

Demonstrated convergence of the equilibrium ensemble for a fast united-residue protein model

F. Marty Ytreberg

Department of Physics, University of Idaho, Moscow, ID 83844-0903

Svetlana Kh. Aroutiounian*

Department of Physics, Dillard University, 2601 Gentilly Blvd.,

New Orleans, LA 70122

Daniel M. Zuckerman†

Department of Computational Biology, University of Pittsburgh,

3501 Fifth Avenue, Pittsburgh, PA 15260

June 29, 2007

Abstract

Due to the time-scale limitations of all-atom simulation of proteins, there has been substantial interest in coarse-grained approaches. Some methods, like “Resolution Exchange,” [Lyman, E.; Ytreberg, F.M.; Zuckerman, D.M. *Phys. Rev. Lett.* **2006**, *96*, 028105–1-4] can accelerate canonical all-atom sampling, but require properly distributed coarse ensembles. We therefore demonstrate that full sampling can indeed be achieved in a sufficiently simplified protein model, as verified by a recently developed convergence analysis. The model accounts for protein backbone geometry in that rigid peptide planes rotate according to atomistically defined dihedral angles, but there are only two degrees of freedom (ϕ and ψ dihedrals) per residue. Our convergence analysis indicates that small proteins (up to 89 residues in our tests) can be simulated for more than 50 “structural decorrelation times” in less than a week on a single processor. We show that the fluctuation behavior is reasonable, as well as discussing applications, limitations, and extensions of the model.

*The first two authors contributed equally

†E-mail: dmz@ccbb.pitt.edu

1 Introduction

How simplified must a molecular model of a protein be for it to allow full canonical sampling? This question may be important to the solution of the protein sampling problem—the generation of protein structures properly distributed according to statistical mechanics—because of the well-known inadequacy of all-atom simulations, which are limited to sub-microsecond timescales. Even small peptides have proven slow to reach convergence.¹ Sophisticated atomistic methods, moreover, which often employ elevated temperatures,^{2–6} have yet to show they can overcome the remaining gap in timescales⁷—which is generally considered to be several orders of magnitude. On the other hand, because of drastically reduced numbers of degrees of freedom and smoother landscapes, coarse-grained models (e.g., refs 8–26) may have the potential to aid the ultimate solution to the sampling problem, particularly in light of recently developed algorithms like “Resolution Exchange”^{27,28} and related methods.^{29–31}

Although the Resolution Exchange approach, in principle, can produce properly distributed atomistic ensembles of protein configurations, it requires full sampling at the coarse-grained level.^{27,28} While the potential for such full sampling has been suggested by some studies of folding and conformational change (e.g., refs 26,32), convergence has yet to be carefully quantified in equilibrium sampling of folded proteins. How much coarse-graining really is necessary? What is the precise computational cost of different approaches? This report represents a first step in answering these questions by studying a united-residue model with realistic backbone geometry.

We will require a quantitative method for assessing sampling. A number of approaches have been suggested,^{1,33–36} but we rely on a recently proposed statistical approach which directly probes the fundamental configuration-space distribution.^{1,37} The method does not require knowledge of important configurational states or any parameter fitting. In essence, the approach attempts to answer the most fundamental statistical question, “What is the minimum time interval between snapshots so that a set of structures will behave as if each member were drawn independently from the configuration-space distribution exhibited by the full trajectory?” This interval is termed the “structural decorrelation time” τ_{dec} , and the goal is to generate simulations of length $t_{\text{sim}} \gg \tau_{\text{dec}}$.

In this report, we demonstrate the convergence of the equilibrium ensemble for several proteins using a fast, united-residue model employing rigid peptide planes. The relative motion of the planes is determined by the *atomistic* geometry embodied in the ϕ and ψ dihedral-angle rotations,




Figure 1: The rigid peptide plane model used this study. Note that, in the coarse-grained simulations, only alpha-carbons are represented, and the only degrees of freedom are ϕ and ψ . Other atoms are shown in the figure only to clarify the geometry and our assumption of rigid peptide planes.

as explained below. We believe such realistic backbone geometry will be necessary for success in Resolution Exchange studies. The use of geometric look-up tables enables the rapid use of only two degrees of freedom per residue (ϕ and ψ), and one interaction site at the alpha-carbon. The simulations are therefore extremely fast. $G\bar{o}$ interactions stabilize the native state while permitting substantial fluctuations in the backbone.

After the model and the simulation approach are explained, the fluctuations are compared with experimental data from X-ray temperature factors and the diversity of NMR structure sets. The simulations are then analyzed for convergence and timing.

2 Coarse-grained model

The coarse-grained model used for this study was chosen to meet several criteria: (i) the fewest number of degrees of freedom per residue; (ii) the ability to utilize lookup tables for enhanced simulation speed; (iii) the stability of the native state along with the potential for substantial non-native fluctuations; and, (iv) the ability to allow the addition of chemical detail, as simply as possible. Thus, we chose a rigid peptide plane model with $G\bar{o}$ interactions^{9,38,39} and sterics based on alpha-carbon interaction sites as shown in Figure 1. The use of such a simple model, we emphasize, is consistent with our goal of understanding both the potential, and the limitations

of coarse models for statistically valid sampling. Once we have understood the costs associated with the present model, we can design more realistic models, as discussed below. In other words, we made no attempt to design the most chemically realistic coarse-grained model, although we believe the use of atomistic peptide geometry is an improvement over a coarse model we considered previously.²⁶

Rigid peptide planes allow the use of only two degrees of freedom per residue, arguably the fewest that one would consider in such a model. Indeed, this is fewer than in a freely rotating chain, although admittedly our model requires somewhat more complex simulation moves, described below.

G \bar{o} interactions were used because they simultaneously stabilize the native state of the protein and also permit reasonable equilibrium fluctuations, as was shown in an earlier study.²⁶ Given our interest in native-state fluctuations and the lack of a *universal* coarse-grained model capable of stabilizing the native state for *any* protein, G \bar{o} interactions are a natural choice for enforcing stability. Further, beyond the reasonable “local” fluctuations shown below, the model also exhibits partial unfolding events which are expected both theoretically and experimentally.^{40–42}

Because we see the present model as only a first step in the development of better models, it is important that it easily allows for the addition of chemical detail, such as Ramachandran propensities which require only the dihedral angles we use explicitly.⁴³ Furthermore, with a rigid peptide plane, the locations of all backbone atoms—and the beta carbon—are known implicitly. Thus hydrogen-bonding and hydrophobic interactions¹⁵ can be included in the model with little effort. In other words, the “extendibility” of the present simple model was a significant factor in its design.

2.1 Potential energy of model system

The total potential used in the model is given by

$$U = U^{\text{nat}} + U^{\text{non}}, \quad (1)$$

where U^{nat} is the total energy for native contacts, and U^{non} is the total energy for non-native contacts.

For the G \bar{o} interactions, all residues that are separated by a distance *less* than a cutoff R_{cut} in the experimental structure are given native interaction

energies defined by a square well:

$$\begin{aligned}
 U^{\text{nat}} &= \sum_{i < j}^{\text{native}} u^{\text{nat}}(r_{ij}), \\
 u^{\text{nat}}(r_{ij}) &= \begin{cases} \infty & \text{if } r_{ij} < r_{ij}^{\text{nat}}(1 - \delta) \\ -\epsilon & \text{if } r_{ij}^{\text{nat}}(1 - \delta) \leq r_{ij} < r_{ij}^{\text{nat}}(1 + \delta) \\ 0 & \text{otherwise} \end{cases}, \quad (2)
 \end{aligned}$$

where r_{ij} is the $C_\alpha - C_\alpha$ distance between residue i and j , r_{ij}^{nat} is the the distance between the residues in the experimental structure, ϵ determines the energy scale of the native Gō attraction, and δ sets the width of the well. All residues that are separated by *more* than R_{cut} in the experimental structure are given non-native interaction energies defined by

$$\begin{aligned}
 U^{\text{non}} &= \sum_{i < j}^{\text{non-native}} u^{\text{non}}(r_{ij}), \\
 u^{\text{non}}(r_{ij}) &= \begin{cases} \infty & \text{if } r_{ij} < (\rho_i + \rho_j)(1 - \delta) \\ +h\epsilon & \text{if } (\rho_i + \rho_j)(1 - \delta) \leq r_{ij} < R_{\text{cut}} \\ 0 & \text{otherwise} \end{cases}, \quad (3)
 \end{aligned}$$

where ρ_i is the hard-core radius of residue i defined as half the C_α distance to the nearest non-covalently-bonded residue in the experimental structure, and h determines the strength of the repulsive interaction.

For this study, parameters were chosen to be similar to those in ref 26, i.e., $\epsilon = 1.0$, $h = 0.3$, $\delta = 0.1$, and $R_{\text{cut}} = 8.0 \text{ \AA}$.

2.2 Monte Carlo simulation

The protein fluctuations were generated using Metropolis Monte Carlo,⁴⁴ although discontinuous molecular dynamics could also be used.²³ Trial configurations were generated by adding a random Gaussian deviate to the values of three sequential pairs of backbone torsions (three ϕ and three ψ angles). We found that changing six sequential backbone torsions maximizes the rate of convergence of the equilibrium ensemble (data not shown). The energy of the trial configuration was then determined using eq (1), and the conformation was accepted with probability $\min(1, e^{-\Delta U/k_B T})$, where ΔU is the total change in potential energy of the system. The width of the Gaussian distribution for generating random deviates was chosen such that the acceptance ratio was about 40% for all simulations. The choice of temperature is discussed below.

2.3 Use of lookup tables

The speed of the coarse-grained simulation was enhanced by using lookup tables to avoid unnecessary computation. In general, lookup tables will increase memory usage while decreasing the number of computations. Since memory is inexpensive and can be expanded easily, utilizing as much memory as possible can be an effective technique for increasing the speed of simulations.

In our model there are only two degrees of freedom per residue (ϕ, ψ), but C_α distances r_{ij} must be computed to determine native and non-native interaction energies given by eqs (2) and (3). All peptide planes are considered to possess ideal, rigid geometry as determined by energy minimization of the all-atom OPLS forcefield⁴⁵ using the TINKER simulation package.⁴⁶

Given a sequence of three residues (alpha carbons), we employed a lookup table to provide the Cartesian coordinates of the third residue—starting from the N-terminus—and its normal vector as a function of ϕ and ψ ; see Figure 1. The table values assume that the first residue is at the origin and the second residue is located on the z-axis. Once the coordinates for the third residue were determined via the lookup table, the fourth residue position was determined using the lookup table in conjunction with a coordinate rotation and shift. Continuing in this fashion, coordinates for the entire protein were determined.

The resolution of the lookup table is an important consideration, i.e., the number of ϕ, ψ values for which Cartesian coordinates are stored. In our simulations, we tried resolutions as high as 0.1° and as low as 1.0° , and found no difference between the results. Thus, all simulation results presented here use lookup tables with a resolution of 1.0° .

2.4 Initial protein relaxation

One perhaps unexpected complication of utilizing a rigid peptide plane model is that great care must be taken to relax the protein before simulations can be performed. Although initial values of ϕ, ψ are obtained from the X-ray or NMR structure, there are slight deviations from planar/ideal geometry in a real protein. These deviations, while small, can accumulate rapidly to become very large differences in the Cartesian coordinate positions of the residues. Thus, the positions of residues near the beginning of the protein will be nearly correct, while the residues near the end of the protein will likely have large errors—compared to the experimental structure being modeled—which can create severe steric clashes or even incorrect

protein topology. The severity of these “errors” necessitates the use of a relaxation procedure to generate a suitable starting structure—i.e., a set of ϕ and ψ angles which, with our ideal-geometry peptide planes, lead to a topologically reasonable and relatively clash-free structure. Proteins larger than 100 residues are difficult to relax.

Before we detail our relaxation procedure, we note that the need for this additional calculation is an artifact of the simplicity of our model—which can be overcome. For instance, it is possible to include *flexible* peptide planes with the use lookup tables (libraries of plane configurations, in this case) without significantly increasing the computational cost of the model. Such an approach, which does not require initial relaxation, is currently under investigation with promising preliminary results (data not shown).

The relaxation procedure employed in the present study first uses the ϕ, ψ values directly obtained from the experimental structure. These dihedrals provide the initial (problematic) structure for a coarse-grained simulation. Due to the deviations from planarity described above, the root-mean-square deviation (RMSD) between the initial structure we create and the experimental structure tends to be large (~ 10 Å was not uncommon for the proteins in this study). To increase the number of native contacts and reduce the number of steric clashes, we next performed what we term “RMSD Monte Carlo” to relax the protein to a low RMSD structure. Trial moves for RMSD Monte Carlo were created as described above, but accepted with probability $\min(1, e^{-\Delta(\text{RMSD})/k_B T_{\text{RMSD}}})$, where $k_B T_{\text{RMSD}} = 10^{-7}$ was chosen so that moves to a higher RMSD were rare. In other words, the energy function itself was not used in this initial phase.

Since residues near the beginning of the protein have less error in the starting structure than residues near the end, we used RMSD Monte Carlo in segments. The first twenty residues were relaxed until the RMSD was constant within a tolerance of 0.0001 Å, followed by the first forty, then the first sixty and so on until the RMSD of the entire protein was relaxed. The RMSD Monte Carlo simulation typically brought the RMSD of the simulated structure to less than 0.5 Å, however, there were generally still steric clashes, and some native contacts were still not present.

The final stage of relaxation was to do regular Metropolis Monte Carlo simulation (i.e., using energy), with a very low temperature. The low temperature Monte Carlo eliminated all steric clashes and allowed the remaining native contacts to form.

Relaxation was performed until four criteria were met: (i) the number of native contacts in the relaxed structure was equal to that in the NMR or X-ray structure; (ii) no steric clashes were present; (iii) no non-native

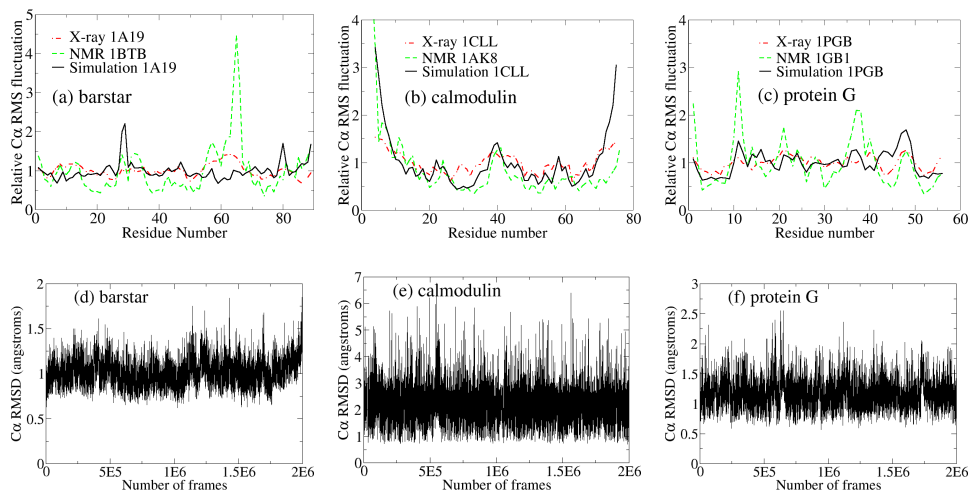


Figure 2: (Color online) Relative alpha-carbon root mean square fluctuations for three different proteins: (a) barstar, (b) calmodulin, and (c) protein G. Each plot shows results from X-ray data (dot-dash), the NMR ensemble (dash), and the coarse-grained simulation (solid). X-ray results were given by $\sqrt{3B/8\pi^2}$, where B is the temperature factor given in the PDB entry. NMR and simulation data were generated using the `g_rmsf` program in the GROMACS molecular simulation package;⁴⁷ each ensemble was aligned to the first structure in the corresponding trajectory. For each coarse-grained simulation, 2×10^9 Monte Carlo steps were performed with snapshots saved every 1000 steps, and the potential energy (1) was set up using the X-ray structure. Panels (d) - (f) show the corresponding whole-structure fluctuations as indicated by the RMSD from the experimental structures.

contacts were present, i.e., $U^{\text{non}} = 0$ in eq (3), and; (iv) the RMSD was less than 1.0 Å. When these criteria were met the structure was saved and used as the starting configuration in all future simulations of the protein.

3 Results and Discussion

Using the coarse-grained protein model described above, we generated and tested equilibrium ensembles for three proteins: barstar (PDB entry 1A19, residues 1-89), the N-terminal domain of calmodulin (PDB entry 1CLL, residues 4-75), and the binding domain of protein G (PDB entry 1PGB, residues 1-56)

For each protein, the initial simulation structure was generated, followed by RMSD and energy relaxation, as described above. Then, production runs of 2×10^9 Monte Carlo moves were performed with snapshots saved every 1000 moves, generating an equilibrium ensemble with 2×10^6 frames.

In an attempt to obtain consistent results for the three proteins, we chose the temperature of the simulation, $k_B T$, to be slightly below the unfolding temperature of the protein. The unfolding temperature was determined by running simulations over a broad range of temperatures and studying the RMSD as a function of simulation time. The temperatures used in the simulations were $k_B T = 0.6$ for barstar, $k_B T = 0.4$ for calmodulin and $k_B T = 0.5$ for protein G.

3.1 Speed of simulations

Due to the use of lookup tables for coordinate transformations, the small number of degrees of freedom, and the use of simple square potentials, equilibrium ensembles were generated very rapidly.

Running on one Xeon 2.4 GHz processor, 2×10^9 Monte Carlo moves with snapshots saved every 1000 steps took roughly 6 days for barstar, 4 days for calmodulin, and 3 days for protein G. Thus, less than a week was required to obtain well-converged simulations of these coarse-grained proteins.

The unfolding temperature for each protein was determined via short simulations of 10^7 Monte Carlo moves for around twenty different temperatures. This additional computational cost to determine the unfolding temperature was roughly 14 hours for barstar, 10 hours for calmodulin and 7 hours for protein G.

3.2 Protein fluctuations

We first sought to determine whether fluctuations in the coarse-grained simulations are reasonable, bearing in mind that the model was designed for speed rather than chemical accuracy. Figure 2 shows the alpha-carbon relative root-mean-square fluctuation for three different proteins. Given the model's minimalist design, the figures show that there is reasonable qualitative agreement between the NMR, X-ray and simulation data — with the notable exception of barstar, where the experimental peak at around residue 65 does not appear in the simulated results. Interestingly, for calmodulin, a standard Pearson correlation analysis indicates that the simulation data agree better with *both* experiments than the experiments do with one another.

We emphasize that the purpose of the current study is to understand the promise and limitations of coarse models for statistically valid sampling. We made no attempt to design the most chemically realistic coarse-grained model, and therefore sought only rough qualitative agreement between experimental and simulated results. The protein systems reported here were not “cherry picked” but intended to provide a representative picture of results obtainable in a minimalist model.

In addition, it should be noted that *none* of the three data sets in Figures 2a, b and c represents the true fluctuations in the protein—for different reasons. The X-ray temperature factor, in addition to thermal fluctuations, includes crystal lattice artifacts and other experimental errors.⁴⁸ NMR ensembles tend to be biased, perhaps severely, toward low energy structures, and thus also do not represent equilibrium ensembles.⁴⁹ Finally, our simulation data is not accurate due to the lack of chemical detail in the forcefield.

The bottom panels of Figure 2 show the whole-molecule fluctuations exhibited throughout the trajectories. In addition to the ability to sample large conformational fluctuations—such as in the case of calmodulin and, to a lesser degree, for protein G—the trajectories are visibly more converged than is typically observed in atomistic simulations, where RMSD values rarely reach a plateau value, let alone sampling around that plateau value multiple times as would be desirable.

3.3 Convergence analysis

The primary purpose of this report is to demonstrate the convergence of the equilibrium ensemble for a coarse-grained protein. The details of the convergence analysis are described in ref 37, so we will only briefly describe the method here.

Previously, Lyman and Zuckerman¹ developed an approach which classifies sampled conformations into the bins of a “structural histogram” using the RMSD as a metric. While promising, the primary limitation of the method was the lack of a quantitative measure of the convergence.

In the method used here, convergence was analyzed by studying the variance of the structural histogram bin populations.³⁷ The new approach allows a rigorous *quantitative* estimation of convergence—via the “structural decorrelation time” τ_{dec} , given by the time between frames required for the variance to reach an analytically computable independent-sampling value. Intuitively, and mathematically, τ_{dec} is the time interval between snapshots for which they behave as if each frame were sampled independently. If simulation times $t_{\text{sim}} \gg \tau_{\text{dec}}$ are obtained, the equilibrium ensemble is considered




Figure 3: Convergence analysis for coarse-grained simulations of three different proteins: (a) barstar, (b) calmodulin, and (c) protein G. Each plot shows the convergence properties for the same trajectories as used for Figure 2, analyzed using the procedure in ref 37. The number of frames required to reach the value of one (the solid horizontal line) is an approximation of the structural decorrelation time τ_{dec} and is shown on each plot. The three curves on each plot are results for different histogram sub-sample sizes³⁷ and demonstrates the robustness of the value of τ_{dec} . The plots predict that the decorrelation times are roughly 40 000 frames for barstar, 20 000 frames for calmodulin and 30 000 frames for protein G. Note that the total number of frames generated for each protein during the simulation was 2×10^6 . Thus, since each simulation was more than $50\tau_{\text{dec}}$ in length, we conclude that the equilibrium ensembles are well-converged. Error bars represent 80% confidence intervals in the expected fluctuations around the ideal value of one, based on the given trajectory length and the numerical procedure used to generate the solid curve.

converged.

Perhaps the most important feature of the convergence analysis for our study is that the method does not require any prior knowledge of important states. Furthermore, there is no parameter-fitting or subjective analysis of any kind.

Figure 3 shows the convergence properties of the coarse-grained simulations using the same trajectories as in Figure 2. The ratio of the observed variance to the ideal variance for independent sampling is plotted as a function of the time between the configurations used to compute the observed variance. When this ratio decreases to one the structural decorrelation time τ_{dec} has been reached, as shown in the figure. The analysis indicates that each simulation is at least 50 times longer than the structural decorrelation time.

We conclude that, in less than a week of single-processor time, the equi-

librium ensembles for these three proteins are well converged.

4 Conclusions

We have demonstrated the convergence of the equilibrium ensemble for a simple united-residue protein model. The model assumes rigid peptide planes, with atomistically correct geometry, and exhibits reasonable residue-level fluctuations based the planes’ geometry, G \ddot{o} interactions, and sterics.

Most importantly, the results indicate *quantitatively* that carefully designed united-residue models have the potential to fully sample protein fluctuations. By using only two degrees of freedom per residue, look up tables for coordinate transforms, and simple square well potentials, we were able to demonstrate that converged equilibrium ensembles can be obtained in less than a week of single processor time. The quantitative convergence analysis indicates that more than 50 “decorrelation times” were simulated in each case, indicating high-precision ensembles. In addition to application in Resolution Exchange sampling of all-atom models,^{27,28} such speed opens up the long-term possibility of large-scale simulation of many proteins.

One important practical limitation of the ideal-peptide-plane geometry in the present model is the need to relax the the initial structure. Proteins larger than 100 residues are difficult to relax. However, we have already begun investigating a flexible-plane model, incorporating pre-calculated libraries of plane configurations, which exhibits no such limitation and remains computationally affordable. We will report on the flexible model in the future.

Although the intrinsic atomistic geometry of the peptide plane was included in our model, it lacks chemical interactions. Yet because we obtained converged ensembles in such a short time, it is clear we can “afford” extensions to the model which include realistic chemistry. For instance, additional potential energy terms such as Ramachandran propensities,⁴³ hydrophobic interactions¹⁵ and hydrogen-bonding can be included at small cost.

Aside from the potential for rigorous atomistic sampling,^{27,28,50} it is important to note the general usefulness of coarse-grained models for generating *ad hoc* atomistic ensembles. Specifically, upon generating a well-sampled ensemble of coarse-grained structures, atomic detail can be added using existing software such as those in refs 51,52. Once minimized and relaxed, these (now) atomically detailed structures form an *ad hoc* ensemble which may be of immediate use in docking^{53,54} and homology modeling applications. Further, in principle, such structures can be re-weighted into the

Boltzmann distribution.⁵⁰

In the long term, one can imagine a day when structural databases will be based not on single (static) structures but rather will collect ensembles—as envisioned in the authors’ scheme for an “Ensemble Protein Database” (<http://www.epdb.pitt.edu/>).

Acknowledgments

We thank Edward Lyman, Bin Zhang, Artem Mamonov, and David Richardson for helpful discussions. Funding was provided by the National Institutes of Health under fellowship GM073517 (to F.M.Y.), and grants GM070987 and ES007318, and by the National Science Foundation grant MCB-0643456.

References

- (1) Lyman, E.; Zuckerman, D. M. *Biophys. J.* **2006**, *91*, 164–172.
- (2) Swendsen, R. H.; Wang, J.-S. *Phys. Rev. Lett.* **1986**, *57*, 2607–2609.
- (3) Hukushima, K.; Nemoto, K. *J. Phys. Soc. Jpn.* **1996**, *65*, 1604–1608.
- (4) Hansmann, U. H. E. *Chem. Phys. Lett.* **1997**, *281*, 140–150.
- (5) Sugita, Y.; Okamoto, Y. *Chem. Phys. Lett.* **1999**, *314*, 141–151.
- (6) Paschek, D.; Garcia, A. E. *Phys. Rev. Lett.* **2004**, *93*, 238105-1–4.
- (7) Zuckerman, D. M.; Lyman, E. *J. Chem. Theory and Comput.* **2006**, *2*, 1200–1202 Erratum: **2**, 1693 (2006).
- (8) Levitt, M.; Warshel, A. *Nature* **1975**, *253*, 694–698.
- (9) Ueda, Y.; Taketomi, H.; Gō, N. *Int. J. Peptide Protein Res.* **1975**, *7*, 445–459.
- (10) Tanaka, S.; Scheraga, H. A. *Proc. Natl. Acad. Sci. USA* **1975**, *72*, 3802–3806.
- (11) Kuntz, I. D.; Crippen, G. M.; Kollman, P. A.; Kimmelman, D. J. *Molec. Bio.* **1976**, *106*, 983–994.
- (12) Miyazawa, S.; Jernigan, R. L. *Macromol.* **1985**, *18*, 534–552.
- (13) Skolnick, J.; Kolinski, A.; Yaris, R. *Proc. Natl. Acad. Sci. USA* **1988**, *85*, 5057–5061.
- (14) Friedrichs, M. S.; Wolynes, P. G. *Science* **1989**, *246*, 371–373.
- (15) Lau, K. F.; Dill, K. A. *Macromolecules* **1989**, *22*, 3986–3997.
- (16) Honeycutt, J. D.; Thirumalai, D. *Proc. Natl. Acad. Sci. USA* **1990**, *87*, 3526–3529.
- (17) Monge, A.; Lathrop, E. J. P.; Gunn, J. R.; Shenkin, P. S.; Friesner, R. A. *J. Molec. Bio.* **1995**, *247*, 995–1012.
- (18) Jernigan, R. L.; Bahar, I. *Curr. Op. Struct. Bio.* **1996**, *6*, 195–209.
- (19) Zhou, Y.; Karplus, M. *Proc. Natl. Acad. Sci. USA* **1997**, *94*, 14429–14432.

- (20) Liwo, A.; Oldziej, S.; Pincus, M. R.; Wawak, R. J.; Rackovsky, S.; Scheraga, H. A. *J. Comput. Chem.* **1997**, *18*, 849–873.
- (21) Liwo, A.; Pincus, M. R.; Wawak, R. J.; Rackovsky, S.; Oldziej, S.; Scheraga, H. A. *J. Comput. Chem.* **1997**, *18*, 874–887.
- (22) Clementi, C.; Jennings, P. A.; Onuchic, J. N. *Proc. Natl. Acad. Sci. USA* **2000**, *97*, 5871–5876.
- (23) Smith, A. V.; Hall, C. K. *Proteins* **2001**, *44*, 344–360.
- (24) Shimada, J.; Shakhnovich, E. I. *Proc. Natl. Acad. Sci. USA* **2002**, *99*, 11175–11180.
- (25) Izvekov, S.; Parrinello, M.; Burnham, C. J.; Voth, G. A. *J. Chem. Phys.* **2004**, *120*, 10896–10913.
- (26) Zuckerman, D. M. *J. Phys. Chem. B* **2004**, *108*, 5127–5137.
- (27) Lyman, E.; Ytreberg, F. M.; Zuckerman, D. M. *Phys. Rev. Lett.* **2006**, *96*, 028105-1–4.
- (28) Lyman, E.; Zuckerman, D. M. *J. Chem. Theory Comp.* **2006**, *2*, 656–666.
- (29) Lwin, T. Z.; Luo, R. *J. Chem. Phys.* **2005**, *123*, 194904-1–10.
- (30) Christen, M.; van Gunsteren, W. F. *J. Chem. Phys.* **2006**, *124*, 154106-1–7.
- (31) Liu, P.; Voth, G. A. *J. Chem. Phys.* **2007**, *126*, 045106-1–6.
- (32) Clementi, C.; Nymeyer, H.; Onuchic, J. N. *J. Molec. Bio.* **2000**, *298*, 937–953.
- (33) Karpen, M. E.; Tobias, D. J.; Brooks III, C. L. *Biochemistry* **1993**, *32*, 412–420.
- (34) Straub, J. E.; Rashkin, A. B.; Thirumalai, D. *J. Am. Chem. Soc.* **1994**, *116*, 2049–2063.
- (35) Smith, L. J.; Daura, X.; van Gunsteren, W. F. *Proteins* **2002**, *48*, 487–496.
- (36) Elmer, S. P.; Pande, V. S. *J. Chem. Phys.* **2004**, *121*, 12760–12771.

- (37) Lyman, E.; Zuckerman, D. M. **2007**, arXiv:q-bio/0607037. arXiv.org ePrint archive. <http://arxiv.org/abs/q-bio/0607037> (accessed Feb 18, 2007).
- (38) Ueda, Y.; Taketomi, H.; Gō, N. *Biopolymers* **1978**, *17*, 1531–1548.
- (39) Gō, N.; Taketomi, H. *Proc. Natl. Acad. Sci. USA* **1978**, *75*, 559–563.
- (40) Careaga, C. L.; Falke, J. J. *J. Mol. Biol.* **1992**, *226*, 1219–1235.
- (41) Bai, Y.; Sosnick, T. R.; Mayne, L.; Englander, S. W. *Science* **1995**, *269*, 192–197.
- (42) Eisenmesser, E. Z.; Millet, O.; Labeikovsky, W.; Korzhnev, D. M.; Wolf-Watz, M.; Bosco, D. A.; Skalicky, J. J.; Kay, L. E.; Kern, D. *Nature* **2005**, *438*, 117–121.
- (43) Lovell, S. C.; Davis, I. W.; Arendall III, W. B.; de Bakker, P. I. W.; Word, J. M.; Prisant, M. G.; Richardson, J. S.; Richardson, D. C. *Proteins* **2003**, *50*, 437–450.
- (44) Metropolis, N.; Rosenbluth, A. W.; Rosenbluth, M. N.; Teller, A. H.; Teller, E. *J. Chem. Phys.* **1953**, *21*, 1087–1092.
- (45) Jorgensen, W. L.; Maxwell, D. S.; Tirado-Rives, J. *J. Am. Chem. Soc.* **1996**, *117*, 11225–11236.
- (46) Ponder, J. W.; Richard, F. M. *J. Comput. Chem.* **1987**, *8*, 1016–1024 <http://dasher.wustl.edu/tinker> (accessed Jan 15, 2005).
- (47) Van Der Spoel, D.; Lindahl, E.; Hess, B.; Groenhof, G.; Mark, A. E.; Berendsen, H. J. C. *J. Comput. Chem.* **2005**, *26*, 1701–1718.
- (48) Northrup, S. H.; Pear, M. R.; McCammon, J. A.; Karplus, M.; Takano, T. *Nature* **1980**, *287*, 659–660.
- (49) Spronk, C. A. E. M.; Nabuurs, S. B.; Bonvin, A. M. J. J.; Krieger, E.; Vuister, G. W.; Vriend, G. *J. Biomolec. NMR* **2003**, *25*, 225–234.
- (50) Ytreberg, F. M.; Zuckerman, D. M. **2006**, arXiv:physics/0609194. arXiv.org ePrint archive. <http://arxiv.org/abs/physics/0609194> (accessed Sept 22, 2006).
- (51) Eyal, E.; Najmanovich, R.; McConkey, B. J.; Edelman, M.; Sobolev, V. *J. Comput. Chem.* **2004**, *25*, 712–724.

- (52) de Bakker, P. I. W.; DePristo, M. A.; Burke, D. F.; Blundell, T. L. *Proteins Struct. Funct. Genet.* **2002**, *51*, 21–40.
- (53) Knegt, R. M. A.; Kuntz, I. D.; Oshiro, C. M. *J. Molec. Bio.* **1997**, *266*, 424–440.
- (54) Shoichet, B. K. *Nature* **2004**, *432*, 862–865.