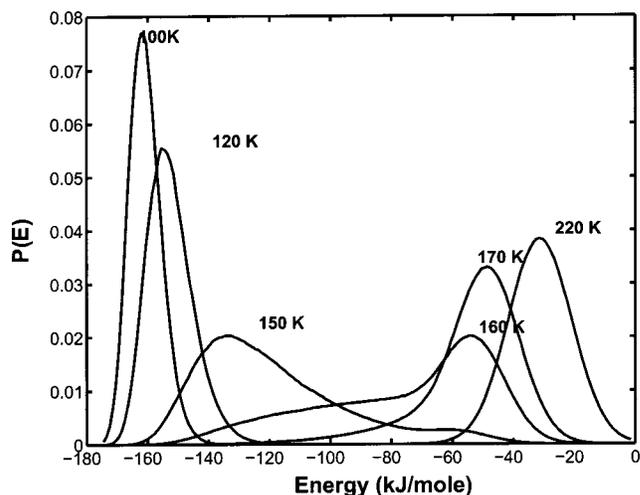
the temperature range of interest (T_{begin}) and the following temperatures $T_i's$ are determined by iteratively solving the equation:

$$\left.\frac{\Delta E}{\sigma_m}\right|_{T_i} = \left[\frac{\Delta E}{\sigma}\right]_{\mathrm{target}}, \tag{15}$$

where $\Delta E = E(T_i) - E(T_{i-1})$ and $\sigma_m = [\sigma(T_i) + \sigma(T_{i-1})]/2$. Since $E(T)$ and $\sigma(T)$ are preestimated, the iteration is relatively fast and simple to implement. One can thus arrive at a distribution that assigns equal $\Delta E/\sigma_m$ to all the adjacent replicas and yield equal acceptance probability of swaps.

We now compare the performance of this scheme with other available empirical approaches. The simplest of these consists of assigning temperatures in a geometric progression. It has been shown that, for the case when the specific heat of the system is independent of temperature, such a distribution does lead to equal acceptances.[27] However, when the specific heat varies with temperature, such as in the case

FIG. 8. Specific heat as a function of temperature for the Go model of protein A as obtained using DOS simulations.
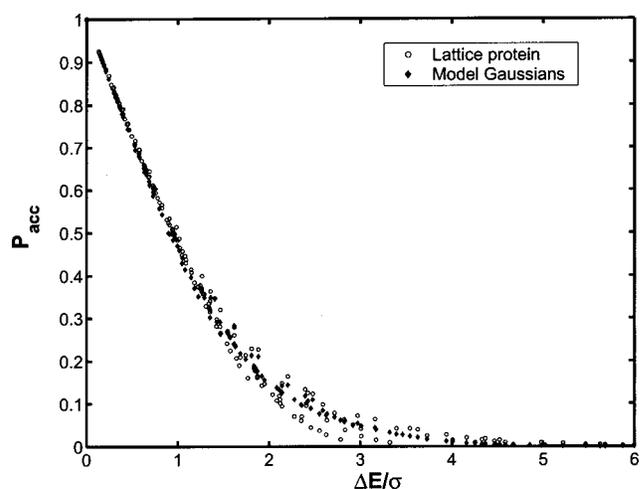

FIG. 10. The probability of acceptance of swap moves as a function of $\Delta E/\sigma$ for a lattice system.

TABLE I. Probability of acceptance of swap as achieved in replica-exchange simulation when the temperatures are assigned based on geometric distribution.

S. No.	T_i/T_{i+1} (geometric)	P_{target}	P_{acc}
1	100/113	0.3	0.43
2	113/128	0.3	0.38
3	128/144	0.3	0.40
4	144/163	0.3	0.35
5	163/184	0.3	0.28
6	184/208	0.3	0.13
7	208/235	0.3	0.18
8	235/265	0.3	0.31
9	265/300	0.3	0.38

TABLE III. Probability of acceptance of swap as achieved in replica-exchange simulation when the temperatures are assigned so as to have equal $\Delta E/\sigma_m$ for all pairs.

S. No.	T_i/T_{i+1} (equal $\Delta E/\sigma$)	P_{target}	P_{acc}
1	100/117	0.3	0.31
2	117/136	0.3	0.34
3	136/157	0.3	0.31
4	157/176	0.3	0.34
5	176/194	0.3	0.33
6	194/211	0.3	0.25
7	211/235	0.3	0.26
8	235/269	0.3	0.31
9	269/305	0.3	0.31

of protein A (see Fig. 8), a geometric distribution may result in highly unequal swap acceptances. Table I shows the P_{acc} obtained for protein A when a geometric distribution is employed. Near the transition temperature, where the specific heat increases, the flux of configurations between replicas is low and the geometric approach is inefficient.

Another scheme employed in the literature[10] does take into account the thermodynamic behavior of the system. Short equilibration runs are performed at fewer temperatures and the data are used to obtain a polynomial fit for average energies $E(T)$ as a function of temperature. This information is then used to determine the target temperatures by iteratively solving the following equation:

$$\exp(\Delta\beta\Delta E) = P_{acc}|_{target}. \qquad (16)$$

The performance of such a scheme is better than that obtained for a geometric progression of temperatures in the respect that it gives more uniform P_{acc}. However for the systems we studied we always obtained a higher value of P_{acc} as compared to the target value used to solve Eq. (16). This is also shown in Table II, where Eq. (16) is solved iteratively for protein A for a target P_{acc} of 0.3. The resultant swap acceptances are consistently higher than the target value.

In Table III results are presented for the optimal distribution determined by setting equal $\Delta E/\sigma_m$ for all adjacent pairs and solving Eq. (15) iteratively. We conclude that this

TABLE II. Probability of acceptance of swaps as achieved in replica-exchange simulation when the temperatures are assigned based on the iterative solutions to $\exp(\Delta\beta\Delta E) = P_{acc}$.

S. No.	T_i/T_{i+1} ($\exp(\Delta\beta\Delta E) = P_{acc}$)	P_{target}	P_{acc}
1	100/113	0.3	0.43
2	113/127	0.3	0.45
3	128/141	0.3	0.44
4	144/157	0.3	0.43
5	157/172	0.3	0.44
6	172/186	0.3	0.48
7	186/198	0.3	0.50
8	198/210	0.3	0.42
9	210/227	0.3	0.41

scheme consistently yields the most uniform acceptance probabilities of swap moves that are closest to the desired value.

D. Optimal number of replicas

The goal of this section is to explore how the performance of a parallel tempering simulation changes as the number of replicas is increased (for a given temperature range). The performance is assessed in terms of the accuracy of simulated specific heat estimates. As the number of replicas is increased, the overlap between the neighboring energy distributions also increases, thereby helping the system relax better. Faster relaxation, however, comes at the cost of added computational demands. Also, as the number of replicas is raised, the round-trip time that is required for a specific configuration to reach the highest temperature and return also increases, thus affecting the performance. As a result, simply increasing the number of replicas does not guarantee an increase in the quality of results that can be achieved in a given amount of computational time.

For systems considered in this work, simulation times of the order of 10–100 ns are sufficient to arrive at a precise estimate of specific heat. Once these "true" values are known, we can determine the deviation of results obtained over a course of multiple short simulations. Specific heat data are generated over a given temperature range by histogram reweighting.[13] For the purpose of determining the optimal number of replicas, short replica-exchange simulations (2 ns for protein A and 0.2 ns for LJ fluid) are performed with varying number of replicas (to collect statistics, six runs are conducted for each replica number).

The performance is reported in Figs. 11 and 12 in terms of the deviation from the true estimate for a fixed total computational time. The total computational time is the cpu time for each replica, times the number of replicas. Figures 11 and 12 show how the specific heat for our Go-type Protein model and Lennard-Jones system approaches the true values for different swap acceptances. The computational time per replica is kept constant when reporting the deviations on the ordinate axis. It is evident from the figures that there exist an optimum (20% swap acceptance) for the number of replicas to be used. Further increases in the number of replicas do not contribute to an increase in the accuracy that can be achieved
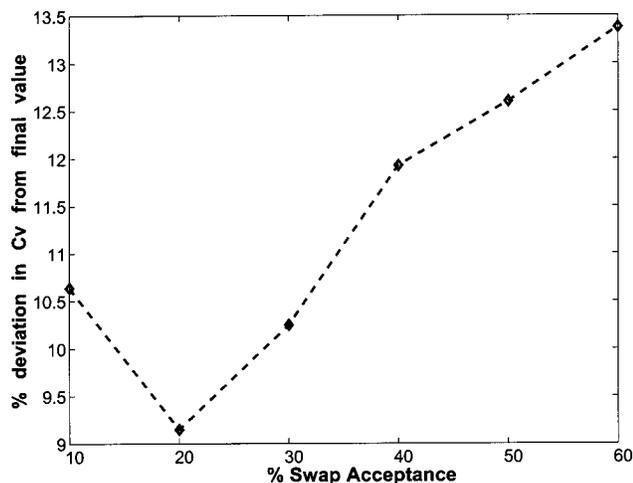
FIG. 11. Percentage deviation in estimated specific heat of Go-type protein as compared to its true value. The true value is computed from 150 ns simulation. The total simulation length (computational time per replica times the number of replicas) is fixed for each data point.

FIG. 13. Percentage deviation in specific heat estimates for Go-type protein A as obtained for two different allocation schemes: the proposed optimal allocation (solid lines) and the geometric distribution of temperatures (in dashed lines). For a range of replicas employed the proposed distribution achieves higher accuracy in a given time.

in a given total computational time. The optimal number of replicas will also depend on the rate at which the configurations are swapped, a parameter which has been kept constant in this study. This dependence, however, does not change the fact that a uniform swap acceptance for all the replica pairs is necessary to achieve ideal performance. As shown in Fig. 13, given a system and a temperature range, an optimal allocation of temperatures always yields estimates of higher quality (than those obtained by other schemes) irrespective of the number of replicas used. As the simulation time is increased, both estimates (one from the optimal and one from the geometric distribution) get better, but our proposed approach leads to smaller errors consistently.

V. CONCLUSIONS

In this work we have presented a protocol for optimal allocation of temperatures in parallel tempering or replica-exchange simulations. These temperatures are assigned based on a recipe that ensures that $\Delta E/\sigma_m$ is equal for all the simulation boxes. The algorithm was devised from a study of model Gaussians, and is extended to a variety of systems, including proteins where the specific heat is a strong function of temperature. Approximate initial guesses for $E(T)$ and $\sigma(T)$ are needed for the algorithm; these can be easily obtained with sufficient accuracy from relatively short simulation runs. The method is shown to be superior to available protocols for obtaining a uniform swap acceptance in multiple replicas.

We have also discussed how the performance of parallel tempering simulations changes as the number of replicas is varied. It is observed that an optimum exists ($\sim$20% swap acceptance for the cases studied in this work), beyond which additional replicas fail to improve the performance of our simulations (evaluated in terms of the accuracy achieved in a given computational time). The optimum may depend on the system under study and the swap frequency employed. However, irrespective of the number of simulation boxes used, the optimal allocation of temperatures yields results of higher accuracy than those obtained using other prescriptions, thus highlighting the need for determining an optimal distribution according to the guidelines proposed in this work.
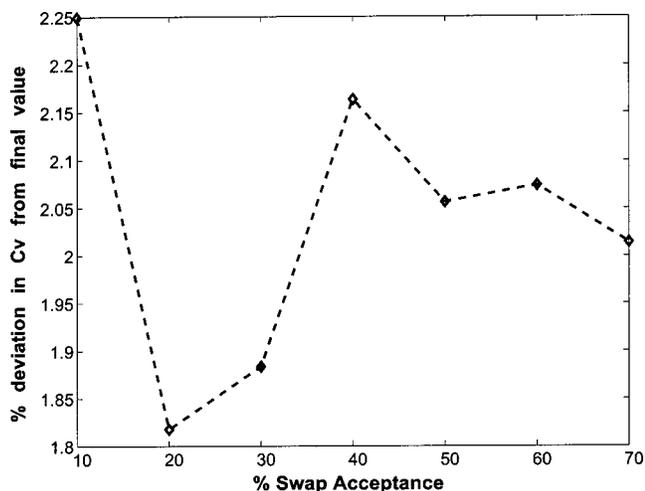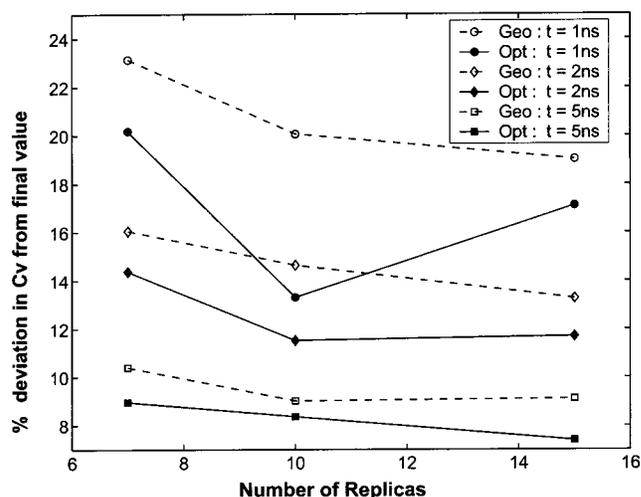

FIG. 12. Percentage deviation in estimated specific heat of LJ fluid as compared to its true value. The true value is computed from 20 ns simulation. The total simulation length (computational time per replica times the number of replicas) is fixed for each data point.

ACKNOWLEDGMENTS

[1] C. J. Geyer, *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface* (American Statistical Association, New York, 1991).

[2] K. Hukushima and K. Nemoto, J. Phys. Soc. Jpn. **65**, 1604 (1996).

[3] R. H. Swendsen and J.-S. Wang, Phys. Rev. Lett. **57**, 2607 (1986).

[4] U. H. E. Hansmann, Chem. Phys. Lett. **280**, 140 (1997).

[5] Y. Sugita and Y. Okamoto, Chem. Phys. Lett. **314**, 141 (1999).

[6] Y. Sugita, A. Kitao, and Y. Okamoto, J. Chem. Phys. **113**, 6042 (2000).

[7] Q. Yan and J. J. de Pablo, J. Chem. Phys. **113**, 1276 (2000).

[8] Q. Yan and J. J. de Pablo, Phys. Rev. Lett. **90**, 035701 (2003).

[9] N. C. Ekdawi-Sever, P. B. Conrad, and J. J. de Pablo, J. Phys. Chem. A **105**, 734 (2001).

[10] K. Sanbonmatsu and A. E. Garcia, Proteins: Struct., Funct., Genet. **46**, 225 (2002).

[11] R. Faller, Q. Yan, and J. J. de Pablo, J. Chem. Phys. **116**, 5419 (2002).

[12] M. K. Fenwick and F. A. Escobedo, J. Chem. Phys. **119**, 11998 (2003).

[13] S. Kumar, D. Bouzida, R. H. Swendsen, P. A. Kolman, and J. M. Rosenberg, J. Comput. Chem. **13**, 1011 (1992).

[14] H. Abe and N. Go, Biopolymers **20**, 1013 (1981).

[15] M. Cieplak, T. X. Hoang, and M. O. Robbins, Proteins: Struct., Funct., Genet. **49**, 114 (2002).

[16] B. R. Brooks, R. E. Bruccoleri, B. D. Olafson, D. J. States, S. Swaminathan, and M. Karplus, J. Comput. Chem. **4**, 187 (1983).

[17] T. Lazaridis and M. Karplus, Proteins **35**, 133 (1999).

[18] N. Rathore, T. A. Knotts, and J. J. de Pablo, Biophys. J. **85**, 3963 (2003).

[19] B. Ilkowski, J. Skolnick, and A. Kolinski, Macromol. Theory Simul. **9**, 523 (2000).

[20] A. Kolinski, L. Jaroszewki, P. Rotkiewicz, and J. Skolnick, J. Phys. Chem. B **102**, 4628 (1998).

[21] A. Kolinski, P. Rotkiewicz, B. Ilkowski, and J. Skolnick, Proteins **37**, 592 (1999).

[22] A. Kolinski and J. Skolnick, Proteins **32**, 475 (1998).

[23] N. Rathore and J. J. de Pablo, J. Chem. Phys. **116**, 7225 (2002).

[24] D. Frenkel and B. Smit, *Understanding Molecular Simulations* (Academic, San Diego, CA, 2002).

[25] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).

[26] N. Rathore, T. A. Knotts, and J. J. de Pablo, J. Chem. Phys. **118**, 4285 (2003).

[27] D. A. Kofke, J. Chem. Phys. **117**, 6911 (2002).