Published in final edited form as:

J Chem Theory Comput. 2008 August 1; 4(8): 1386–1400. doi:10.1021/ct800063d.

Implementation of a Serial Replica Exchange Method in a Physics-Based United-Residue (UNRES) Force Field

Hujun Shen, Cezary Czaplewski, Adam Liwo, and Harold A. Scheraga^{*}

Baker Laboratory of Chemistry and Chemical Biology, Cornell University, Ithaca, New York 14853

Abstract

The kinetic-trapping problem in simulating protein folding can be overcome by using a Replica Exchange Method (REM). However, in implementing REM in molecular dynamics simulations, synchronization between processors on parallel computers is required, and communication between processors limits its ability to sample conformational space in a complex system efficiently. To minimize communication between processors during the simulation, a Serial Replica Exchange Method (SREM) has been proposed recently by Hagan et al. (*J. Phys. Chem. B* **2007**, *111*, 1416–1423). Here, we report the implementation of this new SREM algorithm with our physics-based united-residue (UNRES) force field. The method has been tested on the protein 1E0L with a temperature-independent UNRES force field and on terminally blocked deca-alanine (Ala₁₀) and 1GAB with the recently introduced temperature-dependent UNRES force field. With the temperature-independent force field, SREM reproduces the results of REM but is more efficient in terms of wall-clock time and scales better on distributed-memory machines. However, exact application of SREM to the temperature-dependent UNRES algorithm requires the determination of a four-dimensional distribution of UNRES energy components instead of a one-dimensional energy distribution for each temperature, which is prohibitively expensive. Hence, we assumed that the temperature dependence of the force field can be ignored for neighboring temperatures. This version of SREM worked for Ala₁₀ which is a simple system but failed to reproduce the thermodynamic results as well as regular REM on the more complex 1GAB protein. Hence, SREM can be applied to the temperature-independent but not to the temperature-dependent UNRES force field.

1. Introduction

For a theoretical investigation of protein structure and protein-folding dynamics, it is essential to obtain an adequate sampling of the potential energy surfaces of proteins. However, it is difficult to obtain an accurate canonical distribution from a molecular dynamics (MD) simulation at low temperature because of the rugged energy landscapes of proteins, although the overall shape of foldable proteins is funnel-like.¹ These energy landscapes contain many energy barriers between local minimum-energy states in which a simulation at low temperature is easily trapped. Overcoming these local minima is a key to extending MD simulation to search the conformational spaces of many proteins; therefore, a variety of algorithms have been proposed to overcome this trapping problem. One of them (a generalized-ensemble algorithm²) performs a simulation based on non-Boltzmann probability weight factors to facilitate surmounting energy barriers by a random walk in the energy space. However, the probability weight factors used in this generalized-ensemble algorithm are not known a priori and must be determined by a short trial simulation. Another algorithm, termed the replica-exchange method (REM),^{3–9} is based on the Boltzmann probability weight factor, which is known a priori. Other names for this replica-exchange method are as follows: replica Monte

^{*} Corresponding author phone: (607) 255-4034; fax: (607) 254-4700; has5@cornell.edu.

Carlo method,³ Markov Chain Monte Carlo,⁴ exchange Monte Carlo Method,⁵ multiple Markov chain method,⁶ and parallel tempering.⁷ The replica-exchange method has been tested with the pentapeptide Met-enkephalin system⁸ and has been developed successfully for molecular dynamics simulations.⁹

In the replica-exchange method, a set of simulations of replicas is carried out simultaneously and independently at different temperatures; periodically, exchanges of temperatures between neighboring replicas are attempted and accepted with a well-defined acceptance probability which is consistent with detailed balance; with an ultimate exchange with a high temperature, the simulation at low temperature can surmount energy barriers and thereby search a larger portion of the conformational space of a protein. In addition to overcoming the trapping problem in protein folding, this method takes advantage of the use of canonical ensembles at various temperatures to calculate the thermodynamic characteristics of protein folding.

The replica-exchange method has proven to be a powerful one for studying peptides and small proteins and has been applied to investigate implicit and explicit solvent models in different force fields.⁸⁻¹¹ However, the REM algorithm requires synchronization between different replicas in order to exchange replicas (temperatures) on parallel computers. Therefore, the slowest processor determines the overall performance of the REM simulation. To minimize communication between different processors in the REM simulations, several methods¹²⁻¹⁴ have been developed to facilitate efficient use of distributed computers. One of them is the serial replica exchange method (SREM),¹⁴ which has been tested successfully on the terminally blocked alanine (also termed alanine dipeptide) solvated in 256 TIP4P water molecules by Hagen et al.¹⁴ Because SREM can run asynchronously on distributed computers, it is obvious that SREM has the advantage over REM on an inhomogeneous system. Moreover, in the production phase, SREM can be expected to be more efficient than REM even on homogeneous systems for the following two reasons. First, rigorous implementation of REM requires synchronization of trajectories before a given exchange event; this implies overhead because individual energy and force calculations do not take exactly the same wall-clock time due to the application of cutoff of long-range interactions and other processes running on a given core. These differences in wall-clock time are not substantial, but, when synchronization of hundreds or thousands of trajectories is attempted after 10,000 or 20,000 steps and the protein studied is large, they do matter. Second, communication overhead even on a homogeneous system does matter when frequent communication is required, i.e., when the protein studied is small and many replicas are run.

In the past decade, we have been developing a physics-based united-residue force field, hereafter referred to as UNRES,¹⁵⁻²¹ for off-lattice protein structure simulation for physics-based protein structure prediction. Reduction of the complexity of the problem (i.e., of the number of interaction sites and variables) with UNRES is necessary in order to carry out protein-folding simulations using a reasonable amount of computer time. However, in contrast to most united-residue force fields which are largely knowledge-based potentials, UNRES was carefully derived, based on the physics of interactions, as a cluster-cumulant expansion of the effective free energy of a protein plus the surrounding solvent, in which the secondary degrees of freedom had been averaged out.^{17,20} Moreover, temperature dependence has been introduced recently²² into the force field to reflect the fact that it is a restricted free energy and not a potential energy function.

In this article, we apply this new SREM algorithm to the UNRES force field. In section 2.1, we briefly describe the UNRES model and force field. In section 2.2, we describe the replica exchange method (REM). In section 2.3, we describe the SREM algorithm briefly. In section 3, we report the performance of our implementation of the SREM algorithm for the β -protein 1EOL with the temperature-independent UNRES force field and for the deca-alanine peptide

(Ala₁₀) and 1GAB (a three-helix bundle) with the temperature-dependent UNRES force field. We also carried out regular REM and canonical MD simulations with the UNRES force field for each system in order to compare their performance with that of the SREM algorithm with the UNRES force field under the same conditions. We observed that, with the temperature-independent force field, the performance of SREM is comparable to that of REM. However, SREM could not reproduce results obtained with REM when we tested it on 1GAB with a temperature-dependent UNRES force field. The failure of SREM in this test on 1GAB is attributed to the invalid calculation of the acceptance probability based on the Metropolis criterion, because SREM requires the probability distribution of states as a function of energy and temperature, which is available only in an approximate form for the temperature-dependent version of the UNRES force field.

2. Methods

2.1. UNRES Force Field

In the UNRES model,^{15–28} a polypeptide chain is represented by a sequence of α -carbon atoms linked by virtual bonds with attached united side chains (SC) and united peptide groups (p). Each united peptide group is located in the middle between two consecutive α -carbons. Only these united peptide groups and the united side chains serve as interaction sites; the α -carbons serve only to define the chain geometry, as shown in Figure 1. The UNRES force field has been derived as a restricted free energy (RFE) function of an all-atom polypeptide chain plus the surrounding solvent, where the all-atom energy function is averaged over the degrees of freedom that are lost when passing from the all-atom to the simplified system (viz., the degrees of freedom of the solvent, the dihedral angles for rotation about the bonds in the side chains, and the torsional angles for rotation of the peptide groups about the C ^{α} ...C ^{α} virtual bonds).^{20, 21} The RFE is further decomposed into factors coming from interactions within and between a given number of united interaction sites.²⁰ Expansion of the factors into generalized Kubo cumulants²⁹ enabled us to derive approximate analytical expressions for the respective terms, including the multibody or correlation terms, which are derived in other force fields from structural databases or on a heuristic basis.³⁰ The theoretical basis of the force field is described in detail in our earlier paper.¹⁷ The energy function is expressed by eq 1²²

$$\begin{aligned}
 U = & w_{sc} \sum_{i < j} U_{sc_i sc_j} + w_{scp} \sum_{i \neq j} U_{sc_i p_j} + w_{pp} \sum_{i < j-1} U_{p_i p_j} + \\
 & w_{tor} \sum_i U_{tor}(\gamma_i) + w_{tord} \sum_i U_{tord}(\gamma_i, \gamma_{i+1}) + w_b \sum_i U_b(\theta_i) + \\
 & w_{rot} \sum_i U_{rot}(\alpha_{sc_i}, \beta_{sc_i}) \\
 & + \sum_{m=3}^6 w_{corr}^{(m)} U_{corr}^{(m)} + w_{bond} \sum_i U_{bond}(d_i) + w_{ss} \sum_i U_{ss;i}
 \end{aligned} \quad (1)$$

where $U_{sc_i sc_j}$, $U_{sc_i p_j}$, and $U_{p_i p_j}$ denote the effective free energies of interaction between the side chains in water, between the side chains and the peptide groups, and between the peptide groups, respectively; U_{bond} , U_b , and U_{rot} are the effective potentials for stretching the virtual bonds, for bending the virtual-bond angles, and for the energetics of the rotameric states of virtual side chains, respectively; U_{tor} and U_{tord} are torsional and double-torsional potentials for rotation about the C ^{α} ...C ^{α} virtual bonds; $U_{corr}^{(m)}$, $m = 1, 2, 3, 4$, are correlation or multibody terms pertaining to coupling between backbone-local and backbone-electrostatic interactions²⁰ (where m denotes the order of correlation), and U_{ss} is a term accounting for the energetics of disulfide bonds.³¹ Each of these terms is multiplied by an appropriate weight, w .

The force field was initially applied to solve the protein structure prediction problem as that of finding the global minimum of the effective energy function. However, recently²² we realized that such an approach ignores conformational entropy and, strictly speaking, corresponds to “folding” proteins at the temperature of 0 K; therefore, we redefined²² the protein-structure prediction problem as finding the most probable conformational ensemble at temperatures below the folding-transition and above the glass-transition temperature. Moreover, because UNRES is an effective free-energy and not a potential-energy function, it should depend on temperature explicitly. By taking advantage of the cumulant expansion, this temperature dependence was introduced in our recent work²² by multiplying the energy-term weights by factors $f_n(T)$ (where T denotes absolute temperature) defined by eq 2, where applicable

$$f_n(T) = \frac{\ln(e^1 + e^{-1})}{\ln \left\{ \exp \left[\left(\frac{T}{T_0} \right)^{n-1} \right] + \exp \left[- \left(\frac{T}{T_0} \right)^{n-1} \right] \right\}} \quad (2)$$

where $T_0 = 300$ K.

2.2. Replica Exchange Method (REM)

The replica exchange algorithm based on a molecular dynamics (MD) method is called replica exchange molecular dynamics (REMD).⁹ The basic idea of REMD is to run constant temperature MD simulations on sequential replicas of the system simultaneously and independently, each replica at a different temperature $\{T_0, T_1, \dots, T_i, T_{i+1}, \dots, T_N\}$, where T_N is the maximum temperature at which an MD simulation is run. Usually, those temperatures are arranged in increasing order, e.g., $T_i < T_{i+1}$ ($i = 0, 1, 2, \dots, N-1$). After every fixed number of steps, an attempt is made to exchange temperatures of a pair of neighboring replicas (e.g., i and $i+1$); the velocities are rescaled following the change in temperature, and the exchange processes are then repeated during MD simulations. The neighboring replicas undergo an exchange with the acceptance given by the Metropolis criterion

$$\omega(X_i \rightarrow X_{i+1}) = \begin{cases} 1 & \Delta \leq 0 \\ \exp(-\Delta) & \Delta > 0 \end{cases} \quad (3)$$

where ω is the transition probability and

$$\Delta = (\beta_{i+1} - \beta_i) [U(X_i) - U(X_{i+1})] \quad (4)$$

where $\beta_i = 1/RT_i$, T_i is the temperature corresponding to the i th replica (trajectory), X_i denotes the variables of the UNRES conformation of the i th replica at the attempted exchange point, and $U(X_i)$ denotes the UNRES energy of conformation X_i . If $\Delta \leq 0$, T_i and T_{i+1} are exchanged; otherwise, the exchange is accepted with probability $\exp(-\Delta)$. When the temperature-dependent UNRES force field is applied, eq 4 must be replaced by eq 5;²² this result is obtained directly by inserting the temperature-dependent energy function into eq 14 of ref 9.

$$\Delta = - [\beta_{i+1} U(X_{i+1}, T_{i+1}) - \beta_i U(X_{i+1}, T_i)] + [\beta_{i+1} U(X_i, T_{i+1}) - \beta_i U(X_i, T_i)] \quad (5)$$

The restricted free energy $U(X,T)$ is calculated with eq 5 of ref 22.

To enhance sampling, in this study we used the multiplexing variant of REMD introduced by Rhee and Pande¹² in which m independent replicas are simulated at each temperature; the same extension applied to serial replica exchange is described in the next section.

2.3. Serial Replica Exchange Method (SREM)

The basic idea of the serial replica exchange method is as follows. Let us assume that we know the potential energy distributions, $P_n(E;T_n)$ at the temperatures T_n ($n = 0, 1, 2, \dots, N$). Let us perform a constant temperature molecular dynamics simulation at a temperature T_k , $0 \leq k \leq N$, for a given number of time steps, the final energy being equal to E^0 . From the energy distribution at its neighboring temperature, $T_{k \pm 1}$ (only T_{k-1} if $k = N$, and only T_{k+1} if $k = 0$), we sample an energy E^I . Next, we attempt a move to this neighboring temperature, and the move is accepted or rejected according to the Metropolis criterion (eqs 3 and 4). If this move is accepted, the simulation (trajectory) is continued at the new temperature $T_{k \pm 1}$ with velocities rescaled following the exchange of temperature; otherwise, it remains at the same temperature T_k .

Sampling an energy value from the distribution at a neighboring temperature replaces use of the energy of the conformation at a neighboring temperature in REM but does not require communication with the processor running the trajectory on a neighboring temperature, which is an obvious advantage. On the other hand, the energy distributions at all temperatures must be obtained prior to the production phase of the method which makes it an adaptive method as MUCA and, consequently, introduces an overhead to obtain the energy distribution. Moreover, when SREM is applied to the temperature-dependent UNRES force field, the acceptance probability cannot be computed from eq 5 given only the energy distributions at all temperatures; this problem is discussed in section 2.4.

Now the question is how to obtain the energy distributions to use in the SREM simulation. The algorithm described in ref 14 and also implemented in our work is as follows: (i) perform a short REM simulation to obtain the first approximate energy distribution at each temperature; (ii) with the approximate energy distributions, perform a set of simulations using SREM for a certain amount of time and collect the sampled potential energies at each temperature; (iii) update the potential energy distribution at each temperature by using the sampled potential energies in step (ii); and (iv) repeat steps (ii) and (iii) until the energy distributions converge at each temperature.

To investigate how quickly the potential energy distributions reach equilibrium, we employ the chi-square measure which is defined by eq 6

$$\chi^2(t) = \sum_{i=1}^M \frac{(\overline{P}_i(t) - \overline{P}_i^{ref})^2}{\overline{P}_i^{ref}} \quad (6)$$

Where $\overline{P}_i(t)$ and $\overline{P}_i^{ref}$ are the current and the reference value, respectively, of the energy distribution at the i th bin of the energy histogram averaged over a time window of length t , and M is the number of bins in the energy distribution histogram. In this article, we use the final converged energy distributions obtained by a regular REM simulation as the reference. The energy distributions have become converged when $\chi^2(t)$ becomes stationary. After the energy distributions converge, the update of the energy distributions can be stopped.

To sample the energy distribution, E , we employed the acceptance-rejection method,³² which can be summarized as follows. Suppose that we have an arbitrary probability distribution $P(E;T)$ at temperature T . Let $E^{\max}$, $E^{\min}$, $P^{\max}$, and $P^{\min}$ denote the maximum energy, the minimum energy, the maximum probability, and the minimum probability, respectively, for this energy distribution. Then a point (e^r, y) is generated randomly, where e^r is a random value between $E^{\max}$ and $E^{\min}$, and y is a random number between $P^{\max}$ and $P^{\min}$. If $y \leq P(e^r, T)$, then the value e^r is accepted; otherwise, it is rejected and the sampling step is repeated.

2.4. Modification of the SREM Algorithm for the UNRES Force Field

In the original SREM algorithm,¹⁴ attempts are made to move the temperature of a simulation randomly to its neighboring temperatures during the process. However, this could make the simulation move more frequently to some specific temperatures before the potential energy distributions reach equilibrium. For example, the simulations of Ala₁₀ and 1EOL using this SREM algorithm in Figure 2 show that the trajectories preferentially move to low temperatures for Ala₁₀ and to some lower temperatures for 1EOL during the updating phase. Therefore, to avoid this problem, we modified the original SREM algorithm as follows: (i) the frequency of each replica (simulation) moving (up or down) from its temperature to a neighboring temperature is updated periodically during the MD simulation and (ii) based on the “up or down” frequency for each replica, simulation (trajectory) attempts are made to move to its neighboring temperatures in favor of low frequency. In a word, the distinction between the modification of SREM and the original SREM is the different ways to move a simulation to its neighboring temperatures: the SREM with modification is based on the “up or down” frequency, whereas the original SREM is based on random selection.

However, this modification gives rise to another issue: the modification may violate the detailed balance condition because a simulation is forced to exchange with its neighboring temperatures based on its exchange frequency instead of by a random choice. This feature does not downgrade the algorithm during the equilibration phase, i.e., before the energy probability distributions converge, because the detailed balance condition does not hold in SREM for non-converged energy distributions.¹⁴ However, the detailed balance condition should hold in the production phase when sampling is performed with the converged energy distributions, and the modification is, therefore, not applied in the production phase.

Another problem occurs when a temperature-dependent force field is used with the application of the SREM to UNRES. In the regular REM, eq 4 is simply replaced by eq 5 to compute the acceptance probability. However, the use of eq 5 is not straightforward when SREM is performed with a temperature-dependent force field because this requires knowledge of both $U(X_{i+1}, T_{i+1})$ and $U(X_{i+1}, T_i)$ and, because $U(X_{i+1}, T_{i+1})$ is sampled from the energy distribution at temperature T_{i+1} , the conformation X_{i+1} to which it corresponds is unknown. Consequently, we cannot compute $U(X_{i+1}, T_i)$ given $U(X_{i+1}, T_{i+1})$. Hence, in the applications described in this paper (section 3.2), we assumed that $U(X_{i+1}, T_i) \approx U(X_{i+1}, T_{i+1})$. To use eq 5 accurately, we would need to construct the joint (and hence multidimensional) energy distribution of all groups of energy terms that depend on temperature in the same way; then, after sampling all energy components from the distribution at T_{i+1} , and by taking advantage of eq 5 of ref 22, we could compute the energy at temperature T_i . The energy components fall into 4 groups as far as its temperature dependence is concerned:²²

- i. $\sum_{i < j} U_{SCi}(X)$, $\sum_{i \neq j} U_{scipj}(X)$, $\sum_{i < j-1} U_{PiPj}^{VDW}(X)$, $\sum_i U_b(\theta_i)$, $\sum_i U_{rot}(\alpha_{SCi}, \beta_{SCi})$, and $\sum_{i=1}^{nbond} U_{bond}(d_i)$ are independent of temperature.
- ii. $\sum_{i < j-1} U_{PiPj}^{e1}(X)$ and $\sum_i U_{tor}(\gamma_i)$ are multiplied by $f_2(T)$ of eq 2.

- iii. $U_{\text{corr}}^{(3)}(X)$, $U_{\text{turn}}^{(3)}(X)$, and $\sum_i U_{\text{tord}}(\gamma_i, \gamma_{i+1})$ are multiplied by $f_3(T)$ of eq 2.
- iv. $U_{\text{corr}}^{(4)}(X)$ and $U_{\text{turn}}^{(4)}(X)$ are multiplied by $f_4(T)$ of eq 2.

However, the construction of a four-dimensional distribution of these four groups of energy components requires a much greater, prohibitively expensive, effort to obtain an acceptable convergence compared to obtaining a converged one-dimensional distribution of the total energy.

2.5. Evaluation of Thermodynamic Quantities from Replica Exchange Simulation

To compute thermodynamic quantities (such as the partition function, total energy, and heat capacity) from the results of the simulations carried out at different temperatures, the weighted histogram analysis method (WHAM)³³ was used, as described in ref 22. The expressions for the thermodynamic quantities are given by eqs 17–21 of ref 22.

2.6. Simulation Details

All UNRES/SREM MD simulations were carried out with the Berendsen thermostat,³⁴ the implementation of which in UNRES/MD has been described in our earlier papers.^{27,35} Consequently, the random and stochastic forces were not included. In our earlier work,^{27,35,36} we have shown that Newton's dynamics with the Berendsen thermostat leads to faster simulated folding compared to that of Langevin dynamics, which justifies its use in the present work, where the purpose of the simulations was to obtain converged thermodynamic functions and ensemble averages and not to determine the kinetics of folding. The coupling parameter for the Berendsen thermostat was assumed to be $\tau = 1$ mtu (where mtu stands for molecular time unit, and 1 mtu = 48.9 fs), as in our earlier work.³⁵ The velocity Verlet (VV) algorithm,³⁷ with the variable time step extension developed in our earlier work,²⁸ was used, and the basic time step in integrating the equations of motion was 5 fs. The drawings of the structures of the proteins considered in this work were prepared with the MOLMOL program.³⁸

3. Results and Discussion

3.1. Test of the SREM with the Temperature-Independent UNRES Force Field

First, we tested the SREM algorithm on the 37-residue protein 1E0L³⁹ using the temperature-independent UNRES force field parametrized on 1E0L.²² As in our earlier study,²² we selected the central 28-residue fragment of this protein, which corresponds to a three-stranded β -sheet (Figure 3a).

To generate initial approximate UNRES energy distribution functions, we ran a regular REM simulation for 1,000,000 steps using 25 replicas over a range of temperatures from 200 to 500 K with 4 independent trajectories at each temperature. The folding-transition temperature of 1E0L with the force field used is 339 K.²² Replica exchanges were attempted every 20,000 steps during the simulation. The initial UNRES energy distributions at two representative temperatures (220 and 350 K) are plotted in Figure 4a,b. Then, 100 SREM simulations (4 independent simulations at each temperature) were performed using the initial energy distributions obtained from the short REM simulation. During the SREM simulation, the energy distributions were updated every 500,000 steps for 16,000,000 steps, after which the updated energy distributions roughly converged. A subsequent SREM simulation of another 8,000,000 steps was performed with converged energy distributions to collect the thermodynamic data. To compare the results with the REM and canonical MD, we also ran regular REM simulations and canonical MD simulations at 25 temperatures with 4 independent trajectories at each temperature for 24,000,000 steps.

The converged UNRES energy distributions of the SREM simulations at each temperature nearly converged to those of the REM simulations. For example, the converged UNRES energy distributions of the SREM and REM simulations at two representative temperatures (220 and 350 K) are plotted in Figure 4c,d; the black curves show the converged UNRES energy distributions obtained from the REM run, and the red curves show the converged UNRES energy distributions obtained from the SREM run. However, we observed that the converged UNRES energy distributions from the canonical MD simulations (blue curves) at low temperatures did not converge to those of the REM simulation, but they converged to those of the REM simulation only at high temperatures. For example, at $T = 220$ K, the plot of the converged UNRES energy distribution obtained from the canonical MD run has two peaks (one in the low-energy region and the other in the high-energy region, and the high-energy region made a major contribution to the energy distribution; Figure 4c), which means that canonical MD simulations are trapped in the high-energy region, whereas the canonical MD simulation at 350 K avoided this trapping in the high-energy region.

The convergence of the UNRES energy distributions of the REM, SREM, and canonical MD simulations was assessed quantitatively by using the $\chi^2(t)$ measure defined in eq 6, at these two representative temperatures, 220 and 350 K. The values of $\chi^2(t)$ are plotted against the simulation time in Figure 5. At high temperature in Figure 5b, $T = 350$ K, the values of $\chi^2(t)$ of the canonical MD and SREM simulations have nearly converged to that of the REM simulation after 16 million steps. At lower temperature, $T = 220$ K, the χ^2 values of only the SREM simulation has converged to that of the REM simulation, but the χ^2 values of the canonical MD simulation has not converged to them, which confirms that the canonical MD simulation at lower temperature is trapped (see Figure 4c). To investigate the convergence properties of the SREM UNRES energy distributions at all 25 temperatures, all the simulations and the χ^2 values are averaged over all 25 temperatures, and the results are shown in Figure 6. The χ^2 result in Figure 6 shows that the convergence properties of REM and SREM are the same for all temperatures as those at two representative temperatures, 220 and 350 K (Figure 5), and the averaged result of the canonical MD simulations has not converged yet because of the poor convergence behavior at lower temperatures.

To investigate the thermodynamic properties in these simulations, we have plotted heat-capacity curves. The ensemble averages of the heat capacity calculated from SREM, REM, and canonical MD simulations converge in 20,000,000 steps. The heat capacities, calculated for these simulations with 4 million consecutive steps/trajectories, are shown in Figure 7a-c. When the results of the converged heat capacities obtained from the REM, SREM, and canonical MD simulations were compared, it was observed that the two heat capacity curves of the REM and SREM nearly fit to each other, whereas the heat capacity curve of the canonical MD is obviously shifted from these curves (in Figure 7d).

3.2. Test of the SREM with the Temperature-Dependent UNRES Force Field

To test the SREM with the temperature-dependent UNRES force field,²² we first chose a simple system, a ten-residue polyaniline (Ala₁₀) chain. The folding-transition temperature of Ala₁₀ with the force field used is 311 K. To obtain initial energy distribution functions, we ran a regular REM simulation (it should be noted that eq 4 is replaced by eq 5 in the regular REM run with the temperature-dependent force field) for 400,000 steps using 20 replicas over a range of temperatures from 200 to 440 K with 1 trajectory per temperature. Replica exchanges were attempted every 10,000 steps during the simulation. Then, 20 SREM simulations (1 simulation per temperature) were performed using the energy distributions obtained from the initial REM simulation. Next, the energy distributions were updated every 400,000 steps for 10,000,000 steps, at which point the UNRES energy distributions became roughly steady, which indicates that they converged. A subsequent SREM simulation of another 10,000,000 steps was

performed with converged energy distributions to collect the thermo-dynamic data. To compare the SREM to the REM and the canonical MD simulations, we also ran a regular REM simulation and a canonical MD simulation at 20 temperatures also for 20,000,000 steps each with 1 trajectory per temperature.

The initial UNRES energy distributions obtained from the trial REM simulation at two representative temperatures (210 and 280 K) are plotted in Figure 8a,b. During the updating phase in the SREM simulation, the energy distributions obtained from the SREM simulation of 10,000,000 steps have nearly converged to those of the REM simulation (see Figure 8c,d). The results show that the converged UNRES energy distributions of the REM, MD, and SREM simulations, at low and high temperatures, fit each other very well, and no trapping problem was observed in the MD runs at other low temperatures.

The convergence properties of the energy distributions obtained from the REM, SREM, and canonical MD simulations, by using the $\chi^2(t)$ measure, defined in eq 6, were assessed quantitatively at two representative temperatures, 210 and 280 K. Plots of $\chi^2(t)$ against simulation time are shown in Figure 9. From Figure 9, it is observed that the χ^2 values for SREM and canonical MD at temperatures 210 and 280 K converge very well to those of the REM. The convergence properties of the energy distributions obtained from SREM, canonical MD, and REM, averaged over 20 temperatures, are shown in Figure 10, which illustrates the same behavior (averaged over all temperatures) as those at the two representative temperatures, 210 and 280 K (Figure 9).

To investigate the thermodynamic properties from these simulations, heat-capacity curves for Ala₁₀ were plotted. Heat capacities were calculated for those simulations using 2 million consecutive steps per trajectory, which are shown in Figure 11a-c. The ensemble averages of the heat capacity calculated from the SREM, REM, and canonical MD converged in 12,000,000 steps. When the results of the converged heat capacities obtained from REM, SREM, and canonical MD simulations were compared, it was observed that the three heat capacities for REM, MD, and SREM fit each other very well (as shown in Figure 11d).

However, the Ala₁₀ system is a simple one, so that even canonical MD simulations give as good results as REM does with the temperature dependent force field. Hence, we tried to test SREM on a more complex protein system, 1GAB (see Figure 3b), but we were not able to obtain good results, as shown in Figures 12 and 13. From Figure 12, it is observed, at low temperature (i.e., $T = 220$ K), that the maximum of the UNRES energy distribution obtained from SREM is shifted to higher energies compared to that obtained from REM, while the plot of the UNRES energy distribution obtained from the canonical MD run has two peaks (one in the low-energy region and the other in the high-energy region, and the high-energy region made a major contribution to the energy distribution; Figure 12a). The χ^2 result in Figure 12b shows that the averaged results of the SREM and the canonical MD simulations have not converged to that of the REM simulation because of the poor convergence behavior at lower temperatures. Heat capacities were calculated for those simulations using 4 million consecutive steps per trajectory, which are shown in Figure 13a-c. The ensemble average of the heat capacity calculated from the REM converges in 12,000,000 steps, but the results from the SREM and canonical MD did not converge even in 20,000,000 steps. From Figure 13d, the heat capacities obtained from the SREM and canonical MD simulations for 20,000,000 steps have broad and multiple peaks around the transition temperature.

Therefore, SREM has failed to reproduce the thermodynamics of folding of 1GAB with the temperature-dependent UNRES force field because of the invalid calculation of the acceptance probability in eq 4 for SREM. Apparently, the assumption that $U(X_{i+1}, T_i) \approx U(X_{i+1}, T_{i+1})$ (see section 2.4) works for the simple Ala₁₀ system but does not work for a more complicated 1GAB

system. As mentioned in section 2.4, accurate application of eq 5 would require the determination of a four-dimensional distribution of the UNRES energy components which share the same temperature dependence. Consequently, the advantage of minimizing the cost of communication between processors in the SREM algorithm is overwhelmed by the additional cost to obtain a converged distribution of energy components. Therefore, SREM is applicable only to our temperature-independent UNRES force field.

3.3. Parallel Performance of SREM

To check the performance of SREM, we compared it with REM under the same simulation conditions. During the simulation of 1EOL with a temperature-independent UNRES force field, REM and SREM were run for 20 million steps, and REM spent 10706 s while SREM used 9929 s to finish. During the simulation of polyaniline (Ala₁₀) with the temperature-dependent UNRES force field, REM and SREM were run for 20 million steps, and REM spent 101820 s while SREM used 99864 s to finish.

We also carried out a benchmark assessment of the SREM code using the Cray XT3 computer at the Pittsburgh Supercomputer Center. Weak scaling data, obtained by holding the per-node computational work constant (500,000 steps of MD) for SREM simulations (UNRES/SREM) using the Ala₁₀ system, are presented. For comparison, a set of regular REM simulations (UNRES/REM) carried out using the Cray XT3 computer with synchronization for the same system (Ala₁₀) is presented. Those data are given in Table 1, which shows that the time (i.e., setup time, total time, and nonsetup time) of the SREM simulations and REM simulations changes with a given number of processors. Nonsetup time in the table is the difference between the total time and the setup time. It is straightforward to compare the parallel performance of the SREM simulations and regular REM by plotting the nonsetup time of the simulations against the number of processors (see Figure 14). From the results (Table 1 and Figure 14), we observe that the nonsetup time of the SREM simulations increases slightly with increasing number of processors, whereas the nonsetup time of the REM simulations increases significantly with increasing number of processors. Speedup curves, calculated from the weak scaling data in Table 1 by applying eq 7, are shown in Figure 15

$$s(4n) = 4n \frac{t_{n-s}(4)}{t_{n-s}(4n)} \quad (7)$$

where $s(4n)$ and $t_{n-s}(4n)$ are the speed-up and the nonsetup time with $4n$ processors, respectively, given the same per-processor work load; the calculations were always carried out with a multiple of 4 processors.

We found that the increasing number of processors does not affect the efficiency of SREM but does decrease the efficiency of regular REM, which is shown in Figure 15. Therefore, we can conclude that the parallel performance for SREM is better than that for regular REM at this time.

4. Conclusions

We have implemented UNRES/SREM and tested it on three systems: 1EOL with the temperature-independent and Ala₁₀ and 1GAB with the temperature-dependent force fields. By checking the convergence properties of the energy distributions and the thermodynamic properties of 1EOL, calculated from REM and SREM, we have demonstrated that SREM can reproduce the results of REM, while the canonical MD cannot, which means that the SREM algorithm is comparable to regular REM with the temperature-independent UNRES force field.

For two simple systems studied, 1E0L and Ala₁₀, SREM turned out to be more efficient than REM. However, the data from section 3.3 indicate that the gain in wall-clock time is rather incremental than substantial. Moreover, SREM has the adaptive phase in which energy distributions are constructed, which can take a long time for a larger system.

With the temperature-dependent UNRES force field, SREM can be applicable to some simple systems such as Ala₁₀, but Ala₁₀ is so simple that MD simulations can reproduce the same results as that of REM. To check the performance of SREM with our temperature-dependent UNRES force field, we used 1GAB as the target, but SREM failed to produce good results for this protein because of too-crude an estimation of the acceptance probability under the necessary assumption that the change of the effective energy can be neglected when moving from a current temperature to a neighboring temperature. SREM showed the same poor convergence behavior as canonical MD (Figures 12 and 13). One possible solution would be to use the distribution of groups of energy terms which share common temperature dependences instead of the distribution of energy. This modification would, however, kill the advantage of the SREM algorithm over the regular REM algorithm because of the enormous amount of time required to obtain a converged multidimensional distribution of energy components. Another possible solution is to create a database of $[E(T_{i-1}), E(T_i), E(T_{i+1})]$ triples instead of an energy distribution to sample an energy and translate it to the energy at a neighboring temperature; this is correct in principle, but a problem here might be the large size of such a database. Assuming that 20 trajectories are run at different temperatures, and 500 most recent snapshots are saved from each (as in our study), giving 30,000 energy values to store, which is a reasonable size, this database will grow rapidly with the number of trajectories and number of snapshots to store, which is often necessary to treat larger systems. Therefore, SREM can be applied in a straightforward way only with our temperature-independent UNRES force field but not with our temperature-dependent one.

Acknowledgments

This work was supported by grants from the National Institutes of Health (GM-14312), the National Science Foundation (MCB05-41633), and the NIH John E. Fogarty International Center (TW7193). This research was conducted by using the resources of (a) our 800-processor Beowulf cluster at the Baker Laboratory of Chemistry and Chemical Biology, Cornell University, (b) the National Science Foundation Terascale Computing System at the Pittsburgh Supercomputer Center, (c) the John von Neumann Institute for Computing at the Central Institute for Applied Mathematics, Forschungszentrum Jülich, Germany, (d) our 45-processor Beowulf cluster at the Faculty of Chemistry, University of Gdańsk, (e) the Informatics Center of the Metropolitan Academic Network (ICMAN) in Gdańsk, and (f) the Interdisciplinary Center of Mathematical and Computer Modeling (ICM) at the University of Warsaw.

References

1. Wolynes PG, Onuchic JN, Thirumalai D. *Science* 1995;267:1619–1620. [PubMed: 7886447]
2. Hansmann UHE, Okamoto Y. *J Comput Chem* 1997;18:920–933.
3. Swendsen RH, Wang JS. *Phys Rev Lett* 1986;57:2607–2609. [PubMed: 10033814]
4. Geyer CJ, Thompson EA. *J Am Stat Assoc* 1995;90:909–920.
5. Hukushima K, Nemoto K. *J Phys Soc Jpn* 1996;65:1604–1608.
6. Tesi MC, van Rensburg EJJ, Orlandini E, Whittington SG. *J Stat Phys* 1996;82:155–181.
7. Marinari, E.; Parisi, G.; Ruiz-Lorenzo, JJ. *Spin Glasses and Random Fields*. Young, AP., editor. World Scientific; Singapore: 1998. p. 58-98.
8. Hansmann UHE. *Chem Phys Lett* 1997;281:140–150.
9. Sugita Y, Okamoto Y. *Chem Phys Lett* 1999;314:141–151.
10. García AE, Sanbonmatsu KY. *Proteins* 2001;42:345–354. [PubMed: 11151006]
11. Liu P, Kim BJ, Friesner RA, Berne BJ. *Proc Natl Acad Sci USA* 2005;102:13749–13754. [PubMed: 16172406]

12. Rhee YM, Pandé VS. *Biophys J* 2003;84:775–786. [PubMed: 12547762]
13. Rodinger T, Howell PL, Pomes R. *J Chem Theory Comput* 2006;2:725–731.
14. Hagen M, Kim B, Liu P, Friesner RA, Berne BJ. *J Phys Chem B* 2007;111:1416–1423. [PubMed: 17249714]
15. Liwo A, Pincus MR, Wawak RJ, Rackovsky S, Scheraga HA. *Protein Sci* 1993;2:1697–1714. [PubMed: 7504550]
16. Liwo A, Pincus MR, Wawak RJ, Rackovsky S, Scheraga HA. *Protein Sci* 1993;2:1715–1731. [PubMed: 8251944]
17. Liwo A, Oldziej S, Pincus MR, Wawak RJ, Rackovsky S, Scheraga HA. *J Comput Chem* 1997;18:849–873.
18. Liwo A, Oldziej S, Pincus MR, Wawak RJ, Rackovsky S, Scheraga HA. *J Comput Chem* 1997;18:874–887.
19. Liwo A, Kazmierkiewicz R, Czaplewski C, Groth M, Oldziej S, Wawak RJ, Rackovsky S, Pincus MR, Scheraga HA. *J Comput Chem* 1998;19:259–276.
20. Liwo A, Czaplewski C, Pillardy J, Scheraga HA. *J Chem Phys* 2001;115:2323–2347.
21. Liwo A, Arlukowicz P, Czaplewski C, Oldziej S, Pillardy J, Scheraga HA. *Proc Natl Acad Sci USA* 2002;99:1937–1942. [PubMed: 11854494]
22. Liwo A, Khalili M, Czaplewski C, Kalinowski S, Oldziej S, Wachucik K, Scheraga HA. *J Phys Chem B* 2007;111:260–285. [PubMed: 17201450]
23. Oldziej S, Kozłowska U, Liwo A, Scheraga HA. *J Phys Chem A* 2003;107:8035–8046.
24. Liwo A, Oldziej S, Czaplewski C, Kozłowska U, Scheraga HA. *J Phys Chem B* 2004;108:9421–9438.
25. Oldziej S, Liwo A, Czaplewski C, Pillardy J, Scheraga HA. *J Phys Chem B* 2004;108:16934–16949.
26. Oldziej S, Lagiewka J, Liwo A, Czaplewski C, Chinchio M, Nancias M, Scheraga HA. *J Phys Chem B* 2004;108:16950–16959.
27. Oldziej S, Czaplewski C, Liwo A, Chinchio M, Nancias M, Vila JA, Khalili M, Arnautova YA, Jagielska A, Makowski M, Schafröth HD, Kazmierkiewicz R, Ripoll DR, Pillardy J, Scheraga HA. *Proc Natl Acad Sci USA* 2005;102:7547–7552. [PubMed: 15894609]
28. Khalili M, Liwo A, Rakowski F, Grochowski P, Scheraga HA. *J Phys Chem B* 2005;109:13785–13797. [PubMed: 16852727]
29. Kubo R. *J Phys Soc Jpn* 1962;17:1100–1120.
30. Kolinski A, Skolnick J. *J Chem Phys* 1992;97:9412–9426.
31. Chinchio M, Czaplewski C, Liwo A, Oldziej S, Scheraga HA. *J Chem Theory Comput* 2007;3:1236–1248.
32. Robert, CP.; Casella, G. *Monte Carlo Statistical Methods*. Springer; New York: 1999. Monte Carlo Integration; p. 92–96.
33. Kumar S, Bouzida D, Swendsen RH, Kollman PA, Rosenberg JM. *J Comput Chem* 1992;13:1011–1021.
34. Berendsen HJC, Postma JPM, van Gunsteren WF, DiNola A, Haak JR. *J Chem Phys* 1984;81:3684–3690.
35. Khalili M, Liwo A, Jagielska A, Scheraga HA. *J Phys Chem B* 2005;109:13798–13810. [PubMed: 16852728]
36. Liwo A, Khalili M, Scheraga HA. *Proc Natl Acad Sci USA* 2005;102:2362–2367. [PubMed: 15677316]
37. Swope WC, Andersen HC, Berens PH, Wilson KR. *J Chem Phys* 1982;76:637–649.
38. Koradi R, Billeter M, Wüthrich K. *J Mol Graph* 1996;14:51–55. [PubMed: 8744573]
39. Macias MJ, Gervais V, Civera C, Oschkinat H. *Nat Struct Biol* 2000;7:375–379. [PubMed: 10802733]
40. Johansson MU, de Chateau M, Wikstrom M, Forsen S, Drakenberg T, Bjorck L. *J Mol Biol* 1997;266:859–865. [PubMed: 9086265]

Figure 1.

The UNRES model of the polypeptide chain. Dark circles represent united peptide groups (p) and open circles represent the C^{α} atoms. Ellipsoids represent side chains, with their centers of mass at the SCs. The ps are located halfway between two consecutive C^{α} atoms. The virtual-bond angles, θ , the virtual-bond dihedral angles, γ , and the angles, α_{SC} and β_{SC} , that define the location of a side chain with respect to the backbone, are also indicated.

Figure 2.

Temperature distributions between replicas, expressed as average number of replicas at a given temperature. (a) Ala₁₀ with the temperature-dependent force field and (b) 1E0L with the temperature-independent force field, respectively, in the SREM simulation without modification for 10 million steps. The temperature-independent and temperature-dependent runs are placed in the same figure to show that the bias problem that occurred in the old SREM algorithm is caused by the algorithm itself instead of by the assumption that $U(X_{i+1}, T_i)$ is approximately equal to $U(X_{i+1}, T_{i+1})$.

(a) 1E0L

(b) 1GAB

Figure 3.
Ribbon diagrams of the experimental structures of (a) 1E0L³⁹ and (b) 1GAB.40

Figure 4.

The UNRES energy distributions for 1EOL at two representative temperatures, 220 and 350 K (with the temperature-independent force field);²² curves in the top panels correspond to the initial rough UNRES energy distributions from the trial REM simulations for 1,000,000 steps at (a) 220 K and (b) 350 K; curves in the bottom panels correspond to the converged UNRES energy distributions from the REM (black), SREM (red), and canonical MD (blue) simulations at the two representative temperatures (c) 220 K and (d) 350 K.

Figure 5.

Convergence measure of the UNRES energy distributions for the REM (diamonds), SREM (squares), and canonical MD (triangles) simulations of 1E0L, at temperature (a) 220 K and (b) 350 K (with the temperature-independent force field),²² as a function of the simulation time. The χ^2 curves are plotted by using the converged UNRES energy distributions from the REM run for 24,000,000 steps as the reference.

Figure 6.

Convergence measure of the UNRES energy distributions for REM (diamonds), SREM (squares), and canonical MD (triangles) simulations of 1EOL as a function of the simulation time with the temperature-independent force field.²² The χ^2 curves are plotted by using the converged UNRES energy distributions from the regular REM run for 24,000,000 steps as the reference, and the χ^2 values are averaged over 25 temperatures.

Figure 7.

Plots of heat capacity using 4,000,000 MD consecutive steps/trajectory windows taken from (a) the REM run, (b) the SREM run, and (c) the canonical MD run of 1E0L with the optimized force field (the temperature-independent force field).²² The curves in parts (a)-(c) are colored from red to blue according to the number of MD steps; the color codes are shown in the color bar with numbers indicating the number of millions of MD steps. (d) Two converged heat capacities are obtained from the REM (black) and SREM (red) simulations, and the obvious shift of the heat capacity is obtained from the canonical MD (blue) simulation.

Figure 8.

The UNRES energy distributions for Ala₁₀ at two representative temperatures, 210 and 280 K (with the temperature-dependent force field):²² curves in the top panels correspond to the initial rough UNRES energy distributions for the trial REM simulation for 400,000 steps at temperatures (a) 210 K and (b) 280 K. Curves in the bottom panels correspond to the converged UNRES energy distributions at temperatures (c) 210 K and (d) 280 K obtained from the REM (black), SREM (red), and canonical MD (blue) simulations.

Figure 9.

Convergence measure of the UNRES energy distributions for the REM (diamonds), SREM (squares), and canonical MD (triangles) simulations of Ala₁₀ at temperatures (a) 210 K and (b) 280 K, as a function of the simulation time with the temperature-dependent force field.²² The χ^2 curves are plotted by using the converged UNRES energy distributions from the regular REM run for 20,000,000 steps as the reference.

Figure 10.

Convergence measure of the UNRES energy distributions for the REM (diamonds), SREM (squares), and canonical MD (triangles) simulations of Ala_{10} as a function of the simulation time with the temperature-dependent force field.²² The χ^2 curves are plotted by using the converged UNRES energy distributions from the regular REM run for 20,000,000 steps as the reference, and the χ^2 values are averaged over all 20 temperatures.

Figure 11.

Plots of heat capacity using 2,000,000 consecutive MD steps per trajectory taken from (a) the REM run, (b) the SREM run, and (c) the canonical MD run of Ala₁₀ with the temperature-dependent force field.²² The curves in parts (a)-(c) are colored from red to blue according to the number of MD steps; the color codes are shown in the color bar with numbers indicating the number of millions of MD steps. (d) Three separate converged heat capacities obtained from the REM (black), SREM (red), and canonical MD (blue) simulations.

Figure 12.

(a) The UNRES energy distribution for 1GAB at a temperature 220 K with the temperature-dependent force field.²² The black, red, and blue curves show the UNRES energy distributions obtained from the REM, SREM, and canonical MD simulations for 20 million steps, respectively. (b) Convergence measure of the UNRES energy distributions for the REM (diamonds), SREM (squares), and canonical MD (triangles) simulations of 1GAB as a function of the simulation time. The χ^2 curves are plotted by using the converged UNRES energy distributions from the regular REM run for 20,000,000 steps as the reference, and the χ^2 values are averaged over all 20 temperatures.

Figure 13.

Plots of heat capacity using 4,000,000 consecutive MD steps per trajectory taken from (a) the REM run, (b) the SREM run, and (c) the canonical MD run of 1GAB with the temperature-dependent force field.²² The curves in parts (a)-(c) are colored from red to blue according to the number of MD steps; the color codes are shown in the color bar with numbers indicating the number of millions of MD steps. (d) Three separated heat capacities obtained from the REM (black), SREM (red), and canonical MD (blue) simulations.

Figure 14.

Plots of the weak scaling data, probed by holding the per-node computational work constant for Multiplexing Replica Exchange Molecular Dynamics (REM, squares) and Serial Replica Exchange Molecular Dynamics (SREM, triangles) codes using the Cray XT3 computer.

Figure 15. Speedup plots for Replica Exchange Molecular Dynamics (REM, squares) and Serial Replica Exchange Molecular Dynamics (SREM, triangles) codes using the Cray XT3 computer.

Table 1

Weak Scaling Data of UNRES REM and UNRES SREM for Ala₁₀ Using the Cray XT3 Computer at the Pittsburgh Supercomputer Center

no. of processors	UNRES/REM			UNRES/SREM		
	setup time [s]	total time [s]	nonsetup time [s]	setup time [s]	total time [s]	nonsetup time [s]
4	0.14	67.15	67.01	0.16	88.57	88.41
8	0.17	66.75	66.58	0.17	88.35	88.18
16	0.24	68.11	67.87	0.24	87.51	87.27
32	0.40	67.52	67.12	0.47	88.17	87.70
64	0.61	68.77	68.16	0.66	89.85	89.19
128	1.07	69.75	68.68	1.15	88.73	87.58
256	2.05	73.87	71.82	2.29	89.28	86.99
512	3.94	78.74	74.80	4.20	94.99	90.79
1024	7.64	94.14	86.50	8.65	97.58	88.93