Published in final edited form as:

J Phys Chem B. 2012 September 13; 116(36): 11046–11055. doi:10.1021/jp304720m.

Unrestrained Computation of Free Energy along a Path

Bradley M. Dickson^{*,†}, He Huang, and Carol Beth Post^{*}

Markey Center for Structural Biology, Department of Medicinal Chemistry, Purdue University, 240 S. Martin Jischke Drive, West Lafayette, Indiana 47907-1971, United States

Abstract

We apply the adaptive biasing potential (ABP) method to optimize the principal curve defining a conformational transition between two known end states and to subsequently compute the one-dimensional potential of mean force as a function of arc length along the principal curve. This approach allows the use of the ABP method in a collective variable space of arbitrary dimension and offers several advantages over line-search methods. First, configurations are neither generated along an initial path for the transition nor equilibrated during evolution of the path. Second, and most importantly, the powerful sampling provided by the ABP serves to accelerate the dynamics during the optimization and computation of the free energy. Finally, the free energy is formulated as a potential of mean force that captures changes in the reaction channel along the principal curve, in contrast to the free energy profile evaluated from the local free-energy gradient in restrained path optimization methods. We first demonstrate the ABP formulation of path optimization using a two-dimensional potential surface and then with a more complex system of Src protein tyrosine kinase. The method is shown to be efficient and robust in the case of rugged, free-energy landscapes.

1. INTRODUCTION

In this report we propose an efficient methodology for the study of conformational transitions in biomolecular systems and in general the study of any process that can be

© 2012 American Chemical Society

^{*}Corresponding Author: bdickso@ad.unc.edu; cbp@purdue.edu.

[†]Present Address: Center for Integrative Chemical Biology and Drug Discovery, Division of Chemical Biology and Medicinal Chemistry, University of North Carolina, 120 Mason Farm Road, Chapel Hill, North Carolina 27599.

Notes

The authors declare no competing financial interest.

formulated in terms of a reaction pathway. The problem of finding an ideal reaction pathway is viewed here in three steps: First, to choose suitable collective variables, second to optimize a path through the collective variable space, and finally to compute the free energy along the optimal path. The first of these steps, choosing collective variables (CVs), is crucial but complicated given that each system may be best suited by different CVs. We will assume that this first step is completed and good CVs have been chosen. The second step, to identify a pathway in the CV space that is in some sense optimal, is the focus of this report. Computing the free energy along the path is straightforward once an optimal path is in hand.

Let Ω denote the space defined by the CVs. Our aim is to elucidate a characteristic pathway through Ω and to compute the free energy along that pathway when Ω is a high-dimensional space. There are several existing frameworks for this task,^{1–6} and there is a recent report which describes a number of problems encountered by these approaches in practice.⁶ Reference 6 deals specifically with the string method in collective variables, but the technical difficulties encountered there will be similar for any approach that uses restraints to localize trajectories to a path in Ω . The authors found that an additional sampling technique, Hamiltonian exchange, was required to achieve adequate sampling of the configuration space in a reasonable time. Although not discussed in ref 6, there is also the chance that roughness in the free energy landscape may frustrate a line-search type of path optimization.

It is also worth noting that the line-search methods compute free energy of the optimized path from the local free-energy gradient.^{1–3,5,6} This one-dimensional profile of free energy does not capture important information about the reaction channel width. Changes in the width of the channel will impact the size and location of the barrier and could cause both attributes to differ from the one dimensional profile.

Adaptive biasing force^{7,8} (ABF) and adaptive biasing potential^{9–12} (ABP) methods allow one to compute the free energy in a low dimensional Ω without employing restraints. The notion that restrained sampling in free energy computations might impede convergence was used to motivate ABF early in its development.⁷ An ABF or ABP formulation of the string method would facilitate good sampling along the path in Ω without any additional layers of methodology while still providing a framework for free energy computation.

Here, we follow a natural means of adapting the ABF or ABP paradigm to the problem at hand by considering the potential of mean force (or free energy) along a curve γ in Ω , so that the biasing force or potential is constructed on the one-dimensional arc length along γ . A benefit of this formulation in terms of γ is the use of an ABF or ABP method in a collective variable space of arbitrary dimension. The recent development of the finite temperature string method^{13,14} (FTS) affords a simple recipe for optimizing the curve γ by computing the mean position of trajectories in cross sections along the path and updating γ to coincide with that mean. This optimization scheme converges γ to the “principal curve” for the transition.

To describe the principal curve, first consider a free energy landscape with a single reaction channel, or valley, connecting two stable states A and B . The principal curve is a curve that runs through this channel. One may define a hyperplane at each point along this curve such that the tangent of the curve is orthogonal to the hyperplanes. The principal curve lies at the point in each plane that corresponds to the mean position of the system in that plane. We do caution that “orthogonal” has to be interpreted in the correct coordinate space, as shown in ref 14.

The simplification provided by defining a path with reduced dimensionality in an otherwise complex space dates back at least several decades where it allowed computation of the potential of mean force (PMF) along a path defined by a sequence of points through the

configuration space while avoiding a rigorous parametric definition of the path.¹⁵ The concept has also been used in connection with transition state theory.^{16,17} Jónsson also described the fundamental component of updating the path to coincide with the mean position in a cross section to relax an initially suboptimal path. More recently, the idea of updating the path to coincide with the average position in a cross section was refined as an optimization protocol in the finite temperature string method¹³ and later motivated in the collective variable space as a means of overcoming the expected roughness of the free energy landscape.¹⁴ Moreover, computing the free energy as a function of the one-dimensional arc length allows the computation of a free energy that will reflect changes in reaction channel width. This quantity is more informative than the one-dimensional profile as discussed above.

What we introduce below is an ABP formulation of the finite temperature string^{13,14} or of the reversible work formulation of free energy.^{16,17} This approach automatically, allows trajectories to move along the pathway and does not confine them with a restraint or cell wall. In fact, the biasing potential serves to flatten the free energy along the pathway thereby accelerating the transition. This feature mimics the improved sampling that was observed for restrained trajectories while exchanging restraint potentials⁶ but does so in a natural way without any additional exchange algorithm. Finally, because the method computes the principal curve, it allows γ to be optimized even when the free energy is rugged.

A significant reduction in complexity also occurs with this formulation because no initial path in configuration space must be generated. No computational resources are spent generating (e.g., via targeted, steered or biased molecular dynamics¹⁸) an initial trajectory in the configuration space that spans the conformational transition. Instead, one simply “draws” an initial path through Ω and then launches ABF/ABP from the known stable state(s). Equilibrated structures along the initial path or any subsequent path are not required. Generating meaningful initial paths through the configuration space is an active research area^{6,18,19} and the present approach completely omits this step.

In section 2 we develop the ABP method for sampling the space around a one-dimensional curve γ in the collective variable space Ω . In section 3 we present the algorithm for optimizing γ and computing the potential of mean force by performing simulations with the ABP method. In section 4 we present numerical demonstrations and compare with the maximum flux transition path (MFTP) method for a realistic model of a Src protein tyrosine kinase and a smooth Ω . In section 5 we make our concluding remarks.

2. ADAPTIVELY BIASED DYNAMICS

In this section we present our computational approach. This is broken into four parts. First, the one-dimensional potential of mean force along a curve in Ω is defined. Second, the biased dynamics are proposed and third the biasing force in configuration space is given. Lastly, we discuss discretization and parameter choice.

2.1. Free Energy along a Curve

In the following we denote the configuration space as χ . The collective variables are $\xi(x)$ where $x \in \chi$ and $\xi \in \Omega$. The configuration space is n -dimensional, whereas the CV space is N -dimensional. This corresponds to a chemical or biological system with $n/3$ particles and N collective variables.

Employing a notation close to that of ref 14, we make the following definitions regarding the path. Let the curve γ be parametrized by $\phi(\lambda)$ with $0 \leq \lambda \leq \Lambda$, where Λ is the total length of the curve in Ω . Take $\phi(\lambda = 0) \in A$ and $\phi(\lambda = \Lambda) \in B$, where A and B are some

particular well-defined basins in Ω . For any $\lambda \in [0, \Lambda]$, $\phi(\lambda)$ returns a point in Ω on the curve γ .

The collective variable space can be partitioned according to the arc length λ by associating to each value of λ the set of points in Ω that are closer to $\phi(\lambda)$ than any other point on γ . Partitioning Ω this way defines hyperplanes that intersect γ . The hyperplane at λ is given by the set of points $P_\lambda = \{\xi \mid d(\xi, \phi(\lambda)) < d(\xi, \phi(\lambda')) \forall \lambda' \neq \lambda\}$, with the distance given by

$$d(\xi, \phi(\lambda)) = \sqrt{\sum_{i,j=1}^N (\xi_i - \phi_i(\lambda)) D(\xi)_{ij}^{-1} (\xi_j - \phi_j(\lambda))} \quad (1)$$

where

$$D(\xi)_{ij} = \frac{\langle \delta(\xi(x) - \xi) \nabla_i \xi(x) m^{-1} \nabla_j \xi(x) \rangle}{\langle \delta(\xi(x) - \xi) \rangle} \quad (2)$$

The angle brackets indicate an average of property P with respect to the canonical measure

$$\langle P \rangle = \frac{\int_{\chi} P e^{-\beta V(x)} dx}{\int_{\chi} e^{-\beta V(x)} dx}$$

Throughout we use $V(x)$ to note the potential energy function and $\beta = 1/k_B T$ where k_B is the Boltzmann constant and T is temperature. The justification of this particular choice for $d(\cdot, \cdot)$ is given in ref 14. There it is shown that the level sets of the commitor function are approximated by the level sets of λ near the principal curve of the transition. The hyperplanes P_λ are thus expected to approximate level sets of the commitor in the vicinity of this principal curve.

The configuration space can be associated to P_λ by the argmin function

$$\lambda(\xi(x)) = \text{argmin}[d(\xi(x), \phi(\lambda)), \lambda] \quad (3)$$

This function returns the arc length λ corresponding to the nearest point on γ , which is $\phi(\lambda)$.

The PMF (up to an unimportant additive constant) as a function of λ could now be computed by integrating the Dirac delta function $\delta(\lambda(\xi(x)) - \lambda)$ against the Boltzmann density on χ (see eq 4 below). At each value of λ , the integration would need to be performed over the whole plane P_λ . In addition to practical considerations, the integration over P_λ should be limited to a portion of the hyperplane that is near γ for two reasons. First, the relationship between the commitor function and the level sets of $\lambda(\xi(x))$ only holds in the vicinity of the path.¹⁴ Second, one would expect that far from $\phi(\lambda)$ there are structures in the planes P_λ that do not resemble the set of configurations at $\phi(\lambda)$. Including these structures in the PMF computation will obfuscate detailed information regarding the reaction channel being examined. These points motivate limiting the extent of the planes P_λ away from $\phi(\lambda)$ when computing the PMF associated with a particular reaction channel, which is our current goal.

By defining the sets $B_\lambda = \{\xi \mid d(\xi, \phi(\lambda)) < R\}$ in Ω , one can form a tube around ϕ as the union of these sets, $\mathcal{T}[\phi] = \bigcup_{\lambda=0}^{\Lambda} B_\lambda$. The tube $\mathcal{T}$ in the collective variable space is simply the collection of all of the points in Ω that are within a distance R from at least one point on the curve γ . (Choosing the radius R will be discussed in the next section.) The tube can be broken into slices S_λ by the intersection $S_\lambda = \mathcal{T}[\phi] \cap P_\lambda$. The slice S_λ is the portion of P_λ that is inside the tube. In order to limit the integration to S_λ , we introduce the “tube potential” as

$$V_{\mathcal{T}[\phi]}(\xi(x)) = \begin{cases} 0 & \text{if } \xi(x) \in \mathcal{T}[\phi] \\ \infty & \text{otherwise} \end{cases}$$

By adding this potential to the Hamiltonian, we focus only on those configurations that fall within the tube in Ω . Notice that, when $\lambda = 0$ or Λ , S_λ is a hemisphere, whereas for interior values $0 < \lambda < \Lambda$, the slices S_λ are planes.

Finally, the free energy $A(\lambda)$ (up to an additive, unimportant constant ζ) of a slice S_λ that intersects the curve at $\phi(\lambda)$ is

$$\zeta e^{-\beta A(\lambda)} = \lim_{\alpha \rightarrow 0} Z^{-1} \int_{\mathcal{X}} \delta_\alpha(\lambda(\xi(x)) - \lambda) e^{-\beta[V(x) + V_{\mathcal{T}[\phi]}(\xi(x))]} dx \quad (4)$$

where Z is a normalization constant and

$$\delta_\alpha(\lambda) = \exp\left(-\frac{\lambda^2}{\alpha^2}\right)$$

In practice we will approximate eq 4 by making computations with a finite value of α . Below, we use A_α to indicate this approximation. The strategy to compute $A_\alpha(\lambda)$ and error associated with the finite α approximation are discussed next.

2.2. Biased Dynamics

We compute $A_\alpha(\lambda)$ via the following adaptively biased dynamics. The central feature of these dynamics is that a bias potential, constructed on-the-fly, gradually flattens free energy barriers resulting in an accelerated sampling of the collective variable space. Initially the bias potential is zero everywhere, and over time it adapts to mimic the underlying free energy. These dynamics are a summary of the results in ref 20 where it was shown that these dynamics are equivalent to the well-tempered metadynamics.¹¹

A trajectory x_t is a solution to the biased Langevin equation

$$d\dot{x}_t = -f\dot{x}_t dt - \frac{1}{m} \nabla(V + V_{\mathcal{T}[\phi]} \circ \xi + V_b \circ (\lambda \circ \xi))(x_t) dt + \sqrt{\frac{2f\beta^{-1}}{m}} dB_t$$

where f is the Langevin friction coefficient, B_t is a standard Brownian motion, and m is mass. The bias potential V_b , up to an arbitrary constant, is

$$V_b(\lambda, t) = \beta^{-1} \frac{b}{1-b} \ln[c(1-b)h(\lambda, t) + 1] \quad (5)$$

where

$$h(\lambda, t) = \int_0^t \delta_\alpha(\lambda(\xi(x_s)) - \lambda) ds \quad (6)$$

The histogram $h(\lambda, t)$ counts visits to the region around λ and the bias potential at λ increases with each visit so as to slowly flood the stable states of $A(\lambda)$. The size of the region around λ is controlled by α . The bias potential is built from the sampling history of x_t and in the long-time limit is a good approximation to $-bA(\lambda)$ when α is small. We refer to ref 20 for a derivation of this bias potential. The parameters $0 < c$ and $0 < b < 1$ appear in the definition of the bias potential and control how the bias couples to the dynamics and what fraction of the free energy will be canceled by the bias potential, respectively. The coupling c has inverse time units.

The gradient of V_b in configuration space can be expressed through the chain rule,

$$-\frac{\partial V_b}{\partial x_j} = -\frac{\partial V_b}{\partial \lambda} \sum_{i=1}^N \frac{\partial \lambda}{\partial \xi_i} \frac{d\xi_i}{dx_j} \quad (7)$$

for $j = 1, 2, \dots, n$. The derivative $\partial V_b / \partial \lambda$ is computed from two histograms h and h' accumulated over the evolution of x_t

$$\frac{\partial V_b(\lambda, t)}{\partial t} = \frac{bc\beta^{-1}h'(\lambda, t)}{1+c(1-b)h(\lambda, t)} \quad (8)$$

where

$$h'(\lambda, t) = \int_0^t \partial_\lambda \delta_\alpha(\lambda(\xi(x_s)) - \lambda) ds = \int_0^t \frac{2(\lambda(\xi(x_s)) - \lambda)}{\alpha^2} \delta_\alpha(\lambda(\xi(x_s)) - \lambda) ds \quad (9)$$

At $t = 0$, $\nabla V_b = 0$, and $h = h' = 0$ for all λ . The derivatives $\partial \xi_i / \partial x_j$ are straightforward, whereas $\partial \lambda / \partial \xi_i$ needs some explanation.

For a fixed value of λ , we define the following plane in Ω

$$\vec{n}_\lambda \cdot (\xi - \phi(\lambda)) = 0$$

where the vector $\vec{n}_\lambda$ is the unit normal for the plane. After taking the gradient with respect to ξ

$$\nabla_\xi [\vec{n}_\lambda \cdot (\xi - \phi(\lambda))] = 0$$

and rearranging for $\nabla_\xi \lambda$, we have

$$\nabla_\xi \lambda = \frac{\vec{n}_\lambda}{\vec{n}_\lambda \cdot \frac{\partial \phi(\lambda)}{\partial \lambda} - (\xi - \phi(\lambda)) \cdot \frac{\partial \vec{n}_\lambda}{\partial \lambda}} \quad (10)$$

This expression can be problematic when the denominator is small, as discussed below. The gradient of V_b can now be expressed in Langevin's equation by using eqs 10 and 8 in eq 7.

We approximate the tube potential as a harmonic restraint

$$V_{T|\phi}(x) = \begin{cases} 0 & \text{if } d(\xi(x), \phi(\lambda)) \leq R \\ \frac{\kappa_T}{2} (d(\xi(x), \phi(\lambda)) - R)^2 & \text{otherwise} \end{cases} \quad (11)$$

where $\phi(\lambda)$ is the point on the curve closest to $\xi(x)$. This is a good approximation when the force constant κ_T is large.

The free energy in eq 4 is estimated from the histogram $h(\lambda, t)$ as²⁰

$$A_\alpha(\lambda, t) = -\beta^{-1} \frac{1}{1-b} \ln[h(\lambda, t)] \quad (12)$$

The error associated with finite α in eq 12 is understood as resulting from a convolution with the Gaussian δ_α and is roughly characterized by $\ln[1 + \alpha^2]$.^{12,20} The following section covering discretization motivates a very small α so this error can be ignored in practice with little effect.

In summary, eq 12 can be used to compute the one-dimensional free energy, or PMF, via the biased Langevin equation given in eq 5. The last step to obtain a practical method is to discretize time and γ and to specify values for the parameters α and R .

2.3. Discretization and Parameter Choice

To discretize the curve ϕ , divide the total length into N_Λ equal segments and replace λ with $\lambda_i = (i-1)\Lambda/N_\Lambda$. The set of $\phi(\lambda_i)$ for $i = 1, 2, \dots, N_\Lambda$ are connected with linear segments, giving a piecewise linear approximation to ϕ . In practice only the positions of the points $\phi(\lambda_i)$ are needed, the linear interpolation between them is not. The domain of the histograms h and h' is now discretized and the biasing force eq 7 is computed by approximating $\lambda(\xi(x_t))$ by the λ_i that is closest in value. For a fine grid, this approximation is very small.

Specification of α is based on the discretization of ϕ and can be related to a harmonic force constant¹² k : $\alpha = (2k_B T/k)^{1/2}$, providing some intuition for specifying α by connecting with the more familiar concept of choosing a restraint strength. Imagine that the restraint is used to hold a trajectory near a specific plane and choose k appropriately. Bear in mind that this analogy is simply to understand α . No restraint along λ is used in this approach. In the numerical examples below k is a multiple of $k_B T$, which is common practice for restrained simulations, and N_Λ is fixed so that $2 - \alpha N_\Lambda/\Lambda$.

The condition on N_Λ , $2 - \alpha N_\Lambda/\Lambda$, is such that the Gaussians δ_α will be well represented by the discretization of λ . There is one Gaussian δ_α centered at each discrete “grid” point λ_i for $i = 1, \dots, N_\Lambda$ and this condition on N_Λ ensures that the Gaussians overlap well enough to recover a smooth free energy. Also notice that with the definition of $V_{\mathcal{T}}[\phi]$ in eq 11, N_Λ should be large enough to avoid significant corrugation of the tube walls. (Note that the wall of the tube $V_{\mathcal{T}}[\phi]$ is exactly given by the boundaries of the sets B_λ . If these sets do not overlap by much, the tube walls will be rough.)

To discretize time, divide the interval $[0, t]$ into N_t equal parts of duration $\Delta t = t/N_t$. The Langevin trajectory x_t is replaced by a discrete-time trajectory x_i and the total time is given by $t = N_t \Delta t$. In discrete time, the derivative in eq 8 becomes

$$\frac{\partial V_b(\lambda, N_t)}{\partial \lambda} = \frac{bc\beta^{-1}h'(\lambda, N_t - 1)}{1+c(1-b)h(\lambda, N_t - 1)} = \frac{\Delta t bc\beta^{-1} \sum_{i=0}^{N_t-1} \partial_\lambda \delta_\alpha(\lambda(\xi(x_i)) - \lambda)}{1+\Delta t c(1-b) \sum_{i=0}^{N_t-1} \delta_\alpha(\lambda(\xi(x_i)) - \lambda)} \quad (13)$$

At each time step of the simulation, the histograms h and h' are updated by $\delta_\alpha(\lambda(\xi(x_i)) - \lambda)$ or $\lambda \delta_\alpha(\lambda(\xi(x_i)) - \lambda)$, respectively, for all values of λ within 5 or 6 multiples of α of $\lambda(\xi(x_i))$.

Only three arrays must be stored during this ABP: the summations appearing in the numerator and denominator of eq 13 and the points $\phi(\lambda_j)$ in Ω .

3. OPTIMIZING $\phi(\lambda)$ AND COMPUTING $A_\alpha(\lambda)$

Outlined here is an algorithm to converge an initial ϕ to the principal curve for the transition. By defining the average in discrete time and space with the dynamics driven by $(V + V_{\mathcal{G}_\phi} + V_b)$

$$\langle \xi \rangle_j = \frac{\sum_{i=0}^{N_t} \delta_\alpha(\lambda(\xi(x_i)) - \lambda_j) \xi(x_i) e^{\beta V_b(\lambda(\xi(x_i)), i)}}{\sum_{i=0}^{N_t} \delta_\alpha(\lambda(\xi(x_i)) - \lambda_j) e^{\beta V_b(\lambda(\xi(x_i)), i)}} \quad (14)$$

for each λ_j , the principal curve can be defined as the curve satisfying $\phi(\lambda_j) = \langle \xi \rangle_j$. The reweighting factor is

$$e^{\beta V_b(\lambda, t)} = \left(\frac{c(1-b)h(\lambda, t) + 1}{c(1-b)\max_\lambda [h(\lambda, t)] + 1} \right)^{b/1-b}$$

The bias potential V_b is so far determined only to an additive constant, as only the gradient of the potential is needed for the dynamics. Here we take the liberty of choosing this constant so that the zero of energy is not drifting in time. This relationship between the bias and the histogram h is derived in reference 20. The purpose of the reweighting factors in eq 14 is the following: if the dynamics converge then the probability density on χ is proportional to the Boltzmann factor

$$e^{-\beta(V + V_m \mathcal{F}[\phi] + V_b)}$$

The reweighting in eq 14 ensures that the principal curve computed from the biased dynamics is the same as what would be computed from unbiased dynamics.

When the points $\phi(\lambda_j)$ fall along the principal curve, evaluating eq 14 for each λ_j along ϕ would simply return the point $\phi(\lambda_j)$. However, eq 14 is a conditional mean. It is conditioned on the current ϕ and returns the principal curve for the tube around ϕ . This is only equivalent to the principal curve for the transition when eq 14 returns the current ϕ . Next we discuss constructing the initial path and the computation of some preliminary quantities. Then, we give an algorithm for efficiently computing the principal curve for the transition.

3.1. Initial Path, R , and $D(\xi)$

From unbiased dynamics in A and B compute: $\xi_A = \langle \xi(x) \rangle_A$, $\xi_B = \langle \xi(x) \rangle_B$, $\sigma_A^2 = \langle (\xi(x) - \xi_A)^2 \rangle_A$ and $\sigma_B^2 = \langle (\xi(x) - \xi_B)^2 \rangle_B$. We use $\langle \cdot \rangle_s$ to denote a time average taken over unbiased dynamics in state s . The points ξ_A and ξ_B give an idea of the location of A and B in Ω while

σ_A and σ_B give a basic estimate of a tube radius. One last thing that must be computed from the unbiased trajectories is the tensor $D(\xi)$ defined in eq 2. From the unbiased trajectories one can compute $\langle D(\xi(x)) \rangle_A$ and $\langle D(\xi(x)) \rangle_B$. We make the a priori assumption that $D(\xi)$ is a constant over Ω and checking that the average of $D(\xi)$ is statistically equivalent in A and B validates this somewhat. This assumption was also made in implementing the FTS method and is discussed in ref 14.

As we noticed earlier, S_0 and S_Λ are hemispheres. To get an accurate PMF for the transition, $\phi(0)$ and $\phi(\Lambda)$ should rest outside of A and B , respectively. This practice will keep the volume of the hemispheres from corrupting the configurational density in the planes S_λ near the minima of A and B . To do this, simply draw a straight line from ξ_A to ξ_B and stretch the line by 10–15%. This stretched path serves as the initial ϕ and the points $\{\phi(i)\}_{i=1}^{N_\Lambda}$ are to be equally spaced along this initial path. Only the points $\phi(i)$ must be defined; no equilibration or other dynamics is needed. To be sure the hemispheres do not corrupt the PMF, the end points are fixed throughout the computation of the principal curve. We point out that while we have used a linear path as the initial path, any initial path would suffice. The linear path is used here for convenience.

3.2. Optimizing ϕ

The algorithm employed here is as follows: Given an initial $\phi_0 \in \Omega$, the principal curve ϕ_1 inside $\mathcal{T}_{[\phi_0]}$ is converged. (The sense of “converged” is made precise below.) The principal curve in $\mathcal{T}_{[\phi_0]}$ is then converged and so on. After collecting i updates of the principal curve, the distance between pairs of curves $R(\phi_j, \phi_i)$ (defined below) can be computed for $j = 0, 1, 2, \dots, i-1$. The distance between ϕ_0 and ϕ_i will be large and as j approaches i the distance will shrink. If a set of the i curves are falling in the same reaction channel and differ only in fluctuations, the distance between ϕ_j and any curve in this set will be roughly the same. The graph of $R(\phi_j, \phi_i)$ against j will show a plateau when a series of curves lies in the same channel as ϕ_i . In an ideal case with no statistical noise, the plateau will be at $R(\phi_j, \phi_i) = 0$ for a given reaction channel. Because there will be noise in the statistics giving rise to the curves ϕ , the plateau will occur at a nonzero value.

We define the distance $R(\cdot, \cdot)$ as a Fréchet distance. The ordered sequences $\{\phi_i(\lambda_l)\}_{l=1}^{N_\Lambda}$ and $\{\phi_j(\lambda'_l)\}_{l=1}^{N_\Lambda}$ define polygonal curves in Ω . A “coupling” between ϕ_i and ϕ_j is a sequence of links

$$(\phi_i(\lambda_{l_1}), \phi_j(\lambda'_{l'_1})), (\phi_i(\lambda_{l_2}), \phi_j(\lambda'_{l'_2})), \dots, (\phi_i(\lambda_{l_m}), \phi_j(\lambda'_{l'_m}))$$

where $l_1 = 1$ (the first point on the path), $l_m = N_\Lambda$ (the last point on the path), $l_{q+1} = l_q$ or $l_q + 1$ and similarly for l' . This scheme for indexing the segments is given in ref 21. This choice respects the ordering of the polygonal segments along the two curves and precludes backward coupling along the curves. This coupling also allows a point on one curve at one arc length to be coupled to a point on the other curve that is at a different value of the arc length. In general, the two paths need not have the same number of points. The distance between ϕ_i and ϕ_j is finally

$$F(\phi_i, \phi_j) = \min_C \left[\frac{1}{m} \sum_{q=1}^m d(\phi_i(\lambda_{l_q}), \phi_j(\lambda'_{l'_q})) \right] \quad (15)$$

where C is the space of all couplings between ϕ_i and ϕ_j and m is the number of links in the coupling. This is a variant of the discrete Fréchet distance²¹ between two polygonal curves

where an average distance is used instead of the maximum distance over all connections in the coupling. We found that using an average rather than the maximum results in a smoother, easier to read graph of F . In Figure 5 below an example of the graph of $R(\phi_j; \phi_i)$ can be seen.

To fix a stopping time for the optimization of ϕ , we compute the mean of the curves on the plateau of F , denoted $\langle \phi \rangle$. If an additional observation of ϕ_i does not significantly alter the mean, the algorithm is terminated. The meaning of “significant” is controlled by fixing a tolerance. Below we use $\langle \phi \rangle_{\text{old}}$ and $\langle \phi \rangle_{\text{new}}$ to distinguish between the average of curves on the plateau before and after a new observation, respectively.

Consider the i th iteration of such a scheme. At the beginning of the i th iteration, no blocks of dynamics have been carried out, $N_B = 0$, and the ABP has accumulated no information, $V_b = 0$. The dynamics of the i th iteration are driven by $V + V_{\mathcal{T}_{[\phi_{i-1}]}} + V_b$. The algorithm for the i th iteration is as follows.

1. Evolve M trajectories, or replicas, with the potential $V + V_{\mathcal{T}_{[\phi_{i-1}]}} + V_b$ (where V_b is evolving) for a block of N_t steps, update the number of blocks $N_B = N_B + 1$.
2. Compute an observation of ϕ_i at block N_B

$$\widehat{\phi}_{N_B}(\lambda_i) = \frac{\sum_{l=1}^M \sum_{k=0}^{N_B N_t} \delta_\alpha(\lambda(\xi(x_k(l))) - \lambda_j) \xi(x_k(l)) e^{\beta V_b(\lambda(\xi(x_k(l))), k)}}{\sum_{l=1}^M \sum_{k=0}^{N_B N_t} \delta_\alpha(\lambda(\xi(x_k(l))) - \lambda_j) e^{\beta V_b(\lambda(\xi(x_k(l))), k)}} \quad (16)$$

and apply a smoothing and spacing routine (for example from FTS¹⁴ or as described by eq 18).

3. If all slices $\mathcal{S}_\lambda$ of $\mathcal{T}_{[\phi_{i-1}]}$ have been visited $\mathcal{N}$ times
 - a. save the principal curve $\phi_i \leftarrow \widehat{\phi}_{N_B}$
 - b. save the simulation time of the i th iteration $t_i = N_B \times N_t \times \Delta t$
 - c. set $i = i + 1$, $N_B = 0$, initialize ABP and center the tube potential on ϕ_i
 - d. check the graph of F for plateau
 - if no plateau, go to 1.
 - if there is a plateau but $R(\langle \phi \rangle_{\text{old}}, \langle \phi \rangle_{\text{new}}) > \epsilon$, go to 1
 - if there is a plateau and $F(\langle \phi \rangle_{\text{old}}, \langle \phi \rangle_{\text{new}}) < \epsilon$, stop.

For economy the observations of ϕ_i should be monitored frequently after small intervals of dynamics. The interval here is set by N_t . N_B is the number of these intervals in a given tube and is not predetermined and thus the total computation time is more meaningful than the number of updates of the curve.

It is assumed that there are M initial phase points, ideally half in A and half in B , at thermal equilibrium from which independent ABP simulations can be launched. Each trajectory is independent of the others and evolves with its own adaptive potential V_b . The M trajectories evolve in parallel so the clock τ is representative of the cost of the simulation while $M \times \tau$ is representative of the total sampling.

While optimizing ϕ , the radius of the tube R should be larger than $\max[\sigma_A, \sigma_B]$ to allow the possibility of exploring multiple channels. To avoid the possibility of numerical instabilities resulting from eq 10, we approximate $\nabla_\xi \lambda$ with

$$\nabla_{\xi} \lambda = \frac{\vec{n}_{\lambda}}{\vec{n}_{\lambda} \cdot \frac{\partial \phi(\lambda)}{\partial \lambda}} \quad (17)$$

while optimizing the principal curve for the transition. During the free energy computation, we use a more narrow tube and use the full gradient in eq 10.

3.3. Computing $A_q(\lambda)$

Once we have converged the principal curve for the transition, which is $\langle \phi \rangle$ at the exit of the algorithm in the previous subsection, computing the free energy is a straightforward application of the ABP dynamics. The computation is done in a step separate from optimization of the principal curve so that the histograms h and h' only contain information about the final, converged principal curve for the transition. To compute the free energy we suggest R is set to $\max[\sigma_A, \sigma_B]$ while no other parameters need to be adjusted. One may imagine several schemes for choosing R . We employ a very simple one here that in general is not “bullet-proof”. It is possible to measure collisions with the tube wall, something like the pressure inside the tube, and to choose the radius to minimize or control the number of these collisions. Here, we simply fix the radius.

It may be most efficient to allow the trajectories to combine information about the free energy with some regular frequency but we do not employ these techniques here.²² The free energy is recovered via eq 12.

4. NUMERICAL EXAMPLES

4.1. 2D Demonstration and Further Implementation Details

We use the two-dimensional potential from refs 23 and 24 to give a concrete example of optimizing the curve ϕ and of computing the PMF $A_q(\lambda)$. The surface in Figure 1 proposes two mechanisms for transitioning from state A to state B . There is a direct route which crosses a high barrier and a more curved, indirect route that crosses a lower barrier. The two pathways are separated by a second order saddle indicated by + in Figure 1. At the thermal energy $k_B T = 0.07$, the lower barrier is preferred by unbiased dynamics.²⁴

The initial straight-line path $\phi(\lambda)$ (shown in panel A of Figure 2) was discretized so that $\alpha N_{\lambda}/\Lambda = 20$ and $\alpha^2 = 1/1200$, which can be viewed as a restraint force constant of $2400 k_B T$ ($N_{\lambda} = 2000$). This strong choice for α provides more resolution than what is needed, but it ensures the errors associated with finite α are small and does not incur an increase in overhead. We found that $\sigma_A = 0.2$ and the tube radius was set to $4\sigma_A$ to compute the principal curve.

In the first iteration of the algorithm given in section 3.2, the estimated principal curve $\langle \xi \rangle_{\lambda}$ (shown in panel A of Figure 2) moves from the channel with the higher energy saddle to the channel with the lower energy saddle. This is consistent with what we would expect from a trajectory confined to the plane at the ridge between A and B ; the trajectory should spend more time near the lower energy saddle. The principal curve for the reaction converges after eight updates and is shown in Figure 2 along with the final slices, S_{λ} .

It is possible to sample different mechanisms if, and only if, the tube is wide enough to encompass them and if each S_{λ} can be sufficiently sampled. Local roughness on the free energy surface would likewise be easily overcome. (Sampling two channels of equal probability would force the principal curve to lie between the channels rather than in one channel. See ref 14 for a discussion of this and similar situations.) Of course, if the tube is

very narrow multiple channels cannot be sampled. In general the optimal tube width will be unknown. We present here only one of many possibilities for specifying the width.

Computing the PMF is a straightforward application of the biased dynamics generated by eq 5 where λ is the collective variable and the gradient of λ is given by eq 10. We use the same bias parameters that were used to compute the principal curve for the transition. We set the tube radius to $R = \sigma_A = 0.2$. The free energy is computed from the adaptively biased dynamics by eq 12. We assume convolution errors are small because α is small. As a comparison, eq 4 was computed by direct numerical integration. The results are shown in Figure 3 along with the potential energy evaluated at the principal curve. The small difference between the PMF and the potential energy results from a narrowing of the reaction channel inside the tube. The tube has a finite width and the value of the potential energy across the tube varies significantly compared to β near the dividing surface for the transition. This constitutes a narrowing of the channel. Depending on the shape of the reaction channel within the tube, the barrier in the PMF may not in general align with the barrier on the underlying surface. The entropic contributions to the PMF provide important information that would be overlooked when evaluating the many dimensional free energy profile along the one-dimensional principal curve. Typically, only the one-dimensional profile is computed (using the local free energy gradient) and information about the shape of the reaction channel does not come to bear on the results. It is important to keep in mind that if the tube is more narrow than the reaction channel, the entropic contributions would be missed. The effects of tube width are shown for this simple case in the inset of Figure 3. The width estimated from σ_A is enough to capture the entropic contribution made by the positive curvature mode at the saddle point for transition. This is of course not generally true and one will not be able to ensure that all entropic contributions are captured.

The effects of tube width on the PMF are demonstrated in the inset.

We note that the present method requires a more aggressive smoothing than string methods for two reasons: (1) the number of “nodes” (i.e., the points $\phi(\lambda_i)$) is very large compared to implementations of the string method and (2) we need to minimize the number of intersections among S_λ . The noise in the mean at each plane can result in extremely sharp turns of the estimated curve because the planes are very close together. To smooth the curve, we do not employ the smoothing scheme from ref 1, which has little impact on paths with many images and thus requires a number of iterations to smooth the path. Instead, we choose to mollify the principal curve. This approach results in an easy-to-control algorithm and requires far fewer iterations than the smoothing from the string method. We use the mollification

$$\phi(\lambda_j) = \frac{\sum_{i=0}^{\Lambda} e^{-\left(\frac{\lambda_i - \lambda_j}{n d \lambda}\right)^2} \phi_{\text{raw}}(\lambda_i)}{\sum_{i=0}^{\Lambda} e^{-\left(\frac{\lambda_i - \lambda_j}{n d \lambda}\right)^2}} \quad (18)$$

where $d\lambda$ is the distance between images on the path and n is some integer controlling the strength of the smoothing. We used ϕ_{raw} to indicate the raw, unsmoothed but equally spaced data from simulation. Once one estimate of the principal curve is computed, the smoothing and resampling should be tuned, by considering different values of n in eq 18 to give the curve a satisfactory amount of smoothness. This tuning only needs to be done once in our experience but more iterations could be required if the original data are very rough.

4.2. Src Kinase Gō Model

We consider the conformational transition of the ~30 kDa catalytic domain of a Src family protein kinase, Lyn, between a catalytically active and inactive form. The energy function

used in this study is a double-basin Gō model that features two energy basins corresponding to the active and inactive structures of the molecule. It was constructed by merging two single-basin Gō models V_{act} and V_{ina} via the exponential average:

$$V_{\text{double}} = -\frac{1}{\mathcal{B}} \ln \{ \exp[-\mathcal{B}(V_{\text{act}} + \mathcal{G})] + \exp(-\mathcal{B}V_{\text{ina}}) \} \quad (19)$$

where $\mathcal{G}$ is an energy offset modulating the relative stability of the two states and $\mathcal{B}$ a coupling constant controlling the barrier height. In the present study, $\mathcal{G}$ was set to -4.0 kcal/mol and $\mathcal{B}$ was set to 0.02 mol/kcal. Similar procedures for merging two Gō models have been described elsewhere.^{25,26} The individual single-basin models were built following the Gō potential described by Karanicolas and Brooks,^{27,28} taking the active and inactive homology models of Lyn reported in ref 29 as reference structures. More details of this double-basin Gō model can be found in ref 30. The coarse-grained inactive and active structures are shown in Figure 4, where the main structural features are labeled.

Below we make a comparison between the ABP method described above and the MFTP method.⁵ The MFTP behaves like a minimum free energy path with the addition of a temperature dependent curvature penalty. At zero temperature the penalty vanishes and the MFEP is recovered. Because the MFTP differs from the string method only in a single term, which is only weakly influential in the present case, the following comparison is thus in essence between the ABP method and the string method in collective variables.^{1,6} A number of difficulties faced by the string method were noted in the Introduction above and reviewed in ref 6.

4.2.1. Collective Variables—For this study we use seven collective variables, selected on physical grounds and used previously to study the mechanism of activation in the Src-Gō model described above.³⁰ Each of the CVs is a sum of C_{α} pair distances that connects one region of the molecule to another region of the molecule. Figure 4 identifies the regions, and Figure 6 includes a list of the CV group pairs. The CV are defined by

$$\xi_i = \sum_{j=1}^{m_i} \frac{r_j^{\text{ina}}(i) - r_j^{\text{act}}(i)}{\|r_j^{\text{ina}}(i) - r_j^{\text{act}}(i)\|} r_j(i)$$

Here, m_i is the number of pairs in the i th CV, $r_j(i)$ is the pair distance for the j th pair in the CV. The reference values r_j^{act} and r_j^{ina} are the equilibrium distances computed from the active and inactive reference structures, respectively. The number of pairs in a CV ranges from three to twenty, and the list of the exact pairs is given in ref 30.

Several initial paths were optimized in ref 30 and each one converged to the same basic mechanism, implying the space Ω is smooth. Some small degree of variability was observed in the CV describing the motion of the A-loop near the N-lobe (shown in green in Figure 6). Note that only the half of the A-loop that is closest in residue number to the N-lobe was used in the collective variables. The C-terminal half closest in number to the C-lobe was found to be very mobile in another study and was therefore omitted from the CV.³⁰ This C-terminal section of the A-loop is transparent in Figure 4.

4.2.2. ABP Parameters—For the ABP computation, we find $\sigma_A = 6$ Å and $\sigma_B = 7$ Å and optimize the principal curve with $R = 15$ Å, about 2 times larger than $\max[\sigma_A, \sigma_B]$. We use an α that equates to a force constant of $25k_B T$. This value is 2.5 times larger than the value used to restrain the trajectories in MFTP. For the tube potential, we use a force constant of

$100k_B T$. This choice for α and the distance between ξ_A and ξ_B (the mean positions described in section 3) suggests that we cut the initial linear path into $N_\Lambda = 5000$ segments. This resolution is higher than suggested in the Discretization subsection but it does not sacrifice much in the way of efficiency. We expect the initial linear path to take some curvature during optimization and we should make sure that the optimized path still satisfies the rule $2 - \alpha N_\Lambda / \Lambda$. One could choose to do so by optimizing the number of segments at each update during the principal curve calculation, so that the number of segments is always minimized. This would help efficiency particularly when the dimensionality of Ω is very large. In a large dimensional Ω , the computational overhead is manifest in evaluating $\lambda(\xi(x)) = \operatorname{argmin}[\mathcal{d}(\xi(x), \phi(\lambda)), \lambda]$. In this example, we fix N_Λ for the whole computation.

To compare to MFTP, we run the algorithm from section 3.2 with 30 replicas. The MFTP comprises 40 nodes, and the MFTP parameters (e.g., the path time step, equilibration time and sampling time) were adopted from ref 30. The MFTP converges well with a sampling time (per node on the path) of just 0.05 ns and an equilibration time of 0.005 ns.

For the ABP computations here, we use $\Delta t = 0.055$ ns to match the block time of the MFTP. We use $\mathcal{N} = 100$ and $\epsilon = 0.05$ Å to control convergence of the principal curve. $\mathcal{N}$ designates the minimum number of times that each plane must be visited before accepting an estimate of the principal curve inside a tube. The uncertainty of the principal curve for the transition is controlled by ϵ . These parameters were introduced in section 3.2. These choices were made with no experimentation. The 30 replicas were initially split evenly between the active and inactive states: 15 trajectories begin in the inactive state and 15 trajectories begin in the active state. The initial configurations are simply thermally equilibrated structures in the active and inactive states. The initial path was constructed as a straight line.

The biasing parameters were $b = 0.9$ and $c = 0.1 \Delta t^{-1}$ for all computations. This choice of c limits the time derivative of the bias to roughly $k_B T/11$

$$c = \frac{\Delta t^{-1}}{(11+1)b - 1}$$

This relation comes from equating the maximum of the time derivative of the bias potential with $k_B T/11$. Using $b = 0.9$ implies that about 90% of the free energy barrier along λ will be canceled by the bias.

4.2.3. Principal Curve and PMF—In Figure 5 we show the time required to converge the principal curve for the MFTP and the ABP. While it is not possible to compare the values of the Fréchet distances directly, we may compare the simulation time required to reach a plateau for the two methods. The ABP scheme is seen to be slightly more efficient here but the stepsize of the MFTP updates could potentially be further optimized. That the ABP approach is competitive likely results from the fact that a single update might yield a large displacement of the path, as is evident in the steepness of $F(j,i)$ in Figure 5. Note that we used the Euclidean distance instead of $\mathcal{d}(\cdot, \cdot)$ from eq 1 in eq 15 after verifying that this change of metric does not alter the results. The use of the Euclidean metric is afforded here because all of the CV have the same units. To estimate the computational efficiency of the ABP method, we track the total simulation time rather than using the number of iterations as the computational time (number of blocks) needed to converge the principal curve varies for each iteration. The Fréchet distance to the final path is therefore shown as a function of the total simulation time in Figure 5.

While the purpose of this report is not to discuss details and relevance of the computed path, the principal curve for each CV is displayed in the lower panel of Figure 6 and in the upper panel are results from the MFTP. This path is discussed in detail in ref 30, and we find that the MFTP and the ABP converge to the same basic mechanism.

The PMF from ABP (Figure 7) is computed by using the full biasing gradient in eq 5 (i.e., using eq 10 rather than 17). We compute the PMF with $R = 7 \text{ \AA}$. The PMF is computed by allowing the 30 replicas to evolve independently for $0.11 \text{ }\mu\text{s}$. The PMF $A_\alpha(\lambda, t)$ is computed from

$$A_\alpha(\lambda, t) = -\beta^{-1} \frac{1}{1-b} \ln[h(\lambda, t)/\max_\lambda[h(\lambda, t)]] \quad (20)$$

where the 30 independent histograms are combined as

$$h(\lambda, t) = \sum_{l=1}^M h(\lambda, l, t) \quad (21)$$

In Figure 7 $A(\lambda, t)$ is compared to the one-dimensional “free energy” found with the MFTP. We use “free energy” to highlight the fact that the string method evaluates the multidimensional free energy at points along the path. This “free energy” is a function of the many dimensional collective variable $\xi(x)$. The quantity $A_\alpha(\lambda)$, computed here, is a true potential of mean force and contains information about the shape of the reaction channel that the string or MFTP would be insensitive to.

We estimate the level of convergence by comparing the 30 independent estimates of the free energy with the global estimate given by eq 20. An estimate of the PMF can be computed for each replica of the dynamics by

$$A_\alpha(\lambda, l, t) = -\beta^{-1} \frac{1}{1-b} \ln[h(\lambda, l, t)/\max_\lambda[h(\lambda, l, t)]] \quad (22)$$

If the dynamics are ergodic, $A_\alpha(\lambda, \infty) = A_\alpha(\lambda, l, \infty)$ for all l . The variance

$$\sigma^2 = \frac{1}{29} \sum_{l=1}^{30} (A_\alpha(\lambda, l, t) - A_\alpha(\lambda, t))^2$$

or the standard deviation σ , among the individual PMF therefore gives an indication of convergence. The estimated error in the mean, $\sigma/\sqrt{30}$ is shown in Figure 7. We stress that this estimate is a much more strict notion of convergence than what would be used to estimate error for the string or MFTP, where each node on the path independently produces an error at a point along the path. In contrast, the error estimate for the ABP result requires each replica to produce a reasonable estimate of the PMF, not simply a reasonable estimate at one point on the PMF. With this in mind, we note that the simulation cost associated with the ABP approach, which computes the PMF, was roughly 4.5 times smaller than the computational cost associated with the restrained sampling of the MFTP/ string for the same level of error, 0.5 kcal/mol. This speaks to the efficiency of the adaptively biased sampling.

5. CONCLUSION

We have developed and tested an adaptively biased approach to computing the principal curve and PMF along that curve. The result is an efficient and robust computational tool for studying conformational transitions in biological or chemical systems where the free energy landscape is expected to be rugged. This method affords efficient sampling without the use of additional computational tools, as have been used in conjunction with the string method.⁶ Because the principal curve is the targeted object, rugged free energy landscapes will not trap the optimization.¹⁴ With the use of an ABP-dynamics comes the added benefit that no initial path in the configuration space must be generated and no time is spent equilibrating images in the configuration space at any point in the algorithm. These two features contribute significantly to the efficiency and ease-of-use of the method. In a benchmark of the method we observe this ABP approach is twice as fast as the maximum flux path (or similarly the string method) while neglecting the time required to obtain an initial path for the MFTP. In practice the effort to generate the initial path is significant^{6,19} and is not always repeatable.¹⁸

The present approach computes a potential of mean force along the principal curve. The MFTP,⁵ string method,¹ and swarm method² only evaluate the free energy along a line in the collective variable space and thus omit information about the width of the reaction channel. Changes in the width of the channel can have a great impact on the free energy profile, altering both the height and location of barriers. In the present case, the PMF along the principal curve will correctly capture this additional information.

REFERENCES

1. Maragliano L, Fischer A, Vanden-Eijnden E. *J. Chem. Phys.* 2006; 125:024106–024180.
2. Pan AC, Sezer D, Roux B. *J. Phys. Chem. B.* 2008; 112:3432–3440. [PubMed: 18290641]
3. Gan W, Yang S, Roux B. *Biol. Phys. J.* 2009; 97:L08–L10.
4. Chen M, Yang W. *J. Comput. Chem.* 2009; 30:1649–1653. [PubMed: 19462399]
5. Zhao R, Shen J, Skeel R. *J. Chem. Theory Comput.* 2010; 6:2411–2423. [PubMed: 20890401]
6. Ovchinnikov V, Karplus M, Vanden-Eijnden E. *J. Chem. Phys.* 2011; 134:085103. [PubMed: 21361558]
7. Darve E, Pohorille A. *J. Chem. Phys.* 2001; 115:9169–9183.
8. Darve E, Rodriguez-Gomez D, Pohorille A. *J. Chem. Phys.* 2008; 128:144120–144133. [PubMed: 18412436]
9. Wang F, Landau D. *Phys. Rev. Lett.* 2001; 84:2050–2053. [PubMed: 11289852]
10. Laio A, Parrinello M. *Proc. Natl. Acad. Sci. U.S.A.* 2002; 99:12562–12566. [PubMed: 12271136]
11. Barducci A, Bussi G, Parrinello M. *Phys. Rev. Lett.* 2008; 100:020603–020607. [PubMed: 18232845]
12. Dickson BM, Lelièvre T, Stoltz G, Legoll F, Fleurat-Lessard P. *J. Phys. Chem. B.* 2010; 114:5823–5830. [PubMed: 20380408]
13. E W, Ren W, Vanden-Eijnden E. *J. Phys. Chem. B.* 2005; 109:6688–6693. [PubMed: 16851751]
14. Vanden-Eijnden E, Venturoli M. *J. Chem. Phys.* 2009; 130:194103–194120. [PubMed: 19466817]
15. Elber R. *J. Chem. Phys.* 1990; 93:4312–4321.
16. Schenter GK, Mills G, Jónsson H. *J. Chem. Phys.* 1994; 101:8964–8971.
17. Jóhannesson GH, Jónsson H. *J. Chem. Phys.* 2001; 115:9644–1957.
18. Huang H, Ozkirimli E, Post CB. *J. Chem. Theory Comput.* 2009; 5:1304–1314.
19. Ovchinnikov V, Karplus M. *J. Phys. Chem. B.* submitted.
20. Dickson BM. *Phys. Rev. E.* 2011; 84:037701–037705.
21. Eiter T, Mannila H. Technical Report CD-TR 94/64, Christian Doppler Laboratory for Expert Systems. 1994

22. Minoukadeh K, Chipot C, Lelièvre T. *J. Chem. Phys.* 2009; 130:1008–1017.
23. Chen L, Nash P, Horing N. *J. Chem. Phys.* 2005; 123:091041–091046.
24. Dickson B. *J. Chem. Phys.* 2007; 127:064106–064111. [PubMed: 17705587]
25. Best RB, Chen YG, Hummer G. *Structure.* 2005; 13:1755–1763. [PubMed: 16338404]
26. Yang S, Roux B. *PLoS Comput. Biol.* 2008; 4:e1000047–14. [PubMed: 18369437]
27. Karanicolas J, Brooks CL. *Protein Sci.* 2002; 11:2351–2361. [PubMed: 12237457]
28. Karanicolas J, Brooks CL. *J. Mol. Biol.* 2003; 334:309–325. [PubMed: 14607121]
29. Ozkirimli E, Post CB. *Protein Sci.* 2006; 15:1051–1062. [PubMed: 16597828]
30. Huang H, Zhao R, Dickson BM, Skeel RD, Post CB. *J. Phys. Chem. B.* 2012; 116:4465–4475. [PubMed: 22448785]

Figure 1. There are two stable states labeled A or B, two first order saddles labeled x, and one second order saddle labeled +. Contours are placed at intervals of $k_B T$.

Figure 2.

(A) Initial path shown in black and the first estimate of the principal curve shown in orange. The blue circle indicates the width σ_A as described in the text, and the tube radius for the principal curve calculation was $4\sigma_A$. Every 30th slice, S_λ , is shown in black. (B) The current path (black) and the estimate (orange) after four iterations. (C) The current path (black) and estimate (orange) after eight iterations. (D) The final tube is based on the principal curve after eight iterations. The width of the tube is $\sigma_A = 0.2$ and is highlighted in blue.

Figure 3. Potential energy and the PMF along the principal curve (as shown in panel D of Figure 2). The PMF computed from the biased dynamics is in agreement with the result obtained by direct numerical integration of eq 4.

Figure 4. Inactive and active forms of the kinase domain of a Srcfamily member, Lyn. The main conformational differences between the two forms are changes in the A-loop (yellow), α C helix (green), and relative orientation of the N- (red) and C-lobes (blue). The structures shown are homology modeled from PDB entries 1QCF (inactive) and 3LCK (active).²⁹

Figure 5. Fréchet distance, $F(j,i)$, for ABP and MFTP, where the path at iteration j is compared to the final path at iteration i . The total computational time for the j -th iteration is shown on the abscissa.

Figure 6. Results of the ABP computation and the MFTP computation are shown here. The MFTP is shown with points while the principal curve is shown with smooth lines.

Figure 7. PMF from ABP (shown in red) and the one-dimensional free energy from MFTP (shown in green). The estimated error for the ABP result is shown in blue.