Published in final edited form as:

J Phys Chem B. 2007 May 17; 111(19): 5405–5410. doi:10.1021/jp068826w.

Replica Exchange with Solute Tempering: Efficiency in Large Scale Systems

Xuhui Huang^{†,‡}, Morten Hagen[†], Byungchan Kim[†], Richard A. Friesner[†], Ruhong Zhou^{†,§}, and B. J. Berne^{*,†}

Department of Chemistry, Columbia University, New York, New York 10027, and Computational Biology Center, IBM Thomas J. Watson Research Center, 1101 Kitchawan Road, Yorktown Heights, New York 10598

Abstract

We apply the recently developed replica exchange with solute tempering (REST) to three large solvated peptide systems: an α -helix, a β -hairpin, and a TrpCage, with these peptides defined as the “central group”. We find that our original implementation of REST is not always more efficient than the replica exchange method (REM). Specifically, we find that exchanges between folded (F) and unfolded (U) conformations with vastly different structural energies are greatly reduced by the nonappearance of the water self-interaction energy in the replica exchange acceptance probabilities. REST, however, is expected to remain useful for a large class of systems for which the energy gap between the two states is not large, such as weakly bound protein–ligand complexes. Alternatively, a shell of water molecules can be incorporated into the central group, as discussed in the original paper.

The replica exchange method (REM) has become very popular in the sampling of biomolecular systems.^{1–6} However, the REM is restricted to relatively small systems, since the required number of replicas scales as $O(f^{1/2})$, where f is the number of degrees of freedom in the system.⁷ To overcome this problem, we recently introduced replica exchange with solute tempering (REST)⁸ for all-atom simulations in explicit water. In our original study, we tested REST on an alanine dipeptide dissolved in explicit water, a system with about 1500 atoms, and suggested that the required number of replicas now scales as $O(f_p^{1/2})$, where f_p is the number of degrees of freedom in the central group. In addition, we suggested that the speedup versus the REM,

in terms of converging to the correct underlying distribution, is $O(\sqrt{f/f_p})$. In the current study, we broaden the application of REST to three large solvated peptide systems (α -helix, β -hairpin, and TrpCage) to offer an assessment of the efficiency of REST.

Imposing detailed balance on the standard temperature replica exchange (REM)^{1,2} operation results in the acceptance criterion

$$a(i \rightarrow f) = \begin{cases} 1 & \text{if } \Delta_{nm} \leq 0 \\ \exp(-\Delta_{nm}) & \text{if } \Delta_{nm} > 0 \end{cases} \quad (1)$$

*Corresponding author. E-mail: bb8@columbia.edu.

[†]Columbia University.

[‡]Current address: SimBios Center, Department of Bioengineering, Stanford University School of Medicine, Stanford, CA, 94305.

[§]IBM Thomas J. Watson Research Center.

where $\Delta_{nm} = \Delta\beta\Delta E$, $\Delta\beta = 1/(kT_m) - 1/(kT_n)$, $\Delta E = E(X_n) - E(X_m)$, and $E(X_n)$ is the potential energy of the system for the n th replica with configuration X_n . For a protein system, the energy is composed of three terms:

$$E(X) = E_{pp}(X) + E_{pw}(X) + E_{ww}(X) \quad (2)$$

where E_{pp} , E_{pw} , and E_{ww} are, respectively, the internal energy of the protein, the interaction energy between the protein and water, and the self-interaction energy of the water molecules. Properly solvating a protein system typically requires several thousand water molecules; hence, the self-interaction between the water molecules usually vastly dominates the other terms, that is, $|E_{pp}|, |E_{pw}| \ll |E_{ww}|$.

The main disadvantage of the REM is that the system size causes the required number of replicas to increase rapidly,⁷ roughly as $O(f^{1/2})$, where f is the number of degrees of freedom in the system. However, the presence of the water self-interaction energy, E_{ww} , has two previously underappreciated benefits: (i) Changes in E_{pw} are partially compensated by opposite changes in E_{ww} , since both the magnitude of E_{pw} and the number of water molecules near the protein surface are roughly proportional to the protein surface area. This follows from conservation of the number of water molecules, since those not on the protein surface are by definition in the bulk, where they have stronger average interactions with other water molecules. (ii) Large spontaneous fluctuations in E_{ww} promote exchanges between energetically disparate folded (F) and unfolded (U) conformations at different temperature levels. This is relevant for systems where the average structural energies of folded (F) and unfolded (U) protein conformations are vastly different at neighboring temperatures, that is, $\Delta\beta|\langle E_{pp} + E_{pw} \rangle_{U,n} - \langle E_{pp} + E_{pw} \rangle_{F,m}| \gg 1$. While E_{ww} may reduce Δ_{nm} somewhat, as described above, it is often the fluctuations in the water self-interaction energy that make exchanges between energetically disparate conformations possible, that is,

$$|\langle E_{pp} + E_{pw} \rangle_{U,n} - \langle E_{pp} + E_{pw} \rangle_{F,m}| \ll \sigma_{U,n}(E_{ww}) + \sigma_{F,m}(E_{ww}) \quad (3)$$

To counter the poor scaling of the REM with system size, we recently introduced a new version⁸ of the REM, called REST. In this more general Hamiltonian replica exchange method,⁷ we modified the potential energy surface according to

$$E_m(X) = E_{pp}(X) + \left[\frac{\beta_0 + \beta_m}{2\beta_m} \right] E_{pw}(X) + \left[\frac{\beta_0}{\beta_m} \right] E_{ww}(X) \quad (4)$$

where p denotes the central group, w the bath, and $\beta_m = 1/(kT_m)$. When all of the water molecules are included in the bath (w), the large water self-interaction energy, E_{ww} , vanishes from the replica exchange acceptance criterion. Hence, $\Delta E = \tilde{E}(X_n) - \tilde{E}(X_m)$, where the effective energy is given by $\tilde{E}(X) \equiv E_{pp}(X) + (1/2)E_{pw}(X)$. As demonstrated in our original paper,⁸ this reduces the required number of replicas to $O(f_p^{1/2})$, where f_p is number of degrees of freedom in the central group. Our initial approach emphasized the case where only the protein is included in the central group, but we also discussed the possibility of including waters in the central group. The physical motivation for REST is that the water is already pliable at low temperatures and thus fast ensemble dynamics with “heated” water, that is, water configurations that are sampled according to $\exp\{-\beta_n E_{ww}(X_n)\}$, as in the REM, is overshadowed by the need for substantially more replicas. For systems where the unfolded structures vary greatly and large

conformational changes are required, however, standard REM sampling may be preferable. Sampling of folded structures should be nearly equivalent in the REM and the REST method, since the conformational changes are generally small and the Hamiltonians approach each other at low temperatures.

Contrary to our earlier observation,⁸ when the central group (p) only includes the protein, REST does not always have a speedup relative to the REM of $O(\sqrt{f/f_p})$. Rather, we find that REST is extremely inefficient in certain cases. For this reason, we have not sought to assess the decidedly secondary issue of the convergence rate of this implementation of REST versus the REM. Instead, we examine the deficiencies in the current REST implementation and when they might be important. Significantly, we believe that many of the applications envisioned in our original study,⁸ such as weakly bound protein–ligand systems, remain unaffected.

We attribute the observed inefficiency of our original implementation of REST, with only the protein in the central group (p), to the removal of the full water self-interaction energy, E_{ww} , from the acceptance probability. Specifically, systems for which the folded (F) and unfolded (U) conformations have vastly different effective energies, that is, $\Delta\beta|\langle E_{pp} + (\frac{1}{2})E_{pw}\rangle_{U,n} - \langle E_{pp} + (\frac{1}{2})E_{pw}\rangle_{F,m}| \gg 1$, now present an obstacle. If the fluctuations in the effective energy are substantially smaller than the size of the effective energy barrier, that is,

$$|\langle \tilde{E} \rangle_{U,n} - \langle \tilde{E} \rangle_{F,m}| \gg \sigma_{U,n}(\tilde{E}) + \sigma_{F,m}(\tilde{E}) \quad (5)$$

then exchanges between folded and unfolded conformations at different temperature levels become very infrequent. The failure to exchange folded and unfolded conformations causes the following symptom: When all replicas are started from a folded configuration, replicas at high temperatures unfold and replicas at low temperatures stay permanently folded. Some replicas at intermediate temperatures may sample both folded and unfolded conformations, but will exclusively *exchange* with the same type of conformation.

To provide insight into the role of the bath in replica exchange, we investigate an implementation of REST in which a randomly chosen subset of the water molecules are included in the central group, in addition to the protein. With N_p water molecules in the central group, this leaves $N_w = N - N_p$ waters in the bath (w). We believe that the optimal number of water molecules in the central group, N_p , is the minimum number that allows frequent exchanges between energetically disparate folded (F) and unfolded (U) conformations at different temperature levels. That may still be substantially smaller than the total number of water molecules and thus represent a significant improvement over the REM.

A final thought is to capture more of the compensating power of the water self-interaction energy, E_{ww} , without introducing a large number of additional replicas. Ideally, we would achieve this by including the first few solvation shells of water around the protein in the central group. However, as discussed in our original paper,⁸ the overhead for managing the associated neighbor list may outweigh the benefits. In addition, the number of degrees of freedom in the central group would increase as a result of the included water molecules.

Results

We have used the molecular dynamics program SIM,⁹ developed in our group, to investigate three solvated protein systems, as described below. In each case, we employ the OPLSAA force field¹⁰ for the protein and the P3ME method^{11–13} for the electrostatic interactions. After

a 200 ps equilibration simulation, starting from the native structure of each system, all replicas were simulated for 5 ns (3 ns for the β -hairpin system), with nearest neighbor replica exchanges every 0.4 ps (every 2 ps for the α -helix).

In this section, REST refers to the particular implementation of REST where only the protein is included in the central group (p).

α -Helix

The α -helix 3K(I)¹⁴ is a 16-residue peptide with the sequence AAAAKAAAAKAAAAKA. The N- and C-terminals are capped with Ace and Nme groups, respectively, and the three lysine residues are all positively charged. In addition to the α -helix, the simulation box has 2522 simple point charge (SPC)¹⁵ water molecules and three counterions (Cl^-), about 7700 atoms in total. Six replicas with temperatures of 310, 330, 365, 400, 441, and 488 K were used to generate acceptance ratios ranging from 15 to 30%. Each replica was simulated for 5 ns, with nearest neighbor replica exchanges attempted every 2 ps.

The distributions of the effective energy, $E_{\text{pp}} + (1/2)E_{\text{pw}}$, at the six different temperatures are shown in panel a of Figure 1. The large overlaps between the distributions at neighboring temperatures translate to generous acceptance ratios ranging from 15 to 30%, as mentioned above. In panel b of Figure 1, we display the temperature trajectories of four replicas and observe that they exhibit random walks spanning the whole temperature regime during the 5 ns simulation. This indicates that all ensembles exchange properly with each other. In Figure 2, we show the root mean square deviation (RMSD) of the C- α atoms from the ideal helix as a function of the simulation time and observe that, even at 310 K, the system samples both the folded (RMSD 2–3 Å) and unfolded (RMSD 5–7 Å) states with high frequency.

The temperature and RMSD curves demonstrate that the replicas exchange properly and thus, since REST satisfies detailed balance, the replica at 310 K will eventually sample the distribution of interest.

β -Hairpin

The C-terminus β -hairpin of protein G contains 16 residues with the sequence GEWTYDDATKTFTVTE. We have capped it with the normal Ace and Nme groups, for a total of 256 atoms. As this protein has been studied intensively with the REM by Zhou et al.,^{16,17} also using the OPLSAA force field,¹⁰ we will use their results as a reference for REST on this system. In addition to the β -hairpin, the simulation box contains 1361 SPC water molecules and 3 counterions (Na^+), for a total of 4342 atoms. Each replica was simulated for 3 ns, with nearest neighbor replica exchanges attempted every 0.4 ps.

With 18 replicas in the temperature range from 310 to 684 K, we obtained appreciable overlaps between the distributions of the effective energy, $E_{\text{pp}} + (1/2)E_{\text{pw}}$, as displayed in panel a of Figure 3. The concomitant acceptance ratios are between 20 and 50%, with the minimum value of around 20% achieved at 400 K. However, as shown in panel b, none of the four replicas displayed explore the entire temperature regime during our 3 ns simulation. Indeed, some replicas only sample the region below 400 K, while others always stay above 400 K. It appears that the temperature region around 400 K represents a bottleneck of some sort. This issue is illustrated further in Figure 4, where we investigate the RMSD from the native conformation at three different temperatures. At low temperatures, here 310 K, the hairpin stays folded, with a RMSD less than 4 Å during the whole simulation. At high temperatures, on the other hand, here 684 K, the hairpin quickly unfolds and then exclusively samples unfolded states with a RMSD greater than 5 Å for the rest of the simulation. Hence, folded structures are never

sampled at high temperatures, and vice versa. At intermediate temperatures, here 419 K, both folded and unfolded states with a RMSD in the range 1–8 Å are sampled.

It appears that REST does not sample the β -hairpin system efficiently, even though the exchange acceptance ratios are high. We verify this in Figure 5, where we compare the free energy landscape produced by REST (b) with a REM (a) reference simulation by Zhou et al.^{16,17} Here, the free energy landscape at 310 K is determined as a function of the number of β -sheet H bonds and the hydrophobic core radius of gyration. For easy comparison, both simulations were conducted for 3 ns. Clearly, REST only samples a small region around the folded state, while the standard REM samples a much larger region of phase space. In addition, while the REM results indicate that there is a local minimum at $N_{\text{Hb}}^{\text{core}}=0$ and $R_{\text{g}}^{\text{core}}=10$ Å, corresponding to an unfolded structure, this region is not sampled at all in REST.

TrpCage

TrpCage is a 20-residue miniprotein with the sequence NLYIQWLKDGGPSSGRPPPS, designed by Neidigh et al.¹⁸ The “native” protein structure was obtained by taking the NMR structure (PDB ID: 1L2Y, structure 1 of the total 38 NMR structures) and then subjecting it to a conjugate gradient minimization, followed by a 200 ps equilibration at 300 K. In addition to the TrpCage, the simulation box contains 3809 SPC water molecules and one counterion (Cl^-), for a total of 11 732 atoms. With 10 replicas at 300, 322, 345, 368, 394, 423, 455, 491, 529, and 572 K, we obtained acceptance ratios of around 20%, in accordance with the overlaps between the distributions of the effective energy, $E_{\text{pp}} + (1/2)E_{\text{pw}}$, displayed in panel a of Figure 6. Each replica was simulated for 5 ns, with nearest neighbor replica exchanges attempted every 0.4 ps.

The temperature trajectories of four replica walkers, neither of which manage to visit all of the temperature levels during the 5 ns simulation, are shown in panel b of Figure 6. We address this issue from another perspective in panel a of Figure 7, where we demonstrate that only folded structures are sampled at low temperatures, here 300 K. At high temperatures, here 572 K, folded structures are rarely sampled after a brief initial phase. Only at intermediate temperatures, here 423 K, do we observe a variety of folded and unfolded structures.

In summary, it is clear that REST does not sample the TrpCage system efficiently, as was the case for the β -hairpin system.

TrpCage with Random Central Waters

As suggested in the Introduction, we have modified our implementation of REST to include a randomly chosen subset of water molecules in the central group (p). We call this method REST with random central waters (RESTRCW). It is characterized by the number of waters in the central group, N_{p} , which are chosen at random from the simulation box at time $t = 0$ and stay in the central group during the whole simulation. Here, we perform RESTRCW simulations on the TrpCage system with the number of central water molecules, N_{p} , set to 200, 400, and 1000.

With 200 random central water molecules and a temperature range from 300 to 568 K, we found that 16 replicas are required. The replicas at low temperatures (300–408 K) are initially in the folded state, and the rest of the replicas start in an extended state. We show the RMS deviations from the native structure at different temperatures in panel b of Figure 7. As the extended structures are still not sampled at low temperatures (up to 345 K), it is clear that the inclusion of 200 random central water molecules is not enough. We have repeated the exercise with 400 random central water molecules and observe that the problem persists, as demonstrated in panel c of Figure 7. This system required 18 replicas. Finally, with 1000

random central water molecules, even the lowest temperature (300 K) samples extended structures, as displayed in panel d of Figure 7. However, we found that 30 replicas are required for this system, as the number of degrees of freedom has increased substantially. For comparison, a recent REM study¹⁹ of TrpCage required 50 replicas.

In conclusion, we find that the inclusion of random central water molecules can improve the performance of REST. It is clear that there is an optimal number of central water molecules, that is, such that the number of replicas is at a minimum, yet exchanges between energetically disparate folded (F) and unfolded (U) conformations at different temperature levels are frequent. We expect this optimal number of central water molecules to be substantially smaller than the total number of waters in the simulation box, and thus that REST with random central waters (RESTRCW) represents an improvement over the REM. However, this improvement comes at the expense of poorer scaling with system size.

Discussion

The motivation for the current study was to investigate three new solvated protein systems (α -helix, β -hairpin, and TrpCage) using our original implementation of REST, in which only the protein is included in the central group. REST was designed to reduce the required number of replicas from the $O(f^{1/2})$ of the REM to $O(f_p^{1/2})$, where f and f_p are the number of degrees of freedom in the system and in the central group, respectively. This follows directly from the modified replica exchange acceptance probability.⁸ In addition, we previously observed⁸ that the speedup versus the REM, in terms of converging to the correct underlying distribution, was $O(\sqrt{f/f_p})$. We have not consistently observed this speedup in the present study. Indeed, we find that REST performs significantly worse than the REM in certain cases.

We confirm that REST, with the protein as the central group, performs as expected on a solvated α -helix system, in agreement with our earlier study of a solvated alanine dipeptide.⁸ However, this implementation of REST is very inefficient on two larger systems, a solvated β -hairpin and a solvated TrpCage. We attribute the poor performance of this implementation of REST to the removal of the full water self-interaction energy from the replica exchange acceptance criterion. Specifically, we find that exchanges between folded (F) and unfolded (U) conformations with vastly different structural energies are greatly affected by the removal of the water self-interaction energy.

The presence of the water self-interaction energy in the REM has two previously underappreciated benefits: (i) Changes in the protein–water interaction are partially compensated by opposite changes in the water self-interaction, since both the magnitude of the protein–water interaction and the number of water molecules near the protein surface are roughly proportional to the protein surface area. This follows from conservation of the number of water molecules, since those not on the protein surface are by definition in the bulk, where they have stronger average interactions with other water molecules. (ii) Large spontaneous fluctuations in the water self-interaction energy promote exchanges between energetically disparate folded (F) and unfolded (U) conformations at different temperature levels. We believe that the lack of these two effects is responsible for REST's inability to sample the β -hairpin and TrpCage systems efficiently.

To verify this diagnosis, we have pursued an alternative implementation of REST, called REST with random central waters (RESTRCW), where a randomly chosen subset of water molecules are included in the central group. We find that, with a significant number of random central water molecules, RESTRCW regains the ability to exchange energetically disparate folded (F) and unfolded (U) conformations at different temperature levels. We believe that the optimal number of central water molecules is the minimum number that allows frequent exchanges

between energetically disparate folded (F) and unfolded (U) conformations at different temperature levels. That may still be substantially smaller than the total number of water molecules, and thus represent a significant improvement over the REM. The advantage of this approach is that the neighbor list of water molecules is static. Hence, it may compare favorably with our earlier suggestion⁸ to include the first few solvation shells of water around the protein in the central group. Which of the two approaches will prevail remains to be seen.

In conclusion, we have found that our original implementation of REST, with the protein as the central group, performs poorly on systems where the folded (F) and unfolded (U) conformations have vastly different structural energies. We emphasize that we do not expect this deficiency to affect REST on systems for which the difference in the effective conformational energy, $E_{pp} + (1/2)E_{pw}$, between the folded and unfolded states is on the order of the fluctuations in this quantity. When present, this deficiency can be corrected by including a sufficient number of water molecules in the central group. In this paper, we have demonstrated that REST performs well on systems like the α -helix and the dipeptide. We conjecture that REST is efficient on systems that satisfy the above condition, such as weak protein–ligand binding events.

Acknowledgments

The authors X.H., M.H., and B.K. contributed equally to this work. This work has been supported by grants to B.J.B. from the NSF (CHE16896), the NIH (GM43340), and the National Computational Science Alliance, using the Xeon Linux Cluster under grant MCA95C007N.

References and Notes

- Swendsen RH, Wang JS. Phys Rev Lett 1986;57:2607–2609. [PubMed: 10033814]
- Hukushima K, Nemoto K. J Phys Soc Jpn 1996;65:1604–1608.
- Marinari, E.; Parisi, G.; Ruiz-Lorenzo, JJ. World Scientific. Singapore: 1998. p. 59
- Sugita Y, Okamoto Y. Chem Phys Lett 1999;314:141–151.
- Garcia AE, Sanbonmatsu KY. Proteins 2001;42:345–354. [PubMed: 11151006]
- Rhee Y, Pande V. Biophys J 2003;84:775–786. [PubMed: 12547762]
- Fukunishi H, Watanabe O, Takada S. J Chem Phys 2002;116:9058–9067.
- Liu P, Kim BRF, Berne BJ. Proc Natl Acad Sci USA 2005;102:13749–13754. [PubMed: 16172406]
- Stern H, Rittner F, Pavese M, Harder E, Xu H, Kim B, Liu P. SIM: Molecular Dynamics Simulation Program. 2005
- Jorgensen W, Maxwell D, Tirado-Rives J. J Am Chem Soc 1996;118:11225–11236.
- Hockney, RW.; Eastwood, JW. Computer Simulations Using Particles. IOP; Bristol, U.K.: 1988.
- Deserno M, Holm C. J Chem Phys 1998;109:7678.
- Zhou R, Harder E, Xu H, Berne BJ. J Chem Phys 2001;115:2348.
- Smythe M, Nakie C, Marshall G. J Am Chem Soc 1995;117:10555.
- Berendsen, H.; Postma, J.; van Gunsteren, W.; Hermans, J. Intermolecular Forces. Reidel; Dordrecht, Holland: 1981.
- Zhou R, Berne B. Proc Natl Acad Sci USA 2002;99:12777–12782. [PubMed: 12242327]
- Zhou R, Krilov G, Berne B. J Phys Chem B 2004;108:7528–7530.
- Neidigh J, Fesinmeyer R, Andersen N. Nat Struct Biol 2002;9:425–430. [PubMed: 11979279]
- Zhou R. Proc Natl Acad Sci USA 2003;100:13280–13285. [PubMed: 14581616]

Figure 1.
 (a) Effective energy ($E_{pp} + (1/2)E_{pw}$) distributions at different temperatures for the α -helix system. (b) Temperature trajectories for four (out of six) replica walkers for the α -helix system.

Figure 2.
RMS deviations of REST from the native helix versus the simulation time for the trajectories at three temperatures (310, 400, and 488 K).

Figure 3.

(a) Effective energy ($E_{pp} + (1/2)E_{pw}$) distributions at different temperatures for the β -hairpin system. (b) Temperature trajectories for four (out of 18) replica walkers. None of them sample the whole temperature space.

Figure 4. RMS deviations of REST from the native structure of the β -hairpin peptide versus the simulation time for the trajectories at three temperatures (310, 419, and 684 K).

Figure 5.

Comparison of the free energy (in units of $-kT$) for the β -hairpin at 310 K versus the number of β -sheet H bonds, $N_{\text{HB}}^{\text{core}}$, and the hydrophobic core radius gyration, R_g^{core} . A hydrogen bond is counted if the distance between two heavy atoms (N and O) is less than 3.5 Å and the angle N–HO is larger than 150°. (a) REM; (b) REST. The standard REM results are reproduced from *Proc. Natl. Acad. Sci.* **2001**, 98 (26), 14931–14936.

Figure 6.

(a) Effective energy ($E_{pp} + (1/2)E_{pw}$) distributions at different temperatures for the TrpCage protein system. (b) Temperature trajectories for 4 (out of 10) replica walkers. The temperature series is the following: 300, 322, 345, 368, 394, 423, 455, 491, 529, and 572 K.

Figure 7.

(a) RMS deviations of REST from the native structure of the TrpCage protein versus the simulation time for the trajectories at three different temperatures (300, 423, and 572 K). (b) Same as part a but using REST with random central water at the three temperatures (300, 408, and 548 K) with 200 random central waters. (c) Same as part b but at 300, 412, and 558 K with 400 random central waters. (d) Same as part b but at 300, 405, and 570 K with 1000 random central waters.