

Exchange frequency in replica exchange molecular dynamics

Daniel Sindhikara

Department of Physics, University of Florida, Gainesville, Florida 32611-8435, USA
and Quantum Theory Project, University of Florida, Gainesville, Florida 32611-8435, USA

Yilin Meng and Adrian E. Roitberg^{a)}

Department of Chemistry, University of Florida, Gainesville, Florida 32611-8435, USA
and Quantum Theory Project, University of Florida, Gainesville, Florida 32611-8435, USA

(Received 9 July 2007; accepted 1 November 2007; published online 10 January 2008)

The effect of the exchange-attempt frequency on sampling efficiency is studied in replica exchange molecular dynamics (REMD). We show that sampling efficiency increases with increasing exchange-attempt frequency. This conclusion is contrary to a commonly expressed view in REMD. Five peptides (1–21 residues long) are studied with a spectrum of exchange-attempt rates. Convergence rates are gauged by comparing ensemble properties between fixed length test REMD simulations and longer reference simulations. To show the fundamental correlation between exchange frequency and convergence time, a simple model is designed and studied, displaying the same basic behavior of much more complex systems. © 2008 American Institute of Physics. [DOI: 10.1063/1.2816560]

I. INTRODUCTION

Conformation sampling is an essential concern to the study of complex molecular systems such as proteins. A major obstacle for the correct sampling of such systems is the fact that the potential energy surfaces of proteins are very rugged and contain a large number of local energy minima.¹ This feature of complex systems causes kinetic trapping due to low barrier crossing rates in constant temperature molecular dynamics. Many different techniques have been introduced to deal with this problem. A recent review on the subject of sampling can be found in Ref. 2. In order to overcome kinetic trapping, generalized-ensemble algorithms including multicanonical algorithms,^{3,4} simulated tempering,^{5,6} and parallel-tempering methods^{7–10} are often used. These methods make the system perform a random walk in temperature or energy space which allows the system under study to more easily overcome energy barriers and hence reduces the problem of kinetic trapping. For methods such as multicanonical algorithm and simulated tempering, multicanonical probability factors are found iteratively by running test simulations at multiple temperatures before the production simulation is run. In replica exchange molecular dynamics (REMD),¹¹ however, the multicanonical weight factor is known *a priori* as the product of Boltzmann weight factors.

REMD, the MD version of parallel tempering (PT),⁹ is one of the more frequently used generalized-ensemble methods.^{12–21} During a REMD simulation, several noninteracting replicas of the original system are simulated independently and simultaneously at different temperatures using standard molecular dynamics methods. Periodically (after a given time interval in regular MD) an attempt is made to exchange conformations between two temperature-adjacent

replicas based on a Metropolis criterion. Successful exchanges then move structures between lower and higher temperatures, passing replicas throughout the temperature spectrum and thus each replica contributes to each canonical ensemble. The replicas continuously heat up and cool down—walking over energetic barriers. REMD has been proven to drastically increase rates of convergence towards a proper equilibrium distribution.^{12–16} Recently, extensions to the original REMD algorithm¹¹ such as using nonexponential temperature distribution^{14,22} and attempting to exchange among all pairs of replicas²³ have been proposed to optimize its efficiency.

In order to efficiently accept the attempted exchanges between replicas, there must be enough overlap between potential energy distributions of neighboring temperature.^{24–27} A number of replicas are then needed to span the entire desired temperature range. For the conventional REMD algorithm, the number of replicas increases as $\sim O(f^{1/2})$, where f is the number of degree of freedom of the system.²⁸ Unfortunately, this relation severely restricts the ability of the REMD algorithm to simulate large systems such as proteins in explicit solvent due to the very large number of processors and CPU time needed. Reducing the number of replicas but keeping the same sampling accuracy becomes an important issue. Recently, several new methods such as Hamiltonian REMD algorithm (H-REMD),^{28–31} hybrid solvent model,^{32,33} reservoir REMD method (R-REMD),^{34–37} replica exchange with dynamical scaling (REDS),³¹ and the coupling of multicanonical algorithm with REMD (REMUCA and MUCAREM) (Refs. 38 and 39) have been developed to improve on conventional REMD.

While REMD has been shown to be effective,^{12–16} the question of how often to make these exchange attempts has not yet been adequately explored. The exchange attempt frequency (EAF) is often chosen in an *ad hoc* manner, using the

^{a)}Author to whom correspondence should be addressed. Electronic mail: roitberg@ufl.edu.

same values as in previous papers on the subject. The time between exchange attempts can be found in the literature in a wide range.^{12–21} EAF have been reported as large¹¹ as 100 ps⁻¹ and as small¹⁸ as 0.05 ps⁻¹. Zhang *et al.*, while studying the 21-residue helical peptide Fs21 [Ace-AAAAA(AAARA)₃A-Nme], looked at three values of EAF and found that 1 ps MD interval is an adequate choice for Fs21.⁴⁰ Opps and Schofield found that for various forms of PT, exchange attempts should be performed quite frequently.⁴¹ Still, no clear discussion exists of what elements one should use to make a proper decision for choosing an EAF. As we will show later, substantial improvement in sampling can be accomplished by properly choosing this parameter. Clearly, as the EAF approaches zero, the REMD sampling becomes identical to that of canonical ensemble MD. An argument has been advanced^{17,18,40,42} proposing that if exchanges are attempted too often, equilibration to the local replica temperature will not happen and thus the system would sample conformational space improperly. These two arguments suggest that both small and large EAFs would provide wrong results, for different reasons. We will show that this reasoning is partially flawed and that exchanges should be attempted as often as possible, provided exchanges are done properly.

II. METHODS

A. Replica exchange molecular dynamics (REMD)

A detailed description of the REMD algorithm can be found in the papers of Okamoto and co-workers.^{6,11} In REMD, N noninteracting copies (replicas) of a system are simulated at N different temperatures (one each). Regular molecular dynamics is run and periodically an exchange of conformation between two adjacent temperatures is attempted. Suppose replica i at temperature T_n and replica j at temperature T_m are attempting to exchange; the following satisfies the detailed balance condition:

$$P_n(i)P_m(j)\rho(i,n \rightarrow j,m) = P_m(i)P_n(j)\rho(j,m \rightarrow i,n). \quad (1)$$

Here $\rho(i \rightarrow j)$ is the transition probability between two states i and j . $P_n(i)$ is the population of state i at temperature n (in REMD assumed Boltzmann). If the Metropolis criterion is applied, the exchange probability is obtained as

$$\rho = \min(1, \exp\{(\beta_m - \beta_n)(E(q^{[i]}) - E(q^{[j]}))\}), \quad (2)$$

where positions, momenta, and temperature of one replica are denoted as $\{q[i], p[i], T_n\}$, $i, n = 1, \dots, N$, $\beta = 1/k_B T$, and $E(q^{[i]})$ is the potential energy of structure i . If the exchange between two replicas is accepted, the temperatures of two replicas will be swapped and velocities rescaled to the new temperatures by multiplying all the old velocities by the square root of new temperature to old temperature ratio:¹¹ $v_{\text{new}} = v_{\text{old}} \sqrt{T_{\text{new}}/T_{\text{old}}}$. Velocity rescaling is crucial in REMD simulations. Since the exchange probability is calculated

using potential energy only, we impose the correct kinetic energy corresponding to the target temperature by velocity rescaling when the exchange attempt is accepted. Neglecting the rescaling of the velocities will result in wrong ensemble average even though the thermostat should eventually relax the system to the new temperature. Instead of velocity rescaling, one way to adjust kinetic energy is velocity reassignment.⁴³ However, conformation sampling using velocity reassignment will probably slow diffusion in conformational space especially when the EAF is extremely large.

Upon completion of a simulation, data are collected from the appropriate temperature and compiled into an ensemble average. The weighted-histogram analysis method^{43–45} (WHAM) can also be applied to collect information from all replicas in order to obtain optimal ensemble averages.

B. Simulation details

For our study, four alanine peptides (Ace- A_n -Nme, $n=1, 3, 5, 7$) and Fs21 [Ace-AAAAA(AAARA)₃A-Nme] (A =Alanine, R =Arginine) were simulated. These were chosen to span a range of sizes. Each peptide was blocked by an acetyl group (Ace) at the N -terminus and an N -methylamine (Nme) group at the C -terminus. The temperatures for REMD were distributed exponentially in a manner such that the expected acceptance probability was 20% for the alanine peptides and 15% for Fs21. Each test simulation was run for 10 ns. Sampling was gauged by computing the deviation of conformations between the test runs and REMD runs (see Sec. III C). The number of replicas and the temperature of each replica in each reference simulation were the same as those in the corresponding test runs. The reference calculations started from the same structures as the corresponding test simulations and used an EAF of 0.5 ps⁻¹. The 0.5 ps⁻¹ EAF (2 ps between exchange attempts) was chosen since time intervals of that order have been highly used and tested.^{16,17,46,47} A regular MD simulation would simply take too long to converge. The details of both test and reference simulations such as temperature ranges and reference simulation time can be found in Table I. Thus if the deviation is relatively small, the test simulation approached the behavior of a longer run in less time and hence, would have a relatively high convergence rate.

All simulations were done using the AMBER9 molecular simulation suite⁴⁸ with the AMBER ff99SB force field.⁴⁹ The SHAKE algorithm⁵⁰ was used to constrain the bonds connecting hydrogen and heavy atoms in all the simulations which allowed use of a 2 fs time step. The generalized Born implicit solvent model GB(OBC) (Ref. 51) was used to model water environment in all our calculations. Each calculation was performed in the canonical ensemble (NVT) with a Langevin thermostat, using a collision frequency of 1 ps⁻¹.

C. Conformation deviation between test and reference simulations

In our study, the metric used to evaluate conformational deviations between test and reference simulations was the root-mean-squared deviation (RMSD) of backbone dihedral

TABLE I. Details of the simulated systems and parameters used.

Peptide	EAFs (ps^{-1})	Reference simulation time (ns)	Replicas (count)	Temperature range (K)
ALA1	0.001,0.005,0.01,0.05,0.1,0.5,1,5,10,50,125,250	50	6	190.1–594.7
ALA3	0.001,0.005,0.01,0.05,0.1,0.5,1,5,10,50,125,250	50	8	217.2–672.6
ALA5	0.001,0.005,0.01,0.05,0.1,0.5,1,5,10,50,125,250	50	10	230.4–755.2
ALA7	0.001,0.005,0.01,0.05,0.1,0.5,1,5,10,50,100,125	50	10	238.7–667.4
Fs21	0.001,0.005,0.01,0.05,0.1,0.5,1,5,10,50,100,125	250	14	280.1–683.0

angle populations. The dihedral angle RMSD was computed by constructing 36×36 histogram per residue by binning φ and ψ angle pairs $10^\circ \times 10^\circ$. These histograms were normalized into populations and the RMSD between the histograms from test and reference runs was calculated and later averaged over all residues in each peptide we studied.

III. RESULTS AND DISCUSSION

Figure 1 shows the average dihedral angle RMSD per residue versus EAF, for four peptides. For the smaller peptides (data not shown), any trend is within the noise of the measurement, due to the fact that in the 10 ns of the test simulations conformational space has already been adequately sampled. With increased system size, however, a significant behavior becomes apparent. As expected, over the range of small EAFs, deviation decreases with increasing EAF. Contrary to expectation, however, no upswing exists towards large EAFs. These results imply that in REMD, having a large EAF has clear advantages with no negative affect on sampling.

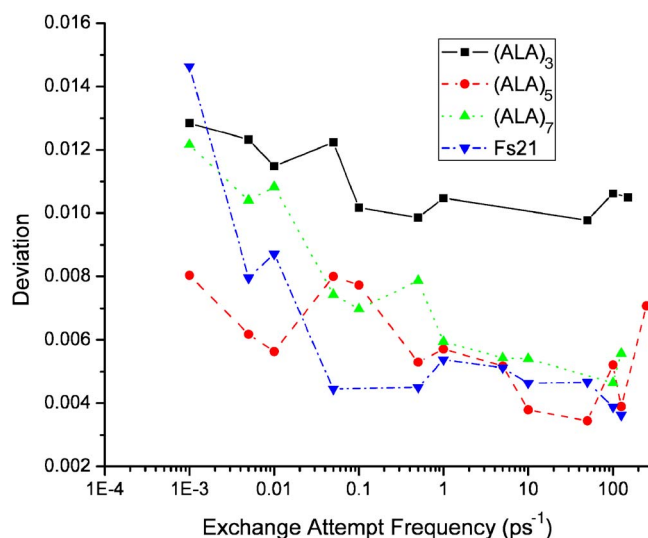


FIG. 1. (Color online) The root-mean-square deviation (RMSD) of test simulations from the corresponding reference simulation as a function of EAF for various peptides. The deviations were calculated between backbone φ and ψ angle distributions (Ramachandran plots). Deviations were normalized by dividing by the number of residues.

For Fs21 in Fig. 1, there is no apparent improvement in final deviation by increasing the EAF above 0.05 ps^{-1} . The RMSDs for these EAFs appear the same (within the noise). However, these are the final deviations after a fixed length (10 ns) of simulation time.

Figure 2 shows the dihedral angle RMSD versus total

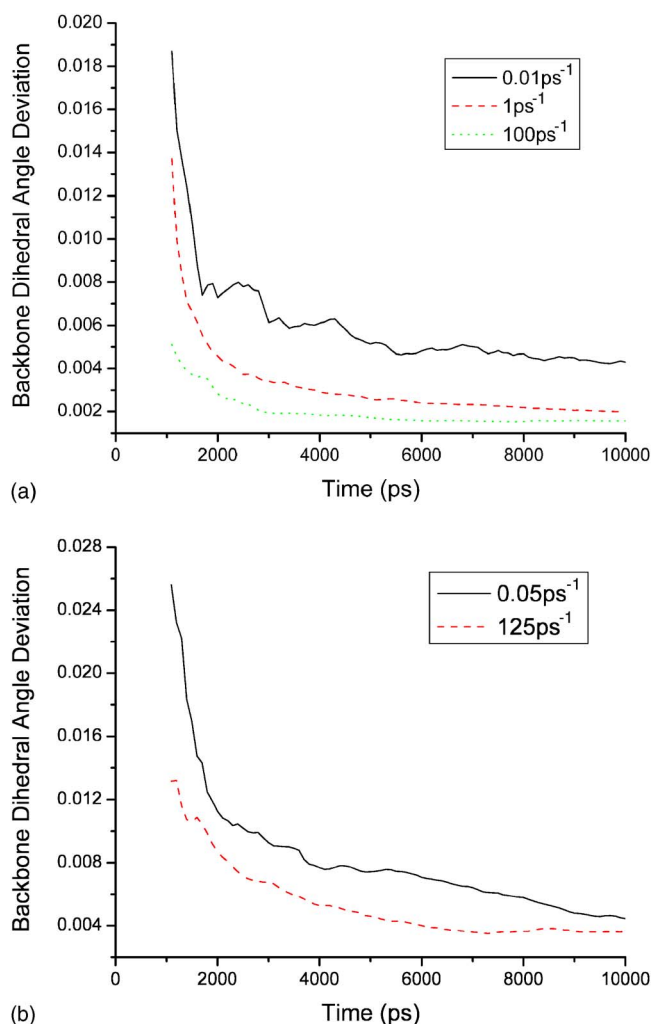


FIG. 2. (Color online) The time evolutions of backbone φ and ψ angle distribution RMSD from corresponding reference simulation: (A) $(\text{ALA})_7$ and (B) Fs21. The calculation of RMSD started from the second nanosecond of 10 ns simulation with a time interval of 100 ps for both $(\text{ALA})_7$ and Fs21.

simulation time. Data are shown for $(\text{ALA})_7$ and Fs21 for various values of EAFs. Deviations were calculated starting from the second nanosecond with time intervals incremented by 100 ps. As expected, the curves decrease with total simulation time. Figure 2 suggests that although the final deviations may be similar between midrange and large EAF simulations, the large EAF simulations clearly converge much faster. After 1 ns simulation, the RMSD of the Fs21 test simulation at 0.05 ps^{-1} was almost twice as much as that from the test simulation at 125 ps^{-1} . We also see a much faster convergence for higher EAF simulations during $(\text{ALA})_7$ simulations. These results indicate that high EAF simulations will achieve convergence in less time than more moderate EAFs.

A. Toy model

To verify the fundamental effect of exchange frequency on REMD, we created a simplified model that can explain the qualitative behavior of convergence versus EAF.

We constructed a model system with two temperatures and three possible configurational states, degenerate in energy. We monitor the convergence of the configurational distribution at the lower temperature (as is typically done for REMD simulations). A transition matrix is used to represent the time-evolution operator of an ensemble as it progresses towards the equilibrium distribution. The distribution, after n iterations of the evolution operator, M , is $v_n = M^n v_0$. Here v_0 is the initial distribution vector, and v_n is the distribution vector after n iterations. We constructed the transition matrix based on Arrhenius-type rates: $k_{ijT} = A \exp(-E_{ij}/k_B T)$ where k_{ijT} is the rate of transition between states i and j at temperature T , E_{ij} is the barrier height between the states, k_B is Boltzmann's constant, T is the absolute temperature of the system, and A is a prefactor that depends on the system and is assumed to be proportional to $T^{1/2}$. At a given temperature, conservation of matter requires that $k_{iiT} = 1 - \sum_{i \neq j} k_{ijT}$. We name k_{Tx} as the rate of exchange between temperatures (replica exchanges) and in this model it is assumed to be constant. With this implementation, $k_{iiT} = 1 - k_{Tx} - \sum_{i \neq j} k_{ijT}$. In practice, k_{Tx} is the product of r_a , the rate of attempted exchanges (equivalent to our EAF in the previous sections), and P_{acc} , the probability of accepting an exchange. We assume $P_{\text{acc}} = \overline{P_{\text{acc}}}$ since the transition matrix represents an averaged time block of simulation. $\overline{P_{\text{acc}}}$ is actually a nontrivial function of the conformation based on the potential energy distribution and the temperature distribution (we will discuss $\overline{P_{\text{acc}}}$ later in the paper). Many REMD simulations are reported in the literature using a P_{acc} between 10% and 40%.^{12–20} An optimal P_{acc} of 20% was proposed for REMD to have the smallest deviations in heat capacity calculations.²⁶

In order to simplify our system, the barriers between all three configurational states are also made equal such that




FIG. 3. (Color online) The RMSD of the final population of the test cases using different barrier heights from the analytical solution as a function of EAF within the toy model. The behavior in this figure agrees with that in Fig. 1 as they both show simulations at higher EAFs are closer to reference simulations after a fixed time interval.

$k_{12T} = k_{23T} = k_{31T} = k_{ijT}$ and $k_{11T} = k_{22T} = k_{33T} = k_{iiT}$. The transition matrix was constructed in the basis $(1h, 2h, 3h, 1l, 2l, 3l)$ where 1, 2, and 3 are the conformations, l is the low temperature, and h is the high temperature. For this calculation, $T_h = 2T_l$. The symmetry of the matrix, M , ensures microscopic reversibility as follows:

$$M = \begin{pmatrix} k_{iiTh} & k_{ijTh} & k_{ijTh} & k_{Tx} & 0 & 0 \\ k_{ijTh} & k_{iiTh} & k_{ijTh} & 0 & k_{Tx} & 0 \\ k_{ijTh} & k_{ijTh} & k_{ijTh} & 0 & 0 & k_{Tx} \\ k_{Tx} & 0 & 0 & k_{iiTl} & k_{ijTl} & k_{ijTl} \\ 0 & k_{Tx} & 0 & k_{ijTl} & k_{iiTl} & k_{ijTl} \\ 0 & 0 & k_{Tx} & k_{ijTl} & k_{ijTl} & k_{iiTl} \end{pmatrix}.$$

Only two unique initial ensemble states (initial distribution vectors) are possible in this system due to the high degeneracy: The state where the high and low temperatures are in the same configuration, and the state where they are not. $(1,0,0,1,0,0)$ is an example of a homopopulated state, and $(1,0,0,0,1,0)$ is a heteropopulated state. Only results from the homopopulated state are shown since there was no significant difference in behavior for the other case.

Based on the system setup, the converged distribution vector is clearly $v_{\infty, i} = \frac{1}{3}$. We quantify the difference between the current configurational ensemble after n iterations and the limiting, equilibrium population. Figure 3 shows the population RMSD from the iterative solution versus k_{Tx} (EAF%) with three different unitless barrier heights, $E_{ij}/k_B T_l$. Final distribution vectors for each k_{Tx} were found by iterating the multiplication of the distribution vector with the corresponding transition matrix ten times. These results

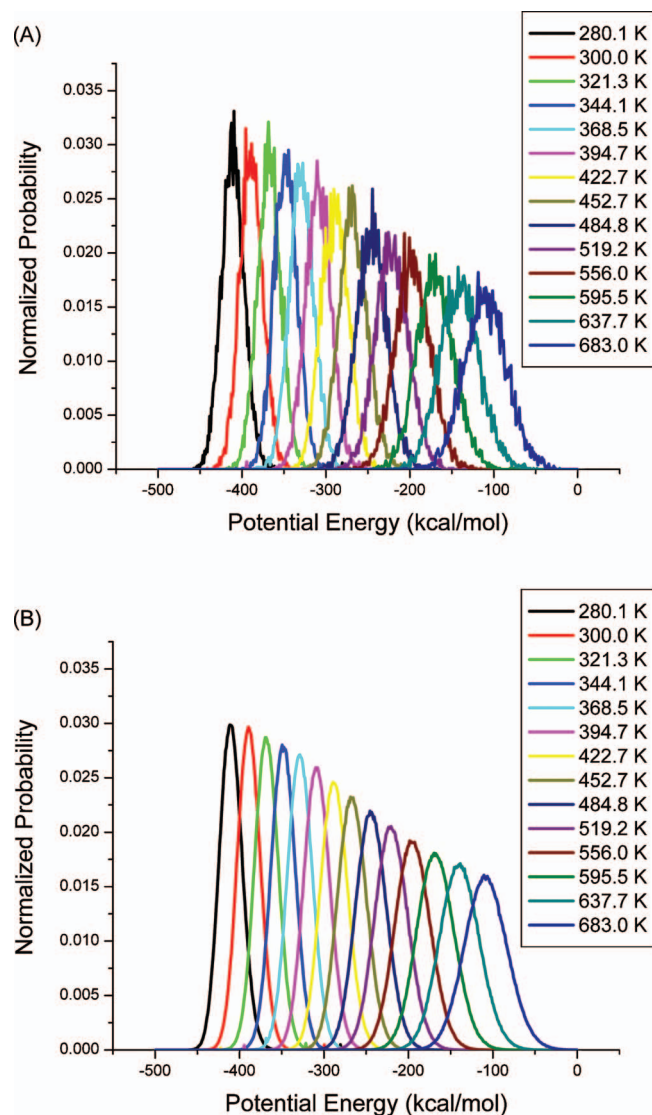


FIG. 4. (Color) Fs21 Potential energy probability distributions at all simulated temperatures: (A) EAF=1 ps⁻¹ and (B) EAF=100 ps⁻¹.

agree nicely with our results from actual REMD simulations (Fig. 1) showing that, indeed, a larger value of EAF accelerates convergence without a turnaround limit. It is also clear that for larger barrier heights, there is an increased benefit of higher EAF.

B. Replica exchange diagnostics at high exchange-attempt frequencies

1. Potential energy and conformational distribution

As discussed in the Introduction, it has been proposed^{17,18,40,42} that if exchanges are made “too often,” thermal equilibration would be inhibited, thus corrupting sampling. If such were the case, there should be a serious

problem with potential energy distributions for different EAFs. In order to test for this problem, potential energy distributions were compared between simulations at high and moderate EAFs. The moderate EAF was used for comparison since it is commonly used and well tested. We plot the distributions of potential energy for the 1 and 100 ps⁻¹ EAF simulations for Fs21 in Fig. 4. The difference in noise simply arises from a difference in the data size for different EAFs and does not affect the results. The distributions for both EAF simulations at a given temperature appear the same (within noise), indicating an identical equilibration of potential energy. The similarity can be quantified by computing the overlap between two distributions at the same temperature but different EAFs, computed as

$$\text{overlap} = \int_{-\infty}^{\infty} dE \cdot P(E)_{\text{EAF}=1 \text{ ps}^{-1}} P(E)_{\text{EAF}=100 \text{ ps}^{-1}}. \quad (3)$$

Table II shows that the overlap at each temperature is nearly unity. We also studied the average potential energy and its standard deviation (related to the specific heat) of (ALA)₇ and Fs21, for different temperatures and EAFs [Figs. 5(a)–5(d)]. These results indicate that high EAF does not disrupt thermal equilibration.

If the potential energy probability distributions follow a Boltzmann distribution, the system’s energy distributions at two canonical temperatures where an overlap exists must obey the following:^{42,52}

$$\ln \left[\frac{P(E, T_2)}{P(E, T_1)} \right] = \left(\frac{1}{k_B T_1} - \frac{1}{k_B T_2} \right) E + \text{const.} \quad (4)$$

We applied the above equation to simulations at both 1 and 100 ps⁻¹ to test canonical equilibration. Figure 6 shows a scatter plot of $\ln[P(E, T_2)/P(E, T_1)]$ versus energy for all adjacent-temperature overlaps for two EAF REMD simulations. Though not shown, the plots were fitted linearly each yielding an r^2 of at least 0.985. Between the two EAF simulations there was a maximum error of 2%. Each EAF had a maximum error of 3% against the ideal slope, $(1/k_B T_1 - 1/k_B T_2)$. These results suggest that high EAF simulations maintain the Boltzmann distribution.

Thermal equilibration, however, does not necessarily imply correct conformational sampling. In addition to the backbone dihedral deviation data (Figs. 1 and 2), we compared helical data for Fs21 against the reference run for various EAFs (Fig. 7).

TABLE II. The overlap between fitted potential energy probability distributions at EAFs of 1 and 100 ps⁻¹.

Temp. (K)	280.1	300.0	321.3	344.1	358.5	394.7	422.7
Overlap	0.997	0.998	0.996	0.998	0.997	0.997	0.994
Temp. (K)	452.7	484.8	519.2	556.0	595.5	637.7	683.0
Overlap	0.996	0.994	0.995	0.994	0.995	0.994	0.995

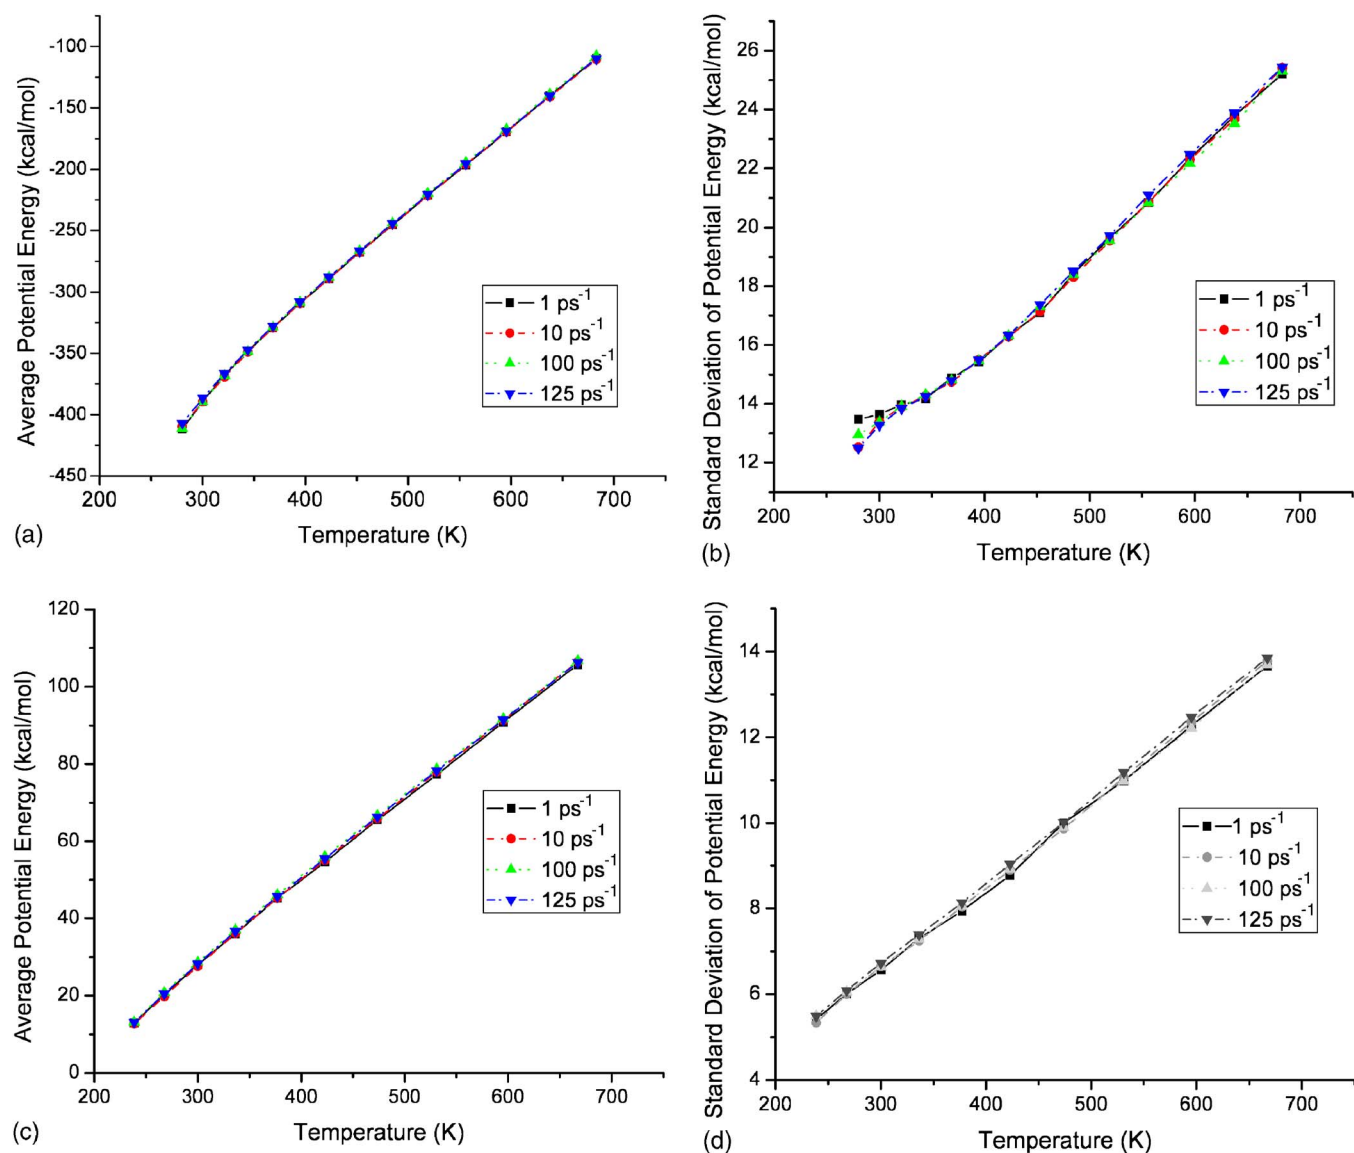


FIG. 5. (Color online) The average potential energy and potential energy standard deviations at EAF=1, 10, 100, 125 ps⁻¹ as functions of temperatures. (a) Average potential energies of Fs21. (b) Potential energy standard deviations of Fs21. (c) Average potential energies of (ALA)7. (d) Potential energy standard deviations of (ALA)7.

Fs21 is known to be partially helical in water.⁵³ The thermodynamics of Fs21 helix-coil transition and helical properties have been studied experimentally and theoretically by several research groups.^{40,53–57} Different force fields and sampling methods and both explicit and implicit water were used in theoretical studies.^{40,56,57} The theoretical helical content highly depends on the force field used during simulations, but a discussion of the reliability of our force field is not an objective of this paper. In our study, Fs21 helical properties at 300 K using various EAFs were also checked in order to observe whether using high EAF will affect peptide thermodynamics. We calculated the average helical content of Fs21 of various EAFs at 300 K using the DSSP definitions.⁵⁸ The results are shown in Fig. 7. The horizontal line in the figure is the helical content given by our reference simulation and the dots are the helical fraction calculated based on test simulations with different EAFs. Uncertainties of average helical fraction at EAFs of 0.05 and 100 ps⁻¹

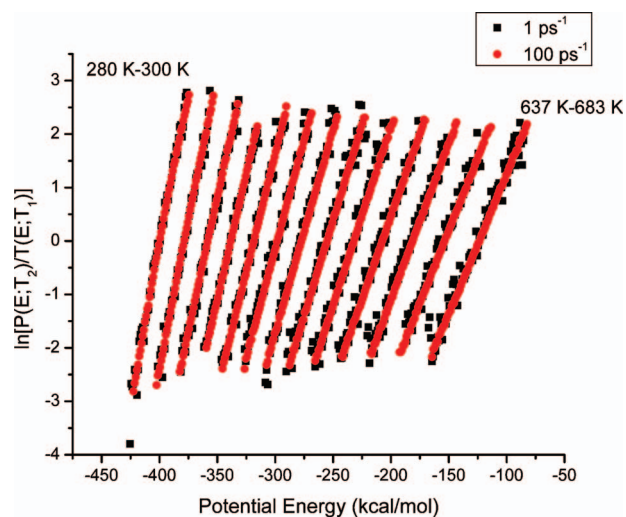


FIG. 6. (Color) Logarithm of energy population ratio for overlap between adjacent temperatures. Data are shown for both 1 and 100 ps⁻¹ simulations.

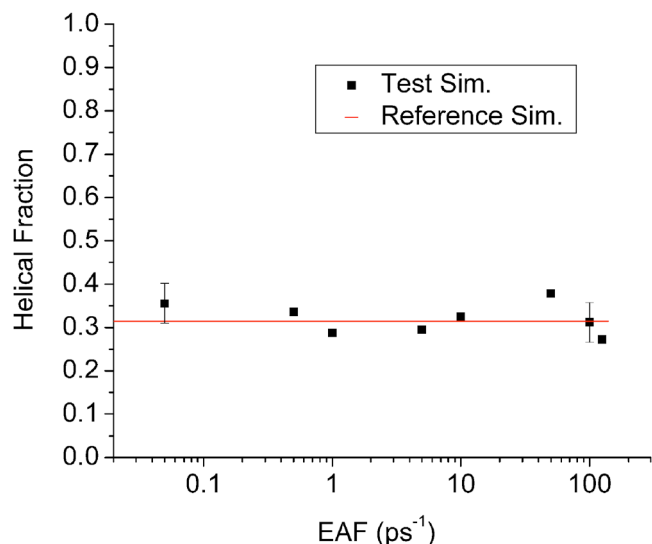


FIG. 7. (Color online) Average helical content of Fs21 for various EAF simulations. The solid line indicates the average helical content of reference simulation. All helical fractions were calculated using DSSP algorithm and at temperature of 300 K.

were also calculated. The last 9 ns of each test simulation were separated into nine pieces and helical contents of each piece of 1 ns trajectory were calculated. The uncertainties were obtained from those nine 1 ns helical fractions. Combining the averages and the uncertainties we obtained, the helical fractions are essentially the same at all EAFs, indicating that the use of a very high EAF does properly sample conformational space. This result is consistent with our potential energy distribution and backbone dihedral angle sampling results.

2. Sampling time

One still might ask the following: If the system is allowed to exchange very often, and hence, each replica potentially spends very little time at a given temperature (continuously), then when is the conformational space sampled? The answer is quite simple: Always. There is always a replica at each temperature and each replica is always sampling some conformational space of its own instantaneous temperature. It should be noted that there is no restriction that barriers must be hopped continuously at one temperature. Figure 8(a) shows the temperature excursions for replica 1 for an EAF of 100 ps^{-1} . This behavior is representative of the other replicas. A full coverage of the temperature space is apparent. The behavior of residue 4 ψ angle of (ALA)₇ versus time for replica 1 is shown in Fig. 8(b). This backbone dihedral angle trajectory shows that the peptide does sample conformational space while moving across temperature space. Though this does not necessarily indicate that the accumulated MD time before an exchange was accepted is sufficient to sample conformations within the local potential energy basin, as in typical Monte Carlo methods, an entire basin need not be sampled continuously. Since the MC step is inexpensive, taking them more frequently results in more opportunities to get accepted moves with negligible drawback. In another perspective, each replica could be consid-

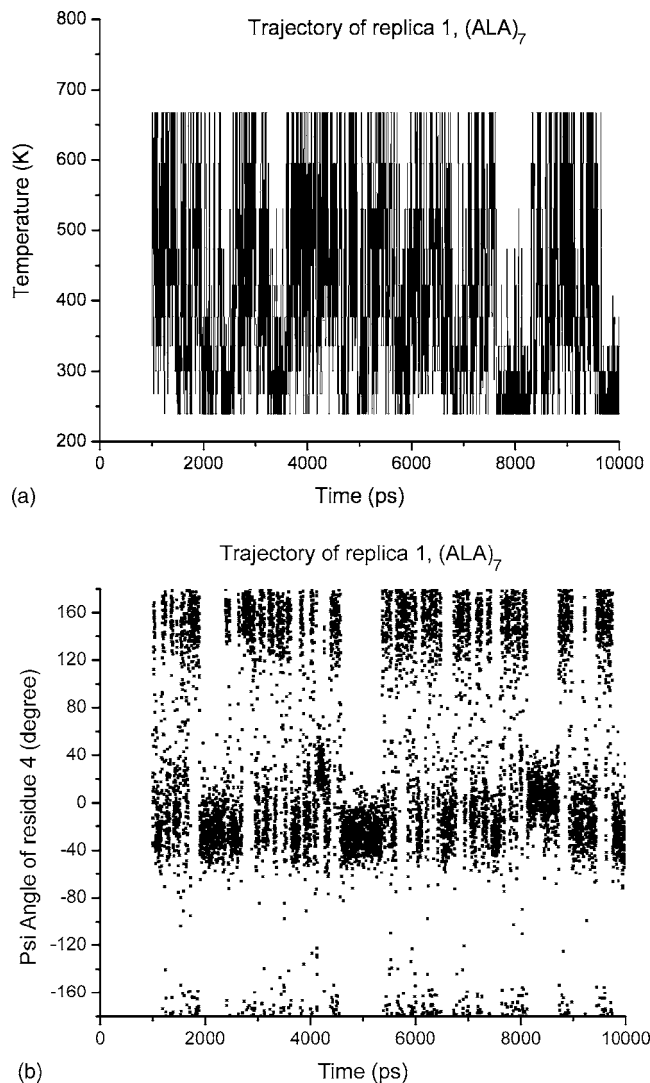


FIG. 8. The time evolutions of replica 1 from (ALA)₇ simulation at EAF = 100 ps^{-1} . (a) Evolution of replica 1 in temperature space; (b) evolution of ψ angles of residue.

ered a particle that continuously samples multicanonical space. Thus as long as exchanges are proper, maintaining detailed balance, increasing EAF will increase the diffusion rate through multicanonical space.

3. Acceptance of MC moves

Increasing *exchange-attempt* frequency does not necessarily imply an increase in the number of accepted moves. The number of accepted moves only increases with an increase in EAF if $(\text{EAF}^{\text{new}}/\text{EAF}^{\text{old}})(P_{\text{acc}}^{\text{old}}/P_{\text{acc}}^{\text{new}}) < 1$. Therefore, we examine the effect of EAF on P_{acc} . The average probability for acceptance of an exchange in the REMD regime with velocity rescaling is^{24,26}

$$\overline{P_{\text{acc}}}(\beta_1, \beta_2) = \int_1 \int_2 P_1(U_1) P_2(U_2) \min(1, e^{\Delta\beta\Delta U}) dU_2 dU_1, \quad (5)$$

where 1 and 2 represent two temperatures at which an exchange attempt is made.

TABLE III. Exchange trapping ratio of selected replicas in (ALA)₇ simulation. EAF has the unit of ps⁻¹.

Replica/EAF	0.005	1	5	10	100	125
1	0.46	0.48	0.55	0.58	0.57	0.58
5	0.35	0.48	0.55	0.58	0.57	0.57
10	0.33	0.50	0.54	0.57	0.57	0.57

Therefore, if the energy distributions are independent of EAF, P_{acc} will also be independent of EAF. This criterion is verified in our system as shown in Fig. 4. The number of accepted Monte Carlo (MC) moves should then increase linearly with EAF: $n_{\text{acc}} = P_{\text{acc}}$ EAF time. This does not mean that the sampling rate will increase linearly with EAF since not all MC moves will necessarily benefit sampling at the same rate.

We propose that at high EAF, a phenomenon we call *exchange trapping* may exist. Exchange trapping would occur when a replica consistently switches back and forth between the same two neighboring temperatures. This might occur due to a shortage of molecular dynamics time to traverse potential energy space. We calculate the exchange trapping as the fraction of exchanges that occur where the new temperature is the same as it was two exchanges before: $T_n = T_{n-2}$. The exchange trapping ratios for (ALA)₇ and Fs21 can be found in Tables III and IV, respectively. The rest of the replicas displayed a similar trend. Based on our calculations, we can see that more than 50% of accepted exchanges were trapped when we attempted EAFs larger than 1 ps⁻¹. Comparing (ALA)₇ and Fs21, we find that the exchange trapping ratios are bigger for Fs21. This is no surprise since Fs21 was simulated with a smaller acceptance ratio than (ALA)₇ (15% and 20%, respectively). Though the exchange trapping ratio does increase with EAF, so does the rate at which a replica traverses temperature space (a usual measure of efficiency and convergence). Note that exchange trapping only occurs ~10% more often in the highest EAF simulations than the ones at commonly used EAFs. Table V shows the number of round-trips between temperature extrema. Still, the exchange trapping does not visibly inhibit the sampling according to the results presented earlier in this paper.

C. High EAF and current program architecture

An important point of caution must be made regarding high EAF simulations on some computer program architectures. Though theoretically, implementing higher exchange attempts should not significantly increase computer time,

certain designs may need to be revamped to allow for rapid exchange attempts. The calculation for exchanges requires only some basic arithmetic and very little information passing (new temperatures are passed, not coordinates or velocities). However, some implementations of REMD (Refs. 43 and 59–62) have taken a regular MD program and simply wrapped it with an outer shell that acts mostly as a script, handling communications and exchange calculations. Under the assumption that the majority of computational time is spent on the MD steps and the exchanges simply add a very small overhead, this architecture holds. However, very large EAFs as those recommended here violate that assumption. Some codes might have to be rewritten to handle communication within the main program instead of through a shell.

IV. CONCLUSIONS

We have closely examined the effect of various values of the exchange-attempt frequency parameter of replica exchange molecular dynamics as applied to peptides of different sizes. We initially expected the backbone dihedral deviation from the correct ensemble to be affected by two major factors functionally dependent on exchange-attempt frequency. The beneficial term that lowers the deviation comes from the fact that with increased exchange attempts, molecules have more opportunity to “see” the broader conformational space made available by sampling at higher temperatures. The hindering term was suggested as due to insufficient sampling from the lack of equilibration.

We have made clear that there is a fundamental benefit of increased exchange frequency through use of a toy model. It was also shown that remarkably consistent potential energy distribution is achieved across EAFs as seen in the overlap we calculated between high and medium EAFs. The conformational distributions and thermodynamic results obtained at equilibrium are independent of EAF. The potential energy distribution, conformational distribution, and thermodynamic results suggested that using high EAFs does not change the canonical ensemble one has chosen. So

TABLE IV. Exchange trapping ratio of selected replicas in Fs21 simulation.

Replica/EAF	0.005	1	5	10	100	125
1	0.44	0.51	0.56	0.59	0.58	0.62
7	0.46	0.49	0.58	0.57	0.59	0.61
14	0.39	0.55	0.57	0.61	0.59	0.60

TABLE V. Number of trips between the two temperature extrema for selected replicas in Fs21 simulation.

Replica/EAF	0.01	0.05	1	5	10	100	125
1	0	0	4	15	22	53	53
4	0	1	12	19	24	34	65
7	0	0	8	13	35	52	72
10	0	1	10	26	28	55	56
14	0	2	12	24	20	50	50

increasing EAF will increase the rate of barrier hopping at lower temperatures, which can lead to faster convergence, due to the fact that the ensembles as well as the sampling rates at the highest temperature are the same for different EAFs. This should hold true for any REMD analog where exchanges are done properly. We then recommend raising the standard EAF to every few time steps for maximum efficiency.

ACKNOWLEDGMENTS

Computer resources and support were provided by the Large Allocations Resource Committee through Grant No. TG-MCA05S010 and the University of Florida High-Performance Computing Center. We thank Carlos Simmerling, Wei Yang, and John Chodera for stimulating discussions. D.S. and Y.M. contributed equally to this work.

- ¹P. G. Wolynes, J. N. Onuchic, and D. Thirumalai, *Science* **267**, 1619 (1995).
- ²A. Roitberg and C. Simmerling, *J. Mol. Graphics Modell.* **22**, 317 (2004).
- ³B. A. Berg and T. Neuhaus, *Phys. Lett. B* **267**, 249 (1991).
- ⁴B. A. Berg and T. Neuhaus, *Phys. Rev. Lett.* **68**, 9 (1992).
- ⁵A. P. Lyubartsev, A. A. Martsinovski, S. V. Shevkunov, and P. N. Vorontsov-Velyaminov, *J. Chem. Phys.* **96**, 1776 (1992).
- ⁶A. Mitsutake, Y. Sugita, and Y. Okamoto, *Biopolymers* **60**, 96 (2001).
- ⁷R. H. Swendsen and J. S. Wang, *Phys. Rev. Lett.* **57**, 2607 (1986).
- ⁸M. C. Tesi, E. J. J. v. Rensburg, E. Orlandini, and S. G. Whittington, *J. Stat. Phys.* **82**, 155 (1996).
- ⁹U. H. E. Hansmann, *Chem. Phys. Lett.* **281**, 140 (1997).
- ¹⁰K. Hukushima and K. Nemoto, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
- ¹¹Y. Sugita and Y. Okamoto, *Chem. Phys. Lett.* **314**, 141 (1999).
- ¹²R. Zhou, B. J. Berne, and R. Germain, *Proc. Natl. Acad. Sci. U.S.A.* **98**, 14931 (2001).
- ¹³A. E. Garcia and K. Sanbonmatsu, *Proteins: Struct., Funct., Genet.* **42**, 345 (2001).
- ¹⁴K. Y. Sanbonmatsu and A. E. Garcia, *Proteins: Struct., Funct., Genet.* **46**, 225 (2002).
- ¹⁵D. Paschek and A. E. Garcia, *Phys. Rev. Lett.* **93**, 238105 (2004).
- ¹⁶J. W. Pitera and W. Swope, *Proc. Natl. Acad. Sci. U.S.A.* **100**, 7587 (2003).
- ¹⁷F. Rao and A. Caffisch, *J. Chem. Phys.* **119**, 4035 (2003).
- ¹⁸M. Cecchini, F. Rao, M. Seeber, and A. Caffisch, *J. Chem. Phys.* **121**, 10748 (2004).
- ¹⁹R. Zhou, *Proc. Natl. Acad. Sci. U.S.A.* **100**, 13280 (2003).
- ²⁰A. E. Garcia and J. N. Onuchic, *Proc. Natl. Acad. Sci. U.S.A.* **100**, 13898 (2003).
- ²¹D. v. d. Spoel and M. M. Seibert, *Phys. Rev. Lett.* **96**, 238102 (2006).
- ²²S. Trebst, M. Troyer, and U. H. E. Hansmann, *J. Chem. Phys.* **124**, 174903 (2006).
- ²³P. Brenner, C. R. Sweet, D. Von Handorf, and J. A. Izaguirre, *J. Chem. Phys.* **126**, 074103 (2007).
- ²⁴D. Kofke, *J. Chem. Phys.* **117**, 6911 (2002).
- ²⁵A. Kone and D. Kofke, *J. Chem. Phys.* **122**, 206101 (2005).
- ²⁶N. Rathore, M. Chopra, and J. J. d. Pablo, *J. Chem. Phys.* **122**, 024111 (2005).
- ²⁷C. Predescu, M. Predescu, and C. V. Ciobanu, *J. Phys. Chem. B* **109**, 4189 (2005).
- ²⁸H. Fukunishi, O. Watanabe, and S. Takana, *J. Chem. Phys.* **116**, 9058 (2002).
- ²⁹P. Liu, B. Kim, R. A. Friesner, and B. J. Berne, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 13749 (2005).
- ³⁰E. Lyman, F. M. Ytreberg, and D. M. Zuckerman, *Phys. Rev. Lett.* **96**, 028105 (2006).
- ³¹S. W. Rick, *J. Chem. Phys.* **126**, 054102 (2007).
- ³²X. Cheng, G. Cui, V. Hornak, and C. Simmerling, *J. Phys. Chem. B* **105**, 8220 (2005).
- ³³A. Okur, L. Wickstrom, M. Layten, R. Geney, K. Song, V. Hornak, and C. Simmerling, *J. Chem. Theory Comput.* **2**, 420 (2006).
- ³⁴A. Okur, D. R. Roe, G. Cui, V. Hornak, and C. Simmerling, *J. Chem. Theory Comput.* **3**, 557 (2007).
- ³⁵A. E. Roitberg, A. Okur, and C. Simmerling, *J. Phys. Chem. B* **111**, 2415 (2007).
- ³⁶H. Li, G. Li, B. A. Berg, and W. Yang, *J. Chem. Phys.* **125**, 144902 (2006).
- ³⁷E. Lyman and D. M. Zuckerman, *Biophys. J.* **91**, 164 (2006).
- ³⁸Y. Sugita, A. Kitao, and Y. Okamoto, *J. Chem. Phys.* **113**, 6042 (2000).
- ³⁹Y. Sugita and Y. Okamoto, *Chem. Phys. Lett.* **329**, 261 (2000).
- ⁴⁰W. Zhang, C. Wu, and Y. Duan, *J. Chem. Phys.* **123**, 154105 (2005).
- ⁴¹S. B. Opps and J. Schofield, *Phys. Rev. E* **63**, 056701 (2001).
- ⁴²Y. M. Rhee and V. S. Pande, *Biophys. J.* **84**, 775 (2003).
- ⁴³J. D. Chodera, W. C. Swope, J. W. Pitera, C. Seok, and K. A. Dill, *J. Chem. Theory Comput.* **3**, 26 (2007).
- ⁴⁴S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman, *J. Comput. Chem.* **13**, 1101 (1992).
- ⁴⁵E. Gallicchio, M. Andreac, A. K. Felts, and R. M. Levy, *J. Phys. Chem. B* **109**, 6722 (2005).
- ⁴⁶H. Lei, C. Wu, H. Liu, and Y. Duan, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 4925 (2007).
- ⁴⁷Y. Z. Ohkubo and C. L. Brooks III, *Proc. Natl. Acad. Sci. U.S.A.* **100**, 13916 (2003).
- ⁴⁸D. A. Case, T. A. Darden, I. T. E. Cheatham, C. L. Simmerling, J. Wang, R. E. Duke, R. Luo, K. M. Merz, D. A. Pearlman, M. Crowley, R. C. Walker, W. Zhang, B. Wang, S. Hayik, A. Roitberg, G. Seabra, K. F. Wong, F. Paesani, X. Wu, S. Brozell, V. Tsui, H. Gohlke, L. Yang, C. Tan, J. Mongan, V. Hornak, G. Cui, P. Beroza, D. H. Mathew, C. Schafmeister, W. S. Ross, and P. A. Kollman, *AMBER9*, University of California, San Francisco, CA, 2006.
- ⁴⁹V. Hornak, R. Abel, A. Okur, B. Strockbine, A. Roitberg, and C. Simmerling, *Proteins: Struct., Funct., Bioinf.* **65**, 712 (2006).
- ⁵⁰J.-P. Ryckaert, G. Ciccotti, and H. J. C. Berendsen, *J. Comput. Phys.* **24**, 327 (1977).
- ⁵¹A. Onufriev, D. Bashford, and D. A. Case, *Proteins: Struct., Funct., Bioinf.* **55**, 383 (2004).
- ⁵²A. E. Garcia, H. Hecce, and D. Paschek, *Annu. Rep. Comp. Chem.* **2**, 83 (2006).
- ⁵³D. J. Lockhart and P. S. Kim, *Science* **260**, 198 (1993).
- ⁵⁴P. A. Thompson, W. A. Eaton, and J. Hofrichter, *Biochemistry* **36**, 9200 (1997).
- ⁵⁵I. K. Lednev, A. S. Karnoup, M. C. Sparrow, and S. A. Asher, *J. Am. Chem. Soc.* **123**, 2388 (2001).
- ⁵⁶A. E. Garcia and K. Y. Sanbonmatsu, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 2782 (2002).
- ⁵⁷H. Nymeyer and A. E. Garcia, *Proc. Natl. Acad. Sci. U.S.A.* **100**, 13934

(2003).

⁵⁸W. Kabsch and C. Sander, *Biopolymers* **22**, 2577 (1983).

⁵⁹M. Feig, J. Karanicolas, I. Charles, and C. L. Brooks III, MMTSB Tool Set, MMTSB NIH Research Resource, The Scripps Research Institute, 2001.

⁶⁰W. Y. Yang, J. W. Pitera, W. C. Swope, and M. Gruebele, *J. Mol. Biol.* **336**, 241 (2004).

⁶¹B. K. Ho and K. A. Dill, *PLOS Comput. Biol.* **2**, 0228 (2006).

⁶²B. Ozkan, A. Wu, J. D. Chodera, and K. A. Dill, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 11987 (2007).