

Temperature Weighted Histogram Analysis Method, Replica Exchange, and Transition Paths[†]

Emilio Gallicchio, Michael Andrec, Anthony K. Felts, and Ronald M. Levy*

Department of Chemistry and Chemical Biology, and BioMaPS Institute of Quantitative Biology, Rutgers University, Piscataway, New Jersey 08854

Received: October 14, 2004; In Final Form: November 22, 2004

We analyzed the data from a replica exchange molecular dynamics simulation using the weighted histogram analysis method to combine data from all of the temperature replicas (T-WHAM) to obtain the room-temperature potential of mean force of the G-peptide (the C-terminal β -hairpin of the B1 domain of protein G) in regions of conformational space not sampled at room temperature. We were able to determine the potential of mean force in the transition region between a minor α -helical population and the major β -hairpin population and identify a possible transition path between them along which the peptide retains a significant amount of secondary structure. This observation provides new insights into a possible mechanism of formation of β -sheet secondary structures in proteins. We developed a novel Bayesian statistical uncertainty estimation method for any quantity derived from WHAM and used it to validate the calculated potential of mean force. The feasibility of estimating regions of the potential of mean force with unfavorable free energy at room temperature by T-WHAM analysis of replica exchange simulations was further tested on a system that can be solved analytically and presented some of the same challenges found in more complex chemical systems.

1. Introduction

Elucidating protein folding pathways in atomic detail is one of the unsolved problems in computational biophysics. The key dynamic processes are difficult to study by computer simulations because they occur on a much longer time scale than it is possible to achieve with standard simulation techniques and because it can be difficult to identify a set of meaningful reaction coordinates. Transition path sampling methods have been developed to search for and characterize such rare events;^{1,2} application of these methods to realistic biological problems is beginning to appear.^{3,4}

The potential of mean force (PMF) yields primarily thermodynamic information but, if expressed along meaningful reaction coordinates, can also be useful for characterizing the kinetics of biological systems.^{5–7} The other challenge is to sample the conformational space of the system adequately enough to generate an accurate PMF. High free energy barriers hinder most sampling protocols. One method capable of surmounting energetic barriers separating conformational minima is the replica exchange molecular dynamics method (REMD).⁸ To accomplish barrier crossings, REMD simulates a series of replicas in a generalized ensemble over a range of temperatures. The replicas at higher temperatures have enough thermal energy to cross barriers. Periodically, coordinates are exchanged using a Metropolis criterion that ensures at any given temperature a canonical distribution is realized. Typically, the results at a given temperature are analyzed using data from only that temperature.^{9–11} Due to the finite number of samples, the range of free energy is limited. This poses a problem when trying to resolve the PMF of high free energy transition regions that separate the low free energy basins. Information about the barriers is contained in the higher temperature data, and, vice versa, more information about the low energy basins can be gleaned from the lower temperature data.

The replicas of higher or lower temperatures than the temperature of interest can be thought of as temperature biased simulations. The weighted histogram analysis method (WHAM)¹² is a general technique to combine data from multiple biased simulations to obtain thermodynamic properties of the unbiased system. WHAM is often used to obtain a PMF from histograms of the reaction coordinate collected from multiple simulations of the system. In each simulation, the system is subjected to a biasing potential designed to restrict sampling of the reaction coordinate in a predetermined region in such a way that over all the simulations good coverage of the reaction coordinate is achieved.^{13,14} Each simulation can also be carried out at a different temperature in addition to a different biasing potential.^{15,16} In fact, one of the earliest applications of WHAM was for the calculation of the density of states of a system over a wide spectrum of energies by combining histograms of the energy from simulations at different temperatures.¹⁷ When the temperature is the only biasing quantity, we refer to the WHAM method as T-WHAM.

Normally the reaction coordinate of interest is also the one on which the biasing potential is based. But it is also possible to use WHAM to unbiased a quantity on which the biasing potential does not explicitly depend.¹⁸ Similarly, when the temperature is the only biasing quantity, T-WHAM can be used not only to obtain the density of states but also to obtain the unbiased PMF along any reaction coordinate.¹⁵ By gathering information from all temperatures of a replica exchange simulation and using it to generate the room-temperature PMF, we expect to be able to expand the observable free energy range and to determine the structure of free energy barriers that are not apparent when using data from a single temperature.

In this paper, we develop a novel Bayesian statistical uncertainty estimation protocol for any quantity derived from WHAM. The WHAM solution corresponds to the maximum of a distribution function for the unbiased population.¹⁸ Using Bayesian techniques, we show that the spread of this distribution

* Corresponding author. E-mail: ronlevy@lutece.rutgers.edu

[†] Part of the special issue "David Chandler Festschrift".

around the maximum is a measure of the statistical uncertainty of the WHAM solution. Such a tool is very useful to validate T-WHAM results from replica exchange data. Statistical noise from histograms collected at temperatures much different than the reference may be greatly amplified. It is therefore necessary to evaluate the statistical uncertainty of the resulting potential of mean force to validate the estimates of the PMF in regions not sampled at the reference temperature.

The test systems we have chosen to demonstrate the utility of T-WHAM for combining data from multiple temperatures are a simple model potential, which can be solved analytically, and the C-terminal peptide of the B1 domain of protein G. This peptide (referred to as the G-peptide¹⁹) has been extensively studied experimentally^{10–22} and theoretically.^{9,10,23–32} This β -hairpin-forming peptide is stabilized by both the formation of secondary structure and the association of hydrophobic residues.³³ There has been great interest in the mechanism of folding of the G-peptide and on the role intermediates may play in the folding process. According to one model,³³ this peptide folds in a zipperlike mechanism, where the β -turn folds first, followed by the sequential formation of β -sheet backbone hydrogen bonds starting from the turn and proceeding toward the termini, bringing the hydrophobic core together. Other models predict that the formation of the hydrophobic core occurs before^{3,4} or simultaneously^{19,22} with the formation of the β -sheet backbone hydrogen bonds.

It was first noted by García and Sanbonmatsu⁹ using REMD simulations in explicit solvent that a small population of α -helical conformations of the G-peptide coexists alongside the major population of β -hairpin conformations. Analysis of our REMD results of the G-peptide in implicit solvent¹¹ shows a similar distribution of conformations and, moreover, that the α -helical population increases at intermediate temperatures higher than room temperature.³⁴ This observation raises the possibility that the α -helix conformational state is an on-path intermediate for some β -hairpin folding paths. If that were the case, then there should be a transition path connecting the α -helical state with the β -hairpin state without complete unfolding of the peptide. In this paper, we demonstrate the usefulness of T-WHAM by obtaining a PMF of the G-peptide that exhibits distinct minima corresponding to α -helix and β -hairpin states and, importantly, also a saddle point of relatively low free energy connecting these minima, suggesting that a suitable mechanism of transition from the α -helix state to the β -hairpin state might exist. We show that by using data only from the room-temperature REMD replica the region of conformational space between the α -helix and β -hairpin regions is unresolved, whereas using T-WHAM to combine the data from the high temperature replicas makes it possible to resolve the potential of mean force in the transition region. By using this PMF as a guide, we formulate a thermodynamically viable mechanism of interconversion between the α -helix and β -hairpin states.

A mechanism similar to the one proposed here may be involved in the formation of β -sheet secondary structures through α -helix-rich intermediates.^{35–40} The interconversion between α -helix and β -hairpin suggests that for some sequences these secondary structures have similar free energies. The ability to interconvert reflects the fact that α -helical and β -sheet secondary structural motifs correspond to competing local free energy minima because the backbone interactions contribute about equally to the free energy. The different interactions involving side chains in the α and β conformations determines which one is more thermodynamically stable.

2. Theory

In this section, we derive the WHAM equations in the context of simultaneous temperature and coordinate bias and show how the uncertainties of the resulting unbiased probabilities can be determined. We follow the maximum likelihood approach used by Bartels and Karplus¹⁸ and extend it to include a full Bayesian treatment of the uncertainties.

Suppose that we have performed S simulations, i ($i = 1, \dots, S$), each at a different inverse temperature β_i and possibly with different biasing potentials $w_i(x)$, where x is a reduced coordinate of interest. We discretize both the reduced coordinate x and the unbiased potential energy E into bins x_j ($j = 1, \dots, B$) and E_k ($k = 1, \dots, \Xi$), respectively. For each simulation, we construct a two-dimensional histogram with respect to the reduced coordinate and the energy and represent the count in the j th x bin and k th E bin in the i th simulation as $n_{i,(j,k)}$. We assume that B and Ξ have been chosen to be sufficiently large so that the resulting histogram can be treated as an accurate estimator of the probability density $P(x, E)$. The probability of finding the system in bin (j, k) during the i th simulation can be written as

$$p_{i,(j,k)} = f_i c_{i,(j,k)} p_{j,k}^{\circ} \quad (1)$$

where $p_{j,k}^{\circ}$ is the unbiased probability of that bin at the reference inverse temperature, β_0

$$c_{i,(j,k)} = \exp[-(\beta_i - \beta_0)E_k] \exp[-\beta_i w_i(x_j)] \quad (2)$$

is the bias factor which accounts for the effect of temperature and any biasing potentials, and f_i is a normalizing constant chosen to ensure that $\sum_{j,k} p_{i,(j,k)} = 1$ (i.e., $f_i^{-1} = \sum_{j,k} c_{i,(j,k)} p_{j,k}^{\circ}$).

To simplify the following equations and to emphasize the fact that the underlying problem can be recast in a one-dimensional form, we will adopt the following change of notation: each occurrence of (j, k) referring to a coordinate bin j and an energy bin k will be replaced by a single index l given by $l = j + B(k - 1)$ ($l = 1, \dots, m$, $m = B\Xi$). In other words, we map the two-dimensional coordinate/energy array into an equivalent one-dimensional vector. Therefore, eqs 1 and 2 can be rewritten as

$$p_{i,l} = f_i c_{i,l} p_l^{\circ} \quad (3)$$

and

$$c_{i,l} = \exp[-(\beta_i - \beta_0)E_l] \exp[-\beta_i w_i(x_l)] \quad (4)$$

where the expressions x_l and E_l refer to the values of the reduced coordinate of bin j and the energy of bin k .

Note that the subsequent derivation is unchanged if there are multiple reduced coordinates, if the temperature is not used as a biasing method (only biasing potentials), or if, as in replica exchange simulations, the temperature is the only biasing method ($w_i(x) = 0$). To handle such cases, one would only need to set the appropriate biasing matrix according to eq 2.

If we assume that each count in each histogram is independent, then the likelihood of observing the i th histogram is given by the multinomial distribution

$$P(n_{i,1}, n_{i,2}, \dots, n_{i,m} | p_{i,1}, p_{i,2}, \dots, p_{i,m}) = \frac{(\sum_l n_{i,l})!}{\prod_l n_{i,l}!} \prod_{l=1}^m (p_{i,l})^{n_{i,l}} \propto \prod_{l=1}^m (f_i c_{i,l} p_l^{\circ})^{n_{i,l}} \quad (5)$$

For real-life simulation methods, however, the samples from which the histograms are constructed may be far from being

independent. Since the effect of correlations is to reduce the effective sample size, the count $n_{i,l}$ is reduced by a factor $g_{i,l} = 1 + 2\tau_{i,l}$, where $\tau_{i,l}$ is the correlation time of the simulation in the l th bin of the i th simulation.¹³ While the $g_{i,l}$ factors may not have a substantial impact on the maximum likelihood estimates (eqs 7 and 8 below),¹³ they will have a significant effect on the estimated uncertainties. For notational simplicity in the subsequent derivations, it will be understood that the $n_{i,l}$ have been scaled by an appropriate $g_{i,l}$ so as to account for correlation.

If the histograms are independent of each other, then the overall likelihood for all S histograms is the product of multinomials

$$P(n_{1,1}, \dots, n_{1,m}; \dots; n_{S,1}, \dots, n_{S,m} | p_1^\circ, p_2^\circ, \dots, p_m^\circ) \propto \prod_{i=1}^S \prod_{l=1}^m (f_i c_{i,l} p_l^\circ)^{n_{i,l}} \quad (6)$$

where the likelihood is conditional only on the unbiased probabilities p_l° , since we can assume that the bias factors $c_{i,l}$ are known parameters, and the normalization constants f_i are known conditional on p_l° . The maximum likelihood estimates of the unbiased probabilities can be found by maximizing eq 6 with respect to $p_1^\circ, \dots, p_m^\circ$ and are given by solutions of the simultaneous nonlinear equations

$$p_l^\circ = \frac{\sum_{i=1}^S n_{i,l}}{\sum_{i=1}^S N f_i c_{i,l}} \quad (\text{for each } l) \quad (7)$$

and

$$f_i^{-1} = \sum_{l=1}^m c_{i,l} p_l^\circ \quad (8)$$

where N_i is the total number of counts in the i th histogram.¹⁸ Equations 7 and 8 are commonly known as the WHAM equations and are usually solved iteratively by choosing arbitrary values of f_i (e.g., $f_1 = f_2 = \dots = f_S = 1$) and calculating $p_1^\circ, \dots, p_m^\circ$ using eq 2. These are then used to calculate new values of f_i , and the procedure is repeated until a fixed point is reached.

The resulting unbiased probabilities can be used to calculate the potential of mean force with respect to the reduced coordinate x or the energy distribution conditional on a given value of x . For the WHAM analysis of replica exchange data for the purpose of calculating a potential of mean force along reduced coordinates x, y , etc., it is sufficient to collect the joint histogram, $n(E_i, x_j, y_k, \dots)$, of the potential energy E and each of the reduced coordinates. Because in the standard implementation of the replica exchange simulations there are no biasing potentials, the biasing matrix (see eq 4) in this case is defined as

$$c_{i,(j,k,\dots)} = \exp[-(\beta_i - \beta_0)E_i] \quad (9)$$

which is independent of the reduced coordinates. The resulting WHAM solution $p^\circ(E_i, x_j, y_k, \dots)$ is then summed over the energy variable to obtain the population in each reduced coordinate bin and the corresponding potential of mean force

$$\text{PMF}(x_j, y_k, \dots) = -(1/\beta_0) \log p^\circ(x_j, y_k, \dots) = -(1/\beta_0) \log \sum_i p^\circ(E_i, x_j, y_k, \dots) \quad (10)$$

We now turn to the question of estimating the statistical uncertainty of the WHAM solution. The source of the statistical noise is the histogram counts. Due to chance, the measured histogram counts will vary if the simulation is repeated from a different initial condition and so would the resulting WHAM estimate of the unbiased probability. We seek an estimate of the variance of the WHAM solution given the known statistics of the histogram counts. Ferrenberg and Swendsen¹² proposed an approximate expression for the uncertainty of the WHAM solution using standard error propagation techniques, which may not be always valid for estimating the statistical uncertainty of the PMF from replica exchange simulation data. We will now derive a novel estimate of the statistical uncertainty for WHAM based on Bayesian statistics and show in Appendix A how the Ferrenberg and Swendsen estimate can be derived from it.

From a Bayesian viewpoint, all of the information that is available about the unknown parameters of a model given the observed data is contained in the posterior probability density function $P(\text{parameters}|\text{data})$, in our case $P(\mathbf{p}^\circ|\mathbf{n})$, where $\mathbf{p}^\circ = (p_1^\circ, \dots, p_m^\circ)$ and $\mathbf{n} = (n_{1,1}, \dots, n_{1,m}; n_{2,1}, \dots, n_{2,m}; \dots; n_{S,1}, \dots, n_{S,m})$. The posterior probability can be related to the likelihood function using Bayes' theorem

$$P(\mathbf{p}^\circ|\mathbf{n}) = \frac{P(\mathbf{n}|\mathbf{p}^\circ)P(\mathbf{p}^\circ)}{P(\mathbf{n})} \quad (11)$$

where $P(\mathbf{p}^\circ)$ is the prior probability of the unknown parameters (reflecting any a priori knowledge the analyst may have) and $P(\mathbf{n})$ is a normalization constant.⁴¹ We take the prior probabilities to be uniform, thereby making the posterior probability proportional to the likelihood function given in eq 6. Given the posterior probability density function, we can estimate our confidence in the estimates of each of the unbiased probabilities p_l° using a measure of how "compact" or "spread out" the posterior probability is with respect to that parameter.

One such measure is the variance $\delta^2 p_l^\circ = \langle p_l^{\circ 2} \rangle - \langle p_l^\circ \rangle^2$ of p_l° calculated from the posterior probability distribution. However, the resulting integrals cannot be solved exactly in closed form. In Appendix A, we describe an approximate solution based on Laplace's integration method⁴² that leads to a closed form expression for the variance. This expression, although generally valid, in practice is not feasible when the number of bins m is large (as in the case of the PMF of the G-peptide) as it involves diagonalizing a $m \times m$ matrix. Sampling approaches, such as those based on the Metropolis algorithm,⁴³ offer an attractive alternative for the applications of WHAM described in this work and are well-suited for characterizing Bayesian posterior density functions.⁴⁴ Although one could generate samples from the posterior density using the Metropolis algorithm directly, a more efficient strategy can be devised by writing the posterior density from eqs 6 and 11 as

$$P(\mathbf{p}^\circ|\mathbf{n}) \propto \left(\prod_{i=1}^S c_i^{N_i} \right) \left(\prod_{l=1}^m p_l^{\circ M_l} \right) = Q(\mathbf{p}^\circ) \left(\prod_{l=1}^m p_l^{\circ M_l} \right) \quad (12)$$

where

$$Q(\mathbf{p}^\circ) = \prod_{i=1}^S \left(\sum_{l=1}^m c_{i,l} p_l^\circ \right)^{-N_i} \quad (13)$$

and M_l is the total number of counts in bin l from all histograms 1 ... S . Furthermore, we can impose the constraint that $\sum_l p_l^\circ = 1$ by setting

$$p_m^\circ = 1 - \sum_{l=1}^{m-1} p_l^\circ \quad (14)$$

and shortening the length of the \mathbf{p}° vector by one (i.e., $\mathbf{p}^\circ = (p_1^\circ, \dots, p_{m-1}^\circ)$). This allows us to rewrite eq 12 as

$$P(\mathbf{p}^\circ | \mathbf{n}) \propto Q(\mathbf{p}^\circ) \left(1 - \sum_{l=1}^{m-1} p_l^\circ\right)^{M_m} \prod_{l=1}^{m-1} p_l^{\circ M_l} \quad (15)$$

In Appendix B, we present an algorithm to efficiently sample the posterior distribution $P(\mathbf{p}^\circ | \mathbf{n})$ from eq 15. The result of this sampling procedure is a series of independent sets of unbiased probabilities p_l° , $l = 1, \dots, m$. The samples can then be used to estimate the statistical uncertainty of the unbiased probability of each bin simply by calculating the standard deviation of the samples corresponding to that bin. To estimate the statistical uncertainty of the potential of mean force PMF(x_j, y_k, \dots), we calculate the potential of mean force from each sample according to eq 10, and we then evaluate the standard deviation of the distribution of potential of mean force values at each reduced coordinate bin. The uncertainties reported in this work are twice the standard deviation.

One issue that arises in the error analysis of the WHAM method as outlined above is the inclusion of bins that have zero counts in all of the simulations. While the maximum likelihood estimate of the unbiased probability for any bin with zero counts is zero, we do not know this with certainty. For example, if we flip a coin 10 times and see no heads, we cannot say with certainty that the probability of heads is 0, since the 11th flip might result in heads. Instead, we would estimate the probability to be less than 0.1. This argument suggests that we should allow the estimates of the unbiased probabilities for zero-count bins to fluctuate in the error analysis sampling. However, if one were to arbitrarily add many zero-count bins, then each such bin will have equal uncertainty. This is because in the multinomial model underlying WHAM each bin is treated equally, and no provision is made for the presence of tails (i.e., that zero-count bins close to non-zero-count bins should be more likely to have a non-zero unbiased probability than zero-count bins far away from any non-zero-count bin). We have found that in practice it is better to force the unbiased probabilities corresponding to zero-count bins to be zero, particularly if there is a large number of such bins (such as when energy and reduced coordinate are strongly correlated).

3. Application of T-WHAM to a Test System

In this section, we test the temperature WHAM method described in the previous section by obtaining the potential of mean force for a system that can be solved analytically. The system has been chosen to highlight the power of the T-WHAM method and some of the issues that will be encountered in the application of the T-WHAM methodology to the G-peptide described in the next section.

The test system is depicted in Figure 1. The degrees of freedom are x , $0 \leq x \leq 1$, and y , $0 \leq y \leq 1$. The potential energy $U(x, y)$ has the following expression

$$U(x, y) = \begin{cases} +\infty & y > x^n \\ a(x + y) & y \leq x^n \end{cases} \quad (16)$$

where $a > 0$ and $n > 1$. We will be concerned with the probability distribution of the x coordinate obtained by integrating over the y coordinate. Small values of x , although energetically favored, are entropically disfavored due to the exponen-

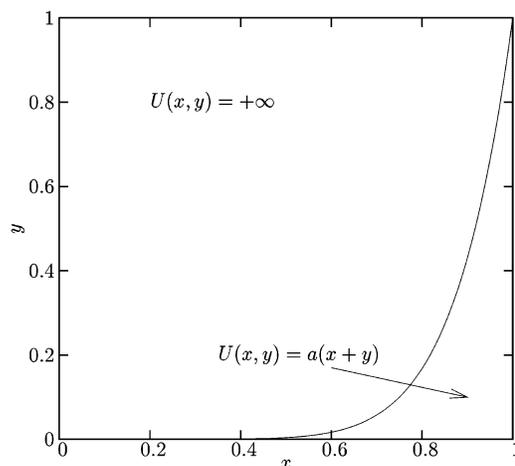


Figure 1. Schematic representation of the test system. The curve has equation $y = x^n$, with $n > 1$. The region of the xy plane above the x^n curve is disallowed. The potential energy below the curve is $U(x, y) = a(x + y)$, $a > 0$.

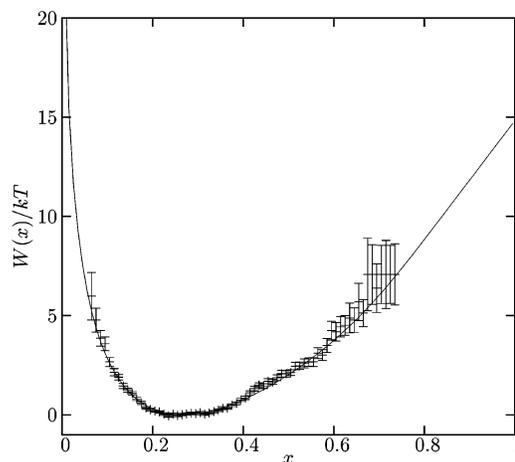


Figure 2. Exact (full line) and estimated (error bars) potential of mean force along the x coordinate of the system from Figure 1 with $n = 8$ and $a = 30$ using 28 000 samples from a single MC simulation at $\beta = 1$.

tially decreasing range of allowed values of y as x approaches 0. Large values of x , on the other hand, are energetically disfavored and entropically favored. The probability distribution $p(x)$ of the x coordinate in the canonical ensemble is therefore expected to have a maximum at an intermediate value of x where the energetic and entropic driving forces are optimally balanced. These features are confirmed by the following expression of $p(x)$ at the inverse temperature $\beta = 1/kT$

$$p(x) = \frac{1}{Q} \int_0^{x^n} dy e^{-\beta a(x+y)} = \frac{e^{-\beta a x}}{\beta a Q} (1 - e^{-\beta a x^n}) \quad (17)$$

where Q is the canonical partition function

$$Q = \int_0^1 dx \int_0^{x^n} dy e^{-\beta a(x+y)} \quad (18)$$

The potential of mean force $W(x) = -(1/\beta) \log p(x)$ for $\beta = 1$, $a = 30$, and $n = 8$ is shown in Figure 2. It presents a minimum at $x = x_{\min} \approx 0.267$ and monotonically increases away from x_{\min} . The potential of mean force tends to infinity for $x \rightarrow 0$. At $x = 1$, the potential of mean force is nearly $15kT$ relative to $W(x_{\min})$.

We seek to reproduce the potential of mean force $W(x)$ for $\beta = 1$, $a = 30$, and $n = 8$ from histograms of x obtained by

sampling the x and y coordinates by Monte Carlo according to the given potential $U(x,y)$. A new proposed position is obtained by generating uniform random displacements $-\Delta < dx < \Delta$ and $-x^n\Delta < dy < x^n\Delta$, respectively. The maximum displacement parameter Δ is adjusted so as to achieve a 50–30% acceptance ratio. Notice that, for optimal sampling efficiency, the largest y step size is x dependent and, consequently, the underlying Monte Carlo (MC) transition probability is not symmetric. Because of this, the Monte Carlo acceptance criterion includes the extra factor $\theta(|y' - y|/x^n\Delta)(x/x')^n$, where (x',y') is the proposed new position and $\theta(\cdot)$ is a function that is equal to 1 if the argument is less than 1 and 0 otherwise. The Monte Carlo walk is started near $(x,y) = (0,0)$. After an equilibration period, histograms of x are collected in 100 bins between $x = 0$ and $x = 1$. The sampling frequency of x is set small enough so to make the correlation between samples negligibly small.

In Figure 2, we show the estimate of the potential of mean force from a single MC simulation at $\beta = 1$ with 28 000 samples. The probability $p(x)$ of the bin centered on x is calculated by dividing the number of counts in that bin by the total number of counts. This can be shown to be the WHAM solution in the case of a single simulation. Error bars are twice the standard deviation of $W(x) = -(1/\beta) \log p(x)$ as obtained from 128 independent samples of $p(x)$ generated as described in the previous section. (In the case of a single simulation, this involves the sampling of a pure Dirichlet distribution and therefore does not require an acceptance/rejection step.)

As Figure 2 shows, the data from the MC simulation allows an accurate estimate of the potential of mean force only up to about $5kT$ from the minimum. Values of x corresponding to the potential of mean force larger than $5kT$ correspond to relative probabilities of less than $1/100$ and are therefore not observed often enough during the simulation to obtain a reliable estimate of the potential of mean force. The potential of mean force cannot even be estimated for values of x less than 0.06 and larger than 0.75 because these values of x are never observed during the 28 000 sample simulation.

We will now obtain the potential of mean force $W(x)$ at $\beta = 1$ by the T-WHAM method. The region beyond $x = 0.6$, which is not well covered by the MC simulation at $\beta = 1$, corresponds to high potential energy and is expected to be better sampled by simulations at high temperatures. On the other hand, the region below $x = 0.2$, which is energetically favored but entropically disfavored, is better sampled at low temperatures. The T-WHAM scheme should therefore include data from simulations at higher as well as lower temperatures than those of the reference. As the potential of mean force at $x = 1$ is $15kT$, it might be concluded that it would be necessary to have a simulation with a temperature 15 times larger than that of the reference in order to sufficiently sample $x \approx 1$. The $x = 1$ state, however, becomes populated at significantly lower temperatures than those expected by this simple analysis because the potential of mean force itself depends on temperature. The minimum of the potential of mean force shifts to higher values of x with increasing temperature. This is a consequence of the fact that larger values of x are favored by entropy. At the highest temperature, the minimum of the potential of mean force is close to $x = 1$, and the $x = 1$ position is visited with high probability.

We performed seven MC simulations at inverse temperatures $\beta = 0.2, 0.4, 0.7, 1.0, 1.5, 2.0$, and 4.0 . Four thousand samples are obtained from each simulation so that the same total number of samples as for the previous single temperature test is collected. The potential of mean force at $\beta = 1$ was then obtained using the T-WHAM method as described above by

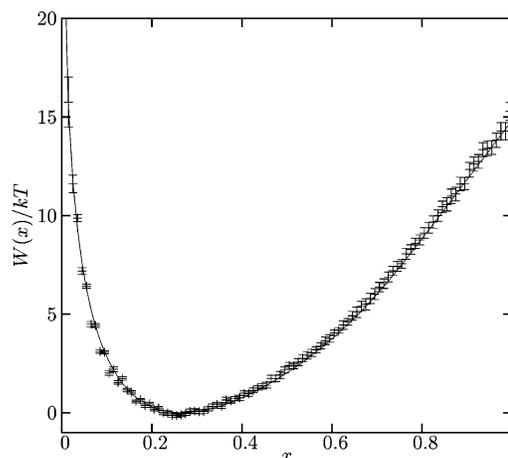


Figure 3. Exact potential of mean force along the x coordinate (full line) and corresponding T-WHAM estimate (error bars) for the system from Figure 1 with $n = 8$ and $a = 30$ using histograms from seven MC simulations of 4000 samples each at $\beta = 0.2, 0.4, 0.7, 1.0, 1.5, 2.0$, and 4.0 .

binning x and the potential energy in 100 bins each. The resulting joint probability of position and energy $p^\circ(x,E)$ is then integrated over the energy coordinate to obtain the marginal distribution of x . The results for the potential of mean force are shown in Figure 3. The potential of mean force estimated by the T-WHAM analysis reproduces well with small statistical uncertainties the exact potential of mean force along the x coordinate. In contrast, the data from the simulation at a single temperature, although seven times longer, allows estimation of only the portion of the potential of mean force corresponding to the range of x sufficiently sampled at that temperature.

This application of the T-WHAM method on this simple system highlights some of the powerful features of the method. Although uniform sampling coverage of phase space cannot be achieved at any single temperature, the T-WHAM method allows us to draw information from simulations at multiple temperatures that overall have a good coverage of phase space. The inclusion of high temperature data makes it possible to sample high potential energy regions. These high potential energy regions are often entropically favored and therefore can become populated at temperatures significantly smaller than would be expected on the basis of the potential energy alone. On the other hand, the inclusion of low temperature data makes it possible to resolve small, entropically disfavored, energy wells which are not observed at higher temperatures.

4. Application of T-WHAM to the G-Peptide

The applicability of the temperature WHAM procedure to a molecular system is demonstrated using the C-terminal β -hairpin of the B1 domain of protein G (the G-peptide) as an example. This peptide has the following sequence: 41-GEWYDDAT-KTFTVTE-56. We studied the G-peptide using replica exchange molecular dynamics (REMD)⁸ with the OPLS all-atom force field and the analytical generalized Born implicit solvent model (OPLS-AA/AGBNP).¹¹ During the REMD calculation, replicas of the system are simulated at 20 different temperatures while at given intervals the temperatures are exchanged between replicas based on a modified Metropolis criterion.⁸ In our previous paper, we reported the potential of mean force (PMF) at a single temperature based on the results of the replica exchange simulation at that temperature (i.e., based on conformations sampled by each replica at that temperature). Here, we demonstrate how the other temperatures can help resolve

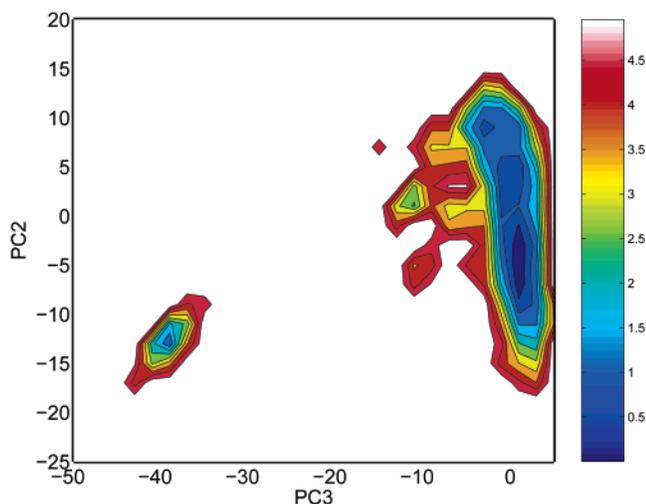


Figure 4. The potential of mean force of the C-terminal peptide from protein G with respect to the third (PC3) and second (PC2) principal components of the interatomic distances. The PMF is generated using the data from one temperature of the REMD simulation. The energy is in kcal/mol. The temperature is 298 K.

features of high free energy not observable when considering only those conformations sampled at a single, low temperature.

Originally,¹¹ we plotted the PMF with respect to the radius of gyration (R_g) of the hydrophobic residues (W43, Y45, F52, and V54) and the number of native β -sheet hydrogen bonds. We chose these particular order parameters in order to compare our results with benchmark explicit solvent simulations.^{9,10} The PMF is divided into two regions corresponding to a major population (88%) of β -hairpin structures (with $R_g < 7 \text{ \AA}$) and a minor population (8%) of α -helices (centered at R_g of 9 \AA). These two regions are separated by a free-energy barrier in the PMF. We want to determine the path(s) by which the α -helix conformation converts to the β -hairpin conformation. Resolving the free energy barrier in the PMF along the R_g and hydrogen bond order parameters is not a good choice for the analysis of possible paths. Using principal component analysis of 42 interatomic distances between nonadjacent alpha carbons (C_α), we determined two principal components (PC) that separate the α and β regions; they correspond to the second and third principal components. The PMF of the peptide at 298 K is plotted in Figure 4 with respect to the second and third PCs. The second principal component (PC2) is dominated by the distance between the N-terminus and the middle of the peptide, and the third PC (PC3) is dominated by the distance between the termini. The region of α -helical conformations is centered around $PC3 = -39$ and $PC2 = -13$. The β -hairpin region lies between -5 and 5 along PC3.

A region of high free energy also separates the α and β regions in Figure 4; however, the separation is greater in this case than it is using the R_g hydrogen bond order parameters. Details in this high free energy region cannot be resolved using the data from a single temperature of the REMD simulation. The largest possible relative free energy that can be measured from N samples collected at temperature T is $-kT \ln(1/N)$ or approximately 6 kcal/mol for our simulation ($N = 40\,000$) at room temperature. The free energy barriers separating the α and β regions are apparently greater than 6 kcal/mol from the minimum.

Information about the higher free-energy regions at 298 K can be gleaned from the other temperatures of the REMD simulation by applying T-WHAM. The 20 temperatures used during the simulation ranged from 270 to 690 K. The higher

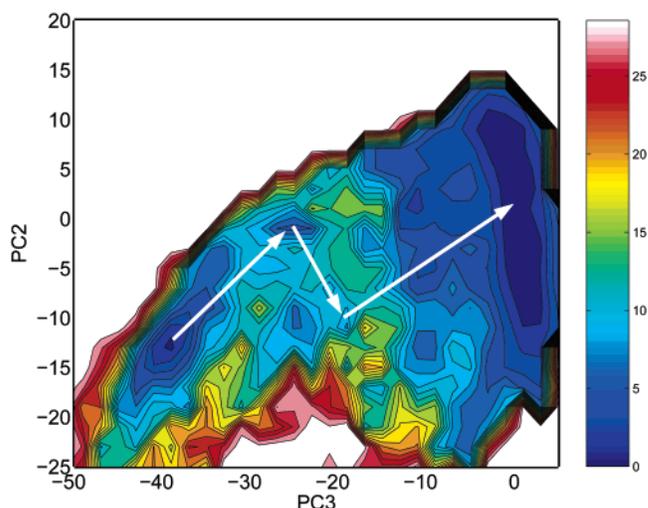


Figure 5. The potential of mean force of the C-terminal peptide from protein G with respect to the third (PC3) and second (PC2) principal components of the interatomic distances. The PMF is generated with the data from all temperatures of the REMD simulation using T-WHAM. The energy is in kcal/mol. The temperature is 298 K. Only the regions with free energies less than 30 kcal/mol are displayed.

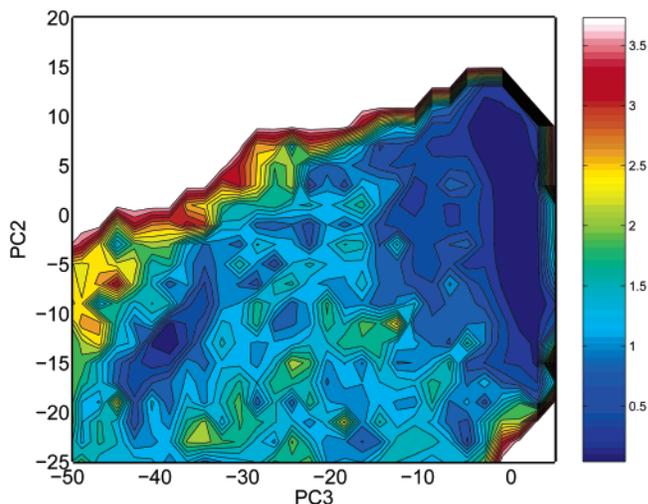


Figure 6. The estimate of the error in the potential of mean force in Figure 5.

free energy regions were resolved by applying T-WHAM using all 20 temperatures to generate a PMF at 298 K, which is shown in Figure 5. Relatively low free-energy pathways connecting the α and β regions are now apparent. The lowest free energy pathway is indicated by the white arrows in Figure 5. The maximum of this free-energy path lies around 13 kcal/mol from the minimum in the PMF.

We performed error analysis of the PMF obtained from the T-WHAM procedure. The estimate of the error in the PMF is shown in Figure 6. The intermediate regions between α and β have small errors well within 2 kcal/mol. (Only along the edges of the PMF are the errors greater than 3 kcal/mol.) The maximum error along the path indicated by the white arrows in Figure 5 is around 1.75 kcal/mol; therefore, the maximum of this path is at most about 15 kcal/mol from the minima. This low error indicates that the free energy changes calculated along this pathway are statistically significant.

We can also use T-WHAM to obtain a better estimate of the free energy difference between the β -hairpin and α -helical conformations. Through use of the data from a single temperature, the calculated free energy difference is -1.567 ± 0.034

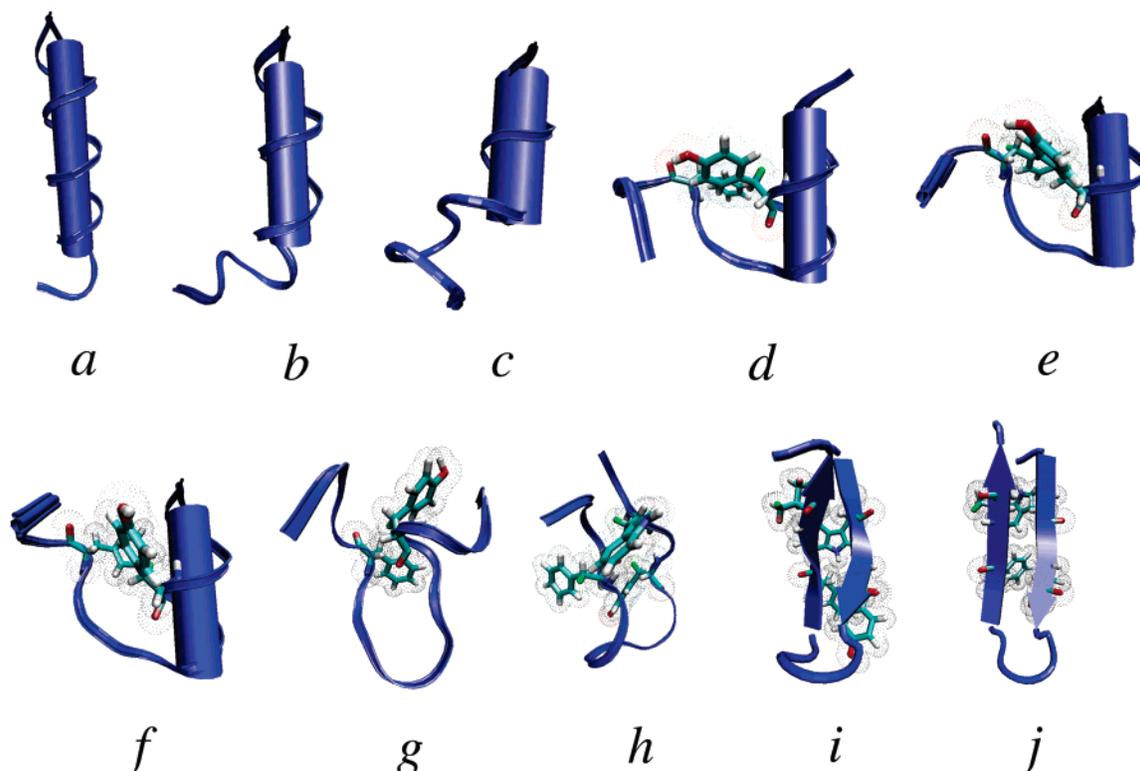


Figure 7. The structures that lie on the minimum free-energy pathway between α and β regions in the PMF in Figure 5. The ribbons follow the trace of the backbone, cylinders indicate the location of α -helices, and arrows indicate β -sheet as determined by STRIDE.⁵⁷ Hydrophobic residues that are in contact with other hydrophobic residues are also shown. The N-terminus of each structure is pointing up and is positioned on the right side in each frame. Structures are drawn using VMD.⁵⁸

kcal/mol in favor of the β -hairpin. Through use of the data from all 20 temperatures, the free energy difference is determined to be -1.573 ± 0.018 kcal/mol. While these two free energies are within the calculated error from each other, the estimation of the error is about a factor of 2 smaller using data from all of the REMD replicas.

To gain some insights into a possible mechanism of the transition between the α -helix and the β -hairpin conformations of the G-peptide, we extract from the replica exchange ensemble representative samples corresponding to the conformations along the path shown in Figure 5. The path was divided into a series of small rectangular regions. The C_{α} root-mean-square deviations (rmsd's) between structures in adjacent regions were calculated. A path was created by finding the sequence of structures that minimizes the cumulative rmsd among the pairs of structures. These conformations are shown in Figure 7 starting from the α -helix conformation (a) and ending in the β -hairpin (j) conformation. This series of structures shows that as the α -helix converts to the β -hairpin the N-terminal segment of the helix remains folded until the C-terminus is oriented nearly parallel to the helix. The N-terminal portion of the helix in structures a through f is stabilized by a charge–polar interaction between the N-terminal amino group on G41 and the hydroxyl group on T44. Later, the remaining segment of the helix unravels as seen with structures f and g. The two strands of the β -hairpin aligned, allowing the β -sheet hydrogen bonds to form in structure i resulting in the final β -hairpin (j). This mechanism demonstrates that the α -helix does not have to completely unfold for the β -hairpin structure to begin to form.

We also tracked the formation of the hydrophobic core during the transformation from the α -helix to the β -hairpin. Part of the hydrophobic core begins to form with contact being made between Y45 and F52 in structure d. Interestingly, the hydrophobic core begins to associate while the conformation is still

significantly α -helical, before any native β -hairpin hydrogen bonds are formed. The phenylalanine and tyrosine remain in contact in all of the conformations from d to j in Figure 5, except for structure g where Y45 and F52 are close but not in contact. The early formation of the hydrophobic core prior to the formation of β -hairpin hydrogen bonds is in agreement with previous theoretical treatments.^{3,25} The mechanism proposed here differs, however, in that the α -helical state is an intermediate in the folding of the hairpin of the G-peptide.

The existence of a minimum free energy path with respect to these reduced coordinates (PC2 and PC3 in Figure 5) does not guarantee that this path describes the dynamic ones actually followed by the system to make the transition from the α -helix to the β -hairpin.⁴⁵ A necessary condition is that the minimum free energy path along the potential of mean force calculated for the reduced coordinates in fact describes a continuous path in conformational space, as does the path shown in Figure 7. Even so, it is possible that there are alternative paths described by other order parameters with lower free energy barriers. The transition path sampling techniques developed by Chandler and co-workers provide a powerful but computationally demanding approach to search for alternative reactive paths.² Another approach is to use discrete path sampling techniques^{46–48} to calculate folding pathways as Wales has done for the G-peptide.⁴⁹ In a forthcoming paper,³⁴ we explore possible folding paths for the G-peptide including the α to β transition by performing a kinetic Monte Carlo walk on the G-peptide replica exchange ensemble of states. In addition to the α -helix to β -hairpin transition path shown in Figure 7, we have found another path where both ends of the helix unwind, leaving approximately one turn of helix in the middle of the molecule. This then serves as a nucleation point for formation of the β -hairpin turn, which is stabilized by hydrophobic interactions between Y45 and F52. By analyzing an ensemble of transition

paths generated using discrete path sampling on the replica exchange state space, we hope to identify alternative order parameters for best characterizing this transition and for construction of the corresponding potentials of mean force using the T-WHAM procedure.

5. Simulation Details

The MD replica exchange canonical sampling method has been implemented in the molecular simulation package IMPACT⁵⁰ following the approach proposed by Sugita and Okamoto.⁸ In this method, a series of structures (the replicas) are simulated in parallel using MD at different temperatures. The temperatures of two replicas are exchanged periodically with a modified Metropolis transition probability.⁸ After the exchange, the velocities of the replicas are resampled at the new given temperatures. In our simulations, 20 replicas are run in parallel at temperatures between 270 and 690 K. Each replica is thermalized at its respective temperature for 100 ps with a time step of 1 fs. MD replica exchange sampling is carried out for a total of 10 ns using a time step of 1 fs. Transitions between adjacent temperatures are attempted every 250 MD steps. Immediately prior to attempting temperature exchanges, the configuration of each replica is stored for later analysis. The transition acceptance ratio we obtain is on average 42%.

MD simulations are carried out using the OPLS-AA force field⁵¹ and the AGBNP implicit solvent model, which is based on a novel pairwise descreening implementation⁵² of the generalized Born model⁵³ and a recently proposed nonpolar hydration free energy estimator. AGBNP employs a parameter-free and conformation-dependent analytical scheme to obtain the pairwise descreening scaling coefficients used in the computation of the Born radii. The agreement between the AGBNP Born radii and exact numerical calculations was found to be excellent. The AGBNP nonpolar model consists of an estimator for the solute–solvent van der Waals interaction energy in addition to an analytical surface area component corresponding to the work of cavity formation.⁵² Because AGBNP is fully analytical with first derivatives, it is well suited for MD sampling. A detailed description of the AGBNP model and its implementation is provided in ref 52. For this study, we have also used a modified generalized Born pair function designed so that the dielectric screening between particular atom pairs can be adjusted as described in ref 11.

6. Conclusions

We have shown that the temperature WHAM (T-WHAM) method can be successfully applied to the calculation of potentials of mean force by combining data from all the replicas of a replica exchange simulation rather than, as is more common, using data only from the temperature of interest. By use of data from the replicas at all of the temperatures, a more accurate representation of the potential of mean force is obtained. We have developed a novel Bayesian scheme to estimate the statistical uncertainty of any quantity derived from WHAM, such as the potential of mean force. This scheme follows from the identification of the WHAM solution with the maximum of the posterior density (which is related to a Dirichlet distribution) and by evaluating alternative solutions by sampling this distribution around the maximum.

We have applied the T-WHAM scheme to obtain the potential of mean force of a test system that can be solved analytically. This system was selected because it presents some of the same challenges found in real chemical systems, which involves a free energy tradeoff between minimizing the energy and

maximizing the entropy. At the reference temperature, the system visits only the relatively small region around the potential of mean force minimum during the simulation. The data from a simulation only at the reference temperature is not sufficient to provide an estimate of the potential of mean force in the high free energy regions. We showed that, by using T-WHAM to combine data from simulations at multiple temperatures, it is possible to calculate the potential of mean force with accuracy and precision along all of the reaction coordinates at the same computational cost as the single temperature simulation. Through use of T-WHAM, the ability to obtain the potential of mean force in a particular region depends on whether that region is properly sampled at any one of the temperatures included in the simulations, not only at the reference temperature. We showed that high temperature simulations are needed to sample regions with high potential energy whereas simulations at lower temperatures than that of the reference are needed to properly sample low entropy regions. We also showed that, because the potential of mean force changes with temperature (favoring low potential energy regions at low temperatures and high entropy regions at high temperatures), regions scarcely populated at the reference temperature become populated at temperatures much closer to the reference temperature than those predicted from the potential of mean force at the reference temperature. This, being a general feature of any chemical system, explains the success of replica exchange techniques in efficiently sampling a wide range of high energy conformations even when employing a relatively modest range of temperatures. We apply the T-WHAM scheme to resolve the potential of mean force along two principal components of the G-peptide from replica exchange molecular dynamics data using the OPLS-AA/AGBNP implicit solvent force field. At room temperature, only two conformational states are significantly populated. The dominant one, with a population fraction around 90%, corresponds to β -hairpin conformations and the minor one to α -helical conformations. At room temperature, the interconversion rate between the β -hairpin state and the α -helix state is too slow to observe transition events during a standard molecular dynamics simulation at room temperature started at either end. In the replica exchange simulation using replicas at temperatures in the range between 270 and 690 K, the two states can interconvert. This enables us to obtain a potential of mean force that spans both states and makes it possible to estimate their relative free energy. In the potential of mean force obtained only from the room-temperature replica, the two stable states are well-resolved, but the transition states that connect them are not resolved. This is because the room-temperature replica does not have enough thermal energy to sample the transition region during the simulation time. In the replica exchange ensemble, the conversion between the β -hairpin state and the α -helix state occurs in high-temperature replicas.

By using data from all of the replicas, we were able to resolve the potential of mean force of a much larger area of conformational space. We were able to determine the potential of mean force in the transition region and identify possible transition states for the conversion of the minor population of α -helical conformations into the major β -hairpin population. One of these transition states has a free energy 13 kcal/mol higher than the minimum of the PMF. By following a connected minimum free energy path from the α -helix region through the transition state and to the β -hairpin basin, we were able to postulate a mechanism of interconversion between the α -helix and β -hairpin conformations. This mechanism reveals how this peptide can undergo a conformational change from an α -helical state to a

β -hairpin state without unfolding, while retaining a significant amount of secondary structure throughout the conversion. A mechanism similar to the one proposed here may be involved in the formation of β -sheet secondary structures in proteins. α -Helical intermediates along the folding pathway of β -sheets have been observed previously^{35–40} in the folding of β -lactoglobulin.

Acknowledgment. This work was supported by National Institutes of Health Grant GM30580. We thank Tateki Ishida for helpful discussions on the WHAM method. This paper has been prepared for the Festschrift in honor of Professor David Chandler's 60th birthday. He has pioneered the development of methods to study rare events in condensed phase chemical and biological systems, and we are very pleased to be able to honor him on this occasion.

Appendix A. Comparison with the Ferrenberg and Swendsen Uncertainty Estimation Formula

In this appendix, we show that our error analysis reduces to the error propagation expression of Ferrenberg and Swendsen (eq 5 in ref 12) under certain approximations.

Let us consider any one particular bin (call it bin number l). The uncertainty in the unbiased probability associated with that bin can be quantified by the variance, $\delta^2 p_l^\circ = \langle p_l^\circ \rangle - \langle p_l^\circ \rangle^2$ of p_l° , calculated from the posterior distribution (eq 15). We have

$$\langle p_l^\circ \rangle = \int \dots \int dp_1^\circ \dots dp_{m-1}^\circ P(\mathbf{p}^\circ | \mathbf{n}) \quad (19)$$

where the limits of integration are given by the condition

$$\sum_{j=1}^{m-1} p_j^\circ \leq 1 \quad (20)$$

and similarly for $\langle p_l^\circ \rangle$. The integral in eq 19 cannot be expressed in closed form. An approximate expression for the variance can be obtained by using Laplace's method. The logarithm of the posterior distribution (eq 15)

$$\ln P(\mathbf{p}^\circ | \mathbf{n}) = C + \sum_s N_s \ln f_s + \sum_l M_l \ln p_l \quad (21)$$

where C is a constant, is expanded up to the second order around the maximum (the WHAM solution) yielding

$$\ln P(\mathbf{p}^\circ | \mathbf{n}) \approx C' + \frac{1}{2} \delta \mathbf{p}^\circ \dagger H \delta \mathbf{p}^\circ \quad (22)$$

where C' is another constant and H is the Hessian matrix, calculated at the maximum, of the logarithm of the posterior whose elements are, from eq 21,

$$H_{ll'} = \frac{\partial^2 \ln P(\mathbf{p}^\circ | \mathbf{n})}{\partial p_l^\circ \partial p_{l'}^\circ} = \sum_s N_s c_{sl} \frac{\partial f_s}{\partial p_l^\circ} - \delta_{ll'} \frac{M_l}{p_l^{\circ 2}} \quad (23)$$

and $\delta \mathbf{p}^\circ$ is the displacement from the WHAM solution. By use of eq 22, the right-hand side of eq 19 reduces to a Gaussian integral that can be solved analytically. This is best done in a coordinate system in which the Hessian is diagonal. If u_{kl} is the l th component of the k th eigenvector and $-1/\sigma_k^2$ the corresponding eigenvalue, then it can be shown that by using the Laplace approximation the variance of p_l° is given by

$$\delta^2 p_l^\circ = \sum_k u_{kl} \sigma_k^2 \quad (24)$$

The Ferrenberg and Swendsen expression for the uncertainties

is recovered by making the further approximation that $\partial f_s / \partial p_l^\circ = 0$ in eq 23. In this approximation, the Hessian is diagonal with diagonal elements $-M_l / p_l^{\circ 2}$, and eq 24 reduces to

$$\delta^2 p_l^\circ = p_l^{\circ 2} / M_l \quad (25)$$

and the relative uncertainty becomes

$$\frac{\delta p_l^\circ}{p_l^\circ} = \frac{1}{\sqrt{M_l}} \quad (26)$$

which is the Ferrenberg and Swendsen result.¹²

From this analysis, it is apparent that the Ferrenberg and Swendsen uncertainty formula is derived from our Bayesian uncertainty estimation protocol under two approximations: the Laplace approximation for calculating the moments of the posterior distribution and the assumption that the free energy factors f_s are constant. The latter approximation is potentially the most problematic. By doing so, the uncertainty in the probability p_l° corresponding to bin l depends only on the total number of counts M_l in that bin. However, the WHAM estimate of p_l° (eqs 7 and 8) depends not only on the histogram count in bin l but also on the counts in all of the other bins through the free energy factors f_s . A full uncertainty analysis should therefore include the statistical noise propagated through the f_s factors, which is neglected by the uncertainty analysis of Ferrenberg and Swendsen. Ferrenberg and Swendsen's estimate of the uncertainty is expected to be least reliable in cases where the free energy factors f_s are very sensitive to statistical noise in the histogram counts. This happens when the overlaps between histograms from different simulations are poor, and as a result, even a small change in the histogram counts in the overlap region could result in a large change in the WHAM free energy estimates. In these cases, Ferrenberg and Swendsen's formula (eq 26) would underestimate the uncertainty of the WHAM solution.

Appendix B. An Algorithm to Sample the Posterior Density

If we ignore for the moment the $Q(\mathbf{p}^\circ)$ factor, then eq 15 is equivalent to a $m - 1$ variate Dirichlet distribution⁵⁴ on the probabilities

$$P_{\text{Dirichlet}}(y_1, \dots, y_{m-1} | \theta_1, \dots, \theta_m) = \frac{\Gamma(\sum_{j=1}^m \theta_j)}{\prod_{j=1}^m \Gamma(\theta_j)} (1 - \sum_{j=1}^{m-1} y_j)^{\theta_m - 1} \prod_{j=1}^{m-1} y_j^{\theta_j - 1} \quad (27)$$

where $y_j \geq 0$ and $\sum_{j=1}^{m-1} y_j \leq 1$. The Dirichlet distribution has the property that all of its univariate conditional probabilities are scaled Beta distributions

$$P_{\text{Dirichlet}}(y_i | y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_{m-1}; \theta_1, \dots, \theta_m) = P_{\text{Beta}}(y_i | \theta_i, \theta_m, 1 - \sum_{j \neq i} y_j) \quad (28)$$

where

$$P_{\text{Beta}}(x | \alpha, \beta, \gamma) = \frac{\Gamma(\alpha + \beta) \gamma^{1-\beta}}{\Gamma(\alpha) \Gamma(\beta)} x^{\alpha-1} (\gamma - x)^{\beta-1} \quad (29)$$

$$0 \leq x \leq \gamma$$

To sample from a Dirichlet distribution (eq 27), one can simply choose a starting point (y_1, \dots, y_{m-1}) and update each y_i in turn according to a Metropolis procedure. If we choose the proposal for y_i according to its conditional distribution as given by eq 28, then the Metropolis acceptance criterion will always be satisfied, and the sampling will be rejectionless (a procedure known as “Gibbs sampling”⁵⁵). This procedure is computationally feasible, since pseudorandom numbers distributed according to the general Beta distribution $P_{\text{Beta}}(x|\alpha, \beta, \gamma)$ can be generated from standard Beta random variates $P_{\text{Beta}}(x|\alpha, \beta, 1)$ by multiplying by γ , while the latter can be generated using established methods.⁵⁶ We also recognize that in the derivation above the choice of bin m , as the dependent variable to ensure that $\sum_l p_l^\circ = 1$, is arbitrary. We therefore use the following sampling algorithm. Beginning with the maximum likelihood estimates of \mathbf{p}° , we update the probability of two randomly chosen bins l and k by generating a random number x distributed according to $P_{\text{Beta}}(x|M_l + 1, M_k + 1, 1)$ and set

$$p_l^{\circ\prime} = (1 - \sum_{j \neq l, k} p_j^\circ) x = (p_l^\circ + p_k^\circ) x \quad (30)$$

and

$$p_k^{\circ\prime} = (p_l^\circ + p_k^\circ) - p_l^{\circ\prime} \quad (31)$$

as a proposal for a new \mathbf{p}° , which we accept with probability

$$\min \left[1, \frac{Q(p_1^{\circ\prime}, \dots, p_l^{\circ\prime}, \dots, p_k^{\circ\prime}, \dots, p_m^{\circ\prime})}{Q(p_1^\circ, \dots, p_l^\circ, \dots, p_k^\circ, \dots, p_m^\circ)} \right] \quad (32)$$

This process is repeated to produce a walk in probability space. Samples are then collected at regular intervals, long enough to ensure statistical independence of consecutive samples.

References and Notes

- Elber, R.; Karplus, M. *Chem. Phys. Lett.* **1987**, *139*, 375–380.
- Bolhuis, P. G.; Chandler, D.; Dellago, C.; Geissler, P. L. *Annu. Rev. Phys. Chem.* **2002**, *53*, 291.
- Bolhuis, P. G. *Proc. Natl. Acad. Sci. U.S.A.* **2003**, *100*, 12129–12134.
- Eastman, P.; Grønbech-Jensen, N.; Doniach, S. *J. Chem. Phys.* **2001**, *114*, 3823–3841.
- Onuchic, J. N.; Luthey-Schulten, Z.; Wolynes, P. G. *Annu. Rev. Phys. Chem.* **1997**, *48*, 545–600.
- Kumar, S.; Ma, B.; Tsai, C.-J.; Sinha, N.; Nussinov, R. *Protein Sci.* **2000**, *9*, 10–19.
- Onuchic, J. N.; Wolynes, P. G. *Curr. Opin. Struct. Biol.* **2004**, *14*, 70–75.
- Sugita, Y.; Okamoto, Y. *Chem. Phys. Lett.* **1999**, *314*, 141–151.
- García, A. E.; Sanbonmatsu, K. Y. *Proteins: Struct. Funct. Genet.* **2001**, *42*, 345–354.
- Zhou, R.; Berne, B. J.; Germain, R. *Proc. Natl. Acad. Sci. U.S.A.* **2001**, *98*, 14931–14936.
- Felts, A. K.; Harano, Y.; Gallicchio, E.; Levy, R. M. *Proteins* **2004**, *56*, 310–321.
- Ferrenberg, A. M.; Swendsen, R. H. *Phys. Rev. Lett.* **1989**, *63*, 1195–1198.
- Kumar, S.; Bouzida, D.; Swendsen, R. H.; Kollman, P. A.; Rosenberg, J. M. *J. Comput. Chem.* **1992**, *13*, 1011–1021.
- Roux, B. *Comput. Phys. Comm.* **1995**, *91*, 275–282.
- Sugita, Y.; Kitao, A.; Okamoto, Y. *J. Chem. Phys.* **2000**, *113*, 6042–6050.
- Shea, J.-E.; Brooks, C. L., III. *Annu. Rev. Phys. Chem.* **2001**, *52*, 499–535.
- Ferrenberg, A. M.; Swendsen, R. H. *Comput. Phys.* **1989**, Sept/Oct, 101–104.
- Bartels, C.; Karplus, M. *J. Comput. Chem.* **1997**, *18*, 1450–1462.
- Honda, S.; Kobayashi, N.; Munekata, E. *J. Mol. Biol.* **2000**, *295*, 269–278.
- Blanco, F. J.; Rivas, G.; Serrano, L. *Nat. Struct. Biol.* **1994**, *1*, 584–590.
- Muñoz, V.; Henry, E. R.; Hofrichter, J.; Eaton, W. A. *Proc. Natl. Acad. Sci. U.S.A.* **1998**, *95*, 5872–5879.
- Kobayashi, N.; Honda, S.; Yoshii, H.; Munekata, E. *Biochemistry* **2000**, *39*, 6564–6571.
- Dinner, A. R.; Lazaridis, T.; Karplus, M. *Proc. Natl. Acad. Sci. U.S.A.* **1999**, *96*, 9068–9073.
- Kolinski, A.; Ilkowski, B.; Skolnick, J. *Biophys. J.* **1999**, *77*, 2942–2952.
- Pande, V.; Rokhsar, D. S. *Proc. Natl. Acad. Sci. U.S.A.* **1999**, *96*, 9062–9067.
- Roccatano, D.; Amadei, A.; Di Nola, A.; Berendsen, H. J. *Protein Sci.* **1999**, *8*, 2130–2143.
- Ma, B.; Nussinov, R. *J. Mol. Biol.* **2000**, *296*, 1091–1104.
- Zagrovic, B.; Sorin, E. J.; Pande, V. *J. Mol. Biol.* **2001**, *313*, 151–169.
- Zhou, R.; Berne, B. J. *Proc. Natl. Acad. Sci. U.S.A.* **2002**, *99*, 12777–12782.
- Zhou, Y.; Linhananta, A. *Proteins: Struct. Funct. Genet.* **2002**, *47*, 154–162.
- Lee, J.; Shin, S. *J. Phys. Chem. B* **2002**, *106*, 8796–8802.
- Tsai, J.; Levitt, M. *Biophys. Chem.* **2002**, *101–102*, 187–201.
- Muñoz, V.; Thompson, P. A.; Hofrichter, J.; Eaton, W. A. *Nature* **1997**, *390*, 196–199.
- Andrec, M.; Felts, A. K.; Gallicchio, E.; Levy, R. M. In preparation.
- Hamada, D.; Segawa, S.; Goto, Y. *Nat. Struct. Biol.* **1996**, *3*, 868–873.
- Kuwajima, K.; Yamaya, H.; Sugai, S. *J. Mol. Biol.* **1996**, *264*, 806–822.
- Hamada, D.; Goto, Y. *J. Mol. Biol.* **1997**, *269*, 479–487.
- Arai, M.; Ikura, T.; Semisotnov, G.; Kihara, H.; Amemiya, Y.; Kuwajima, K. *J. Mol. Biol.* **1998**, *275*, 149–162.
- Forge, V.; Hoshino, M.; Kuwata, K.; Arai, M.; Kuwajima, K.; Batt, C. A.; Goto, Y. *J. Mol. Biol.* **2000**, *296*, 1039–1051.
- Kuwata, K.; Shastry, R.; Cheng, H.; Hoshino, M.; Batt, C. A.; Goto, Y.; Roder, H. *Nat. Struct. Biol.* **2001**, *8*, 151–155.
- Sivia, D. S. *Data Analysis: A Bayesian Tutorial*; Oxford University Press: Oxford, 1996.
- De Bruijn, N. G. *Asymptotic Methods in Analysis*; Dover: New York, 1981.
- Metropolis, N.; Rosenbluth, A. W.; Rosenbluth, M. N.; Teller, A. H.; Teller, E. *J. Chem. Phys.* **1953**, *21*, 1087–1091.
- Hastings, W. K. *Biometrika* **1970**, *57*, 97–109.
- Dellago, C.; Bolhuis, P. G.; Csajka, F. S.; Chandler, D. *J. Chem. Phys.* **1998**, *108*, 1964–1977.
- Okzan, S. B.; Dill, K. A.; Behar, I. *Protein Sci.* **2002**, *11*, 1958–1970.
- Okzan, S. B.; Dill, K. A.; Behar, I. *Biopolymers* **2003**, *68*, 35–46.
- Chekmarev, D.; Ishida, T.; Levy, R. M. *J. Phys. Chem. B* **2004**, *108*, 19487–19495.
- Evans, D. A.; Wales, D. J. *J. Chem. Phys.* **2004**, *121*, 1080–1090.
- Kitchen, D. B.; Hirata, F.; D. A. Kofke, D. A.; Westbrook, J. D.; Yormush, M.; Levy, R. M. *J. Comput. Chem.* **1990**, *11*, 1169–1180.
- Jorgensen, W. L.; Maxwell, D. S.; Tirado-Rives, J. *J. Am. Chem. Soc.* **1996**, *118*, 11225–11236.
- Gallicchio, E.; Levy, R. M. *J. Comput. Chem.* **2004**, *25*, 479–499.
- Qiu, D.; Shenkin, P. S.; Hollinger, F. P.; Still, W. C. *J. Phys. Chem. A* **1997**, *101*, 3005–3014.
- Johnson, N. L.; Kotz, S. *Distributions in Statistics: Continuous Multivariate Distributions*; John Wiley & Sons: New York, 1972.
- Gelfand, A. E.; Smith, A. F. M. *J. Am. Stat. Assn.* **1990**, *85*, 398–409.
- Bratley, P.; Fox, B. L.; Schrage, L. E. *A Guide to Simulation*; Springer-Verlag: New York, 1987.
- Frishman, D.; Argos, P. *Proteins* **1995**, *23*, 566–579.
- Humphrey, W.; Dalke, A.; Schulten, K. *J. Mol. Graphics* **1996**, *14*, 33–38.