

Method for Computing Protein Binding Affinity

CHARLES F. F. KARNEY,¹ JASON E. FERRARA,¹ STEPHAN BRUNNER^{2,*}

¹Sarnoff Corporation, Princeton, New Jersey 08543-5300

²Locus Pharmaceuticals, Inc., Blue Bell, Pennsylvania 19422-2700

Received 14 July 2004; Accepted 30 October 2004

DOI 10.1002/jcc.20167

Published online in Wiley InterScience (www.interscience.wiley.com).


Abstract: A Monte Carlo method is given to compute the binding affinity of a ligand to a protein. The method involves extending configuration space by a discrete variable indicating whether the ligand is bound to the protein and a special Monte Carlo move, which allows transitions between the unbound and bound states. Provided that an accurate protein structure is given, that the protein–ligand binding site is known, and that an accurate chemical force field together with a continuum solvation model is used, this method provides a quantitative estimate of the free energy of binding.

© 2004 Wiley Periodicals, Inc. J Comput Chem 26: 243–251, 2005

Key words: free energy; binding affinity; Monte Carlo methods; equilibrium constants; proteins

Introduction

Many drugs work by binding to a target protein in an organism and affecting the action of this protein. The binding of the drug molecule, the ligand L , to the protein P under physiological conditions is usually reversible (characterized by weak chemical interactions rather than covalent bonds),



and, in equilibrium and in the dilute limit, the concentration of the ligand–protein complex $[LP]$ is given by the dissociation constant

$$K_d = \frac{[L][P]}{[LP]}.$$

It is convenient to define the binding affinity as

$$pK_d = -\log_{10}\left(\frac{K_d/N_A}{1 \text{ kmol m}^{-3}}\right),$$

where N_A is the Avogadro constant. A high value for pK_d is crucial to obtaining a good drug molecule and, consequently, the ability to compute pK_d accurately would greatly accelerate drug discovery by allowing many molecules to be screened *in silico* before any time-consuming syntheses and assays are done. The dissociation constant can also be expressed in terms of the binding free energy ΔF ,

$$K_d = \exp(\beta\Delta F)/V_0, \quad (1)$$

where V_0 is the system volume, $\beta = 1/(kT)$, T is the temperature, and k is the Boltzmann constant. Similarly, pK_d is given by

$$pK_d = -\frac{\beta\Delta F}{\ln 10} + \log_{10}(V_0 N_A \times 1 \text{ kmol m}^{-3}).$$

The quantity ΔF is the free energy difference of the ligand and the protein forming a bound complex LP , the “bound system,” compared to the ligand and the protein isolated from one another $L + P$, the “unbound system.”

To compute pK_d , we require (1) a sufficiently accurate model of the protein and the ligand and their interaction, and (2) a good way to compute the resulting value of pK_d . In this article, we assume the first requirement is fulfilled and instead focus on meeting the second.

The standard methods of computing free energies^{1,2} are not capable of computing ΔF directly because the unbound and bound systems are too dissimilar, which hinders transitions between these systems. Instead, typically, two close ligands L_a and L_b , are compared separately unbound and bound to the protein, thereby obtaining the difference in the free energies $\Delta F_a - \Delta F_b$.

We present here a practical method for directly computing ΔF , and hence pK_d . The method consists of: (1) formulating the problem in an extended phase space, which allows the unbound

Correspondence to: C. F. F. Karney; e-mail: ckarney@sarnoff.com

*Current address: CRPP, Ecole Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland

Contract/grant sponsor: U.S. Army Medical Research and Materiel Command; contract/grant number: DAMD17-03-C-0082

and bound systems to be treated as a single system and K_d to be expressed as the ratio of two canonical averages; (2) introducing a new Monte Carlo move, the “wormhole move,” to make transitions between the unbound and bound states in this extended system; and (3) a method to determine the “portals” needed for the wormhole move.

Formulation

Consider a system of volume V_0 consisting of a ligand molecule L and a protein molecule P in a solvent. The state of the system is given by $\Gamma = [\Gamma_L, \Gamma_P]$ where Γ_M represents the phase space configuration (position, orientation, and conformation) of molecule M . In equilibrium, the system obeys the Boltzmann distribution³

$$f(\Gamma) = \frac{\exp[-\beta E(\Gamma)]}{\int \exp[-\beta E(\Gamma)] d\Gamma},$$

where $E(\Gamma)$ is the energy of the system with the equilibrium effects of the solvent implicitly included as a solvation free energy. In this article, we will assume that $E(\Gamma)$ is given.

Typical molecular interactions have a short range. In view of this, let us define Σ_0 to represent all accessible Γ space (i.e., L and P somewhere in the system volume V_0), and Σ_1 to represent that portion of Σ_0 where there is an appreciable interaction between L and P , which therefore form the complex LP . In the dilute limit, $V_0 \rightarrow \infty$, the dissociation constant can then be written as^{1,4,5}

$$K_d = \frac{1}{V_0} \frac{\int_{\Sigma_0} \exp[-\beta E_0(\Gamma)] d\Gamma}{\int_{\Sigma_1} \exp[-\beta E_1(\Gamma)] d\Gamma}, \quad (2)$$

where $E_1(\Gamma) = E(\Gamma)$ is the full energy and $E_0(\Gamma)$ is the unbound energy given by ignoring the interaction between the protein and the ligand. The Helmholtz free energy of the unbound and bound systems is³

$$F_\lambda = -\frac{1}{\beta} \ln \left(\int_{\Sigma_\lambda} \exp[-\beta E_\lambda(\Gamma)] d\Gamma \right),$$

for $\lambda = 0$ and 1 , and eq. (1) is obtained from eq. (2) with $\Delta F = F_1 - F_0$. The definition of K_d , eq. (2), is strictly independent of V_0 because of translational symmetry (ignoring boundary effects). It is also independent of the precise definition of Σ_1 , provided that Σ_1 includes the protein–ligand binding site and does not include a “macroscopic” volume beyond this.

In this formulation, we have assumed that the system volume V_0 is fixed. However, in most physiological systems, the pressure is held constant and the binding affinity is then related to the differences in the Gibbs free energy, which introduces a correction term that is the product of the pressure and the change in the volume caused by the formation of the LP complex.⁶ We expect this correction to be small for typical ligand–protein interactions in solution.

We would like to cast eq. (2) as the ratio of canonical averages that can be computed using the canonical-ensemble Monte Carlo method.⁷ To achieve this, we combine the unbound and bound systems by extending phase space to include a discrete index $\lambda \in \{0, 1\}$ and consider a system in this extended space with energy $E_\lambda^*(\Gamma)$ for which the canonical average is defined by

$$\langle X \rangle = \frac{\sum_\lambda \int d\Gamma \exp[-\beta E_\lambda^*(\Gamma)] X_\lambda(\Gamma)}{\sum_\lambda \int d\Gamma \exp[-\beta E_\lambda^*(\Gamma)]}.$$

We take $E_\lambda^*(\Gamma)$ to be infinite for $\Gamma \notin \Sigma_\lambda$, and finite otherwise. Now eq. (2) can be rewritten as

$$K_d = \frac{1}{V_0} \frac{\langle \delta_{\lambda 0} \exp(-\beta[E_0(\Gamma) - E_0^*(\Gamma)]) \rangle}{\langle \delta_{\lambda 1} \exp(-\beta[E_1(\Gamma) - E_1^*(\Gamma)]) \rangle}, \quad (3)$$

where $\delta_{\lambda\mu}$ is the Kronecker delta. If $E_\lambda^*(\Gamma) \approx E_\lambda(\Gamma)$, the terms being averaged are $O(1)$. Because the definition of K_d is independent of V_0 , we can pick $V_0 \sim 1/K_d$ so that approximately the same number of samples contribute to each of the canonical averages. We show later, eq. (7), that this choice minimizes the error in the estimate of K_d .

We can evaluate $E_\lambda^*(\Gamma)$ with short energy cutoffs allowing it to be computed more quickly than $E_\lambda(\Gamma)$ and the terms contributing to the averages in eq. (3) can be accumulated every hundred steps, for example. Because there is typically a high correlation between successive steps in a Monte Carlo simulation, this method allows the averages to be computed to a given degree of accuracy in less time than if we had used $E_\lambda^*(\Gamma) = E_\lambda(\Gamma)$.

The extension of phase space has been used in other free energy calculations, to combine, for example, systems at several different temperatures⁸ or to treat the “reaction coordinate” controlling the transition between two chemical species as a dynamic variable.^{9,10} In our case, the use of the wormhole Monte Carlo (described in the next section) allows us to include just the two systems of interest without the need to compute the properties of (possibly unphysical) intermediate systems.

Wormhole Monte Carlo

We can compute the canonical averages in eq. (3) using the Monte Carlo method⁷ to make steps from $[\Gamma, \lambda]$ to $[\Gamma', \lambda']$ with probability

$$\min[1, \exp(-\beta[E_\lambda^*(\Gamma') - E_\lambda^*(\Gamma)])].$$

However, the estimate of the ratio in eq. (3) will be very poor, because transitions between $\lambda = 0$ and 1 will be extremely rare—typically, E_0^* is shallow and wide, while E_1^* is deep and narrow; see Figure 1a. One possible way of remedying this is to treat λ as a continuous variable,^{9,10} providing a suitably interpolated definition of $E_\lambda(\Gamma)$, and allowing Monte Carlo steps with small changes in λ . In practice, this approach only “works” if the two end points are sufficiently similar, thus limiting this method to the comparison of chemically close molecules.




Figure 1. (a) Schematic representation of $E_0^*(\Gamma)$ (shallow and wide) and $E_1^*(\Gamma)$ (deep and narrow). Conventional Monte Carlo moves between $\lambda = 0$ and 1 (shown as dashed lines) are nearly always rejected because they lead to large increases in energy. (b) Schematic representation of typical portals for the wormhole moves for the case illustrated in (a). The large ratio of the volume of the unbound ($\lambda = 0$) portal compared to the bound ($\lambda = 1$) portals compensates for the higher energy of the unbound configurations. This results in accepted wormhole moves (dashed lines) between all the portals.

Here we propose an alternative way of carrying out a Monte Carlo simulation of the $E_\lambda^*(\Gamma)$ system. We restrict the standard moves to changes in Γ only, and allow changes in λ via “wormhole moves” that connect otherwise disconnected regions of configuration space. (The term “wormhole” is employed to evoke the idea of interstellar travel via wormholes.^{11a} This wormhole move is not to be confused with a different move of the same name used in polymer simulations.^{11b}) This obviates the need to treat intermediate values of λ , permitting us to compute the free-energy differences needed to determine the absolute binding affinity.

Assume we have some system defined on a phase space Y whose equilibrium distribution is proportional to $g(Y)$. The canonical average of a quantity $X(Y)$ is defined by

$$\langle X \rangle = \frac{\int dY g(Y) X(Y)}{\int dY g(Y)}.$$

In our application, we make the identification $Y = [\Gamma, \lambda]$, $\int dY = \sum_\lambda \int d\Gamma$, and $g(Y) = \exp[-\beta E_\lambda^*(\Gamma)]$.

Let us define a set of “portal functions,” w, w', w'', \dots , on Y , with properties

$$0 \leq w(Y) \leq 1/v < \infty,$$

$$\int dY w(Y) = 1,$$

where v is a representative phase-space volume of the portal function. A wormhole move consists of the following steps: select a pair of portals (w, w') with probability $p_{ww'}$; reject the move with probability $1 - vw(Y)$, where Y is the current state; otherwise, with probability $vw(Y)$, pick a configuration Y' with probability $w'(Y')$; and accept the move to Y' with probability $P_{ww'}(Y, Y')$. If the move is rejected, Y is retained as the new state. We determine $P_{ww'}(Y, Y')$ by demanding that the rate of making transitions from Y to Y' via portals (w, w') is balanced by the reverse rate from Y' to Y via portals (w', w), that is,

$$R_{ww'}(Y, Y') = R_{w'w}(Y', Y)$$

(the condition of detailed balance), where the rates are given by

$$R_{ww'}(Y, Y') = p_{ww'} \frac{g(Y)}{\int dY g(Y)} vw(Y) w'(Y') P_{ww'}(Y, Y').$$

A possible solution for the acceptance probability is

$$P_{ww'}(Y, Y') = \min\left(1, \frac{p_{w'w} g(Y') v'}{p_{ww'} g(Y) v}\right),$$

where we have made use of the identity $\alpha \min(1, \alpha^{-1}) = \min(1, \alpha)$, which is valid for $\alpha > 0$.

To apply this move, let us use a specific rather simple form for the portal functions, $w(Y)$, namely

$$w(Y) = \begin{cases} 1/v, & \text{for } Y \in w, \\ 0, & \text{otherwise,} \end{cases}$$

where we now denote a portal by w , which defines an arbitrary subset of Y space of volume v . In principle, there is no restriction on the choice of the portals; however, practical considerations, discussed below, dictate how they are chosen. Furthermore, for simplicity, we will assume that the wormhole probabilities are all equal, $p_{ww'} = \text{const}$.

Let us now describe the wormhole move in $[\Gamma, \lambda]$ space. Starting with a configuration $[\Gamma, \lambda]$, first pick a random portal w . If $[\Gamma, \lambda] \notin w$, then reject the move. Otherwise, pick a random configuration $[\Gamma', \lambda']$ uniformly in a randomly chosen portal w' , and accept the move to $[\Gamma', \lambda']$ with probability

$$\min\left(1, \frac{\exp[-\beta E_\lambda^*(\Gamma')] v'}{\exp[-\beta E_\lambda^*(\Gamma)] v}\right). \quad (4)$$

This differs from the standard Boltzmann acceptance probability by the ratio of volumes, v'/v , which can be large enough to compensate for the difference in the mean energies in the portals. Indeed, if the mean energy of configurations in w scales as $\beta^{-1} \ln v + \text{const.}$, the acceptance probability, eq. (4), is $O(1)$, thereby allowing moves between shallow wide wells and deep narrow ones; see Figure 1b. In practice, each portal will occupy one of the energy wells of either $E_0^*(\Gamma)$ or $E_1^*(\Gamma)$. This implies that the $\lambda = 0$ portals will permit unrestricted translation and rotation of the ligand and the protein subject to the V_0 constraint, and the $\lambda = 1$ portals will allow unrestricted motion of one of the molecules.

The wormhole move embodies two concepts that have been used separately in other works: stretching or shrinking Γ space when making the move compensates for possibly large energy differences between wells;^{12,13} and jumping between disconnected regions of phase space enables the Markov chain to explore regions of phase space separated by large energy barriers.^{14,15} In the following sections, we will show how these elements may be combined to permit the computation of protein binding affinities with realistic force fields.

Finding the Portals

For the wormhole method to be practical, we need a reliable way of choosing the portals. We describe this process first in the general case. Let us write $\Gamma = [\Gamma_L, \Gamma_P]$, where $\Gamma_M = [X_M, \Xi_M]$ represents the state of molecule M , X_M represents its position and orientation, and Ξ_M represents its conformation.

We assume that Ξ_M is expressed in such a way that any constraints on the positions of the atoms in M (e.g., bond lengths and bond angles) are implicitly accounted for, so that the dimensionality of Ξ_M reflects the number of degrees of conformational freedom, n_M , for this molecule.

Because of the translational and rotational symmetry of the system, certain components of Γ are ignorable. We can therefore write $E_\lambda^*(\Gamma) = E_\lambda^*(\Xi_\lambda)$, with

$$\Xi_0 = [\Xi_L, \Xi_P],$$

$$\Xi_1 = [Y, \Xi_L, \Xi_P],$$

where $Y = X_L - X_P$ denotes the position and orientation of the ligand relative to the protein. The dimensionality of Ξ_λ is n_λ with $n_0 = n_L + n_P$ and $n_1 = n_0 + 6$.

The strategy for determining the portals is to carry out conventional canonical Monte Carlo simulations with $E_\lambda^*(\Xi_\lambda)$ separately for $\lambda = 0$ and 1. For each λ , we obtain a canonical set of configurations $\{\Xi\}$ (suppressing the λ subscripts for brevity), to which we fit n -dimensional ellipsoids, as follows. First we try to fit a single ellipsoid to $\{\Xi\}$. The center of the ellipsoid is given by the mean configuration $\langle \Xi \rangle$. For each configuration in $\{\Xi\}$, we determine the deviation from the mean, $\delta \Xi = \Xi - \langle \Xi \rangle$, and compute a covariance matrix for the configurations which can be diagonalized as

$$\langle \delta \Xi \delta \Xi \rangle = \mathbf{P} \mathbf{\Lambda} \mathbf{P}^T,$$

where \mathbf{P} is the matrix of (column) eigenvectors and $\mathbf{\Lambda}$ is the diagonal matrix of eigenvalues. Because of the properties of the covariance, \mathbf{P} is real and orthogonal, $\mathbf{P}^{-1} = \mathbf{P}^T$, and the eigenvalues are real and nonnegative. If there are no hidden constraints on the motion, we additionally can assume that the eigenvalues are strictly positive. We find it convenient to define

$$\mathbf{B} = \mathbf{P} \mathbf{\Lambda}^{1/2}, \quad \mathbf{B}^{-1} = \mathbf{\Lambda}^{-1/2} \mathbf{P}^T,$$

so that we can write

$$\langle \delta \Xi \delta \Xi \rangle = \mathbf{B} \mathbf{B}^T.$$

We take the semiaxes of the ellipsoid to be the columns of $\sqrt{n} \mathbf{B}$. The multiplier here, \sqrt{n} , is chosen to ensure that $O(1)$ of the configurations in $\{\Xi\}$ lie within the ellipsoid. This choice is motivated by considering a symmetric n -dimensional Gaussian

$$f(\mathbf{r}) = \frac{\exp(-\frac{1}{2} r^2)}{(2\pi)^{n/2}},$$

for which we have

$$\langle r^2 \rangle = \int r^2 f(\mathbf{r}) d^n \mathbf{r} = n.$$

Ellipsoids are a natural choice to use to fit the set of configurations for the following reasons. (1) The iso-density contours of the distribution in a harmonic well are ellipsoids. (2) It is easy to sample points randomly from an ellipsoid. (3) Conversely, it is easy to test that a point lies inside an ellipsoid. (4) The volume of an n -dimensional ellipsoid is given by

$$v_n = \frac{\pi^{n/2}}{(n/2)!} \prod_{i=1}^n a_i,$$

where a_i is the length of i th semiaxis. Note well the degenerate case of this result, $v_0 = 1$, which is used if the protein and ligand are both rigid ($n_0 = 0$). The sampling and testing of points, (2) and (3), can either be accomplished by transforming with the matrix \mathbf{B} . Alternatively, we can use the simpler Cholesky decomposition of the covariance matrix

$$\langle \delta \Xi \delta \Xi \rangle = \mathbf{C} \mathbf{C}^T,$$

where \mathbf{C} is a lower triangular matrix. Both \mathbf{C} and \mathbf{C}^{-1} may be computed by direct (noniterative) methods.¹⁶

Having defined an ellipsoid in this way, we test its suitability as a portal by demanding that $O(1)$ of the configurations sampled uniformly from it have energies close to its mean energy $\langle E^*(\Xi) \rangle$. If this test fails, we split $\{\Xi\}$ into two sets according to the sign of $\delta \Xi$ projected along the largest semiaxis of the ellipsoid and construct new trial portals from each of these sets.

The ellipsoids that result from this process constitute our portals. When computing the volumes of the portals for use in eq. (4), we need to account for the freedom to place $2 - \lambda$ molecules at arbitrary positions and orientations in the volume V_0 . In practice, this means we multiply the volume of the $\lambda = 0$ portals by σV_0 where V_0 is the translational volume and σ is the orientational volume (given below).

To complete the specification of the portals, we need to describe how $\langle \Xi \rangle$ and $\delta \Xi$ are formed, because, as a consequence of working in the subspace where the molecular constraints are implicitly satisfied, Ξ is not simply a point in \mathbb{R}^n . We wish to represent each ellipsoid on a locally Cartesian space \mathbb{R}^n , which lets us use familiar formulas for defining the ellipsoid. We also demand that the mapping to \mathbb{R}^n have constant Jacobian to maintain detailed balance when sampling from the ellipsoids.

We illustrate this by considering the case of the protein and the ligand being made up of $n_M + 1$ rigid fragments simply connected by n_M bonds, each of which allow rotation only (i.e., the bond lengths and bond angles are fixed). The relative position and orientation of L with respect to P , Y , can be defined in terms of one of the rigid fragments of each molecule. Finally, we use unit quaternions¹⁷ to represent orientation. Because the quaternions q and $-q$ represent the same rotation, orientations are given by an *axis*¹⁸ of the sphere S^3 .

The coordinates making up Ξ are then: (a) the position component of Y , a point in \mathbb{R}^3 ; (b) the orientation component of Y , an axis of the sphere S^3 ; and (c) the dihedral angles of the rotatable bonds of the ligand and protein, points on the circle S^1 . The definition of $\langle \Xi \rangle$ and $\delta \Xi$ is straightforward for (a), because we use the normal arithmetic definitions. For (c), we define the mean for each angle¹⁸ by the direction of the mean of the points on S^1 embedded in \mathbb{R}^2 . We form the deviation in this case by subtraction modulo 2π so that the result lies in $[-\pi, \pi]$.

To find the mean of the orientations (b), we similarly embed S^3 in \mathbb{R}^4 . We define the mean orientation¹⁸ as the axis in \mathbb{R}^4 about which the moment of inertia of the sample axes is minimum. To find the deviation of the orientation, we compute the differential rotation $\delta q = q \langle q \rangle^*$, which takes the mean orientation to the sample orientation, and we project this into a “turn” vector \mathbf{u} in the unit ball in \mathbb{R}^3 so as to preserve the metric. This is achieved by a generalization of the Lambert azimuthal equal-area projection as described in the Appendix. In this representation, the volume of orientational space (which contributes to the multiplier for the unbound volumes) is $\sigma = \frac{4}{3} \pi$.

Occasionally, the point sampled from w' corresponds to one of the coordinates “wrapping around;” that is, the change in the dihedral lies outside $[-\pi, \pi]$ or the turn \mathbf{u} lies outside the unit ball. To preserve detailed balance, we reject the resulting move.

There are four ways in which we can improve the quality of the portals obtained by this method so as to make successful wormhole moves more likely. (1) The Monte Carlo runs used to obtain the samples from which the portals are defined should begin with an “annealing” phase where the temperature is started at some high value and slowly reduced to T at which point we start gathering samples. This allows the Monte Carlo sampling to explore phase space more thoroughly. (2) Other methods of finding conformational energy minima¹⁹ can be applied to provide additional starting points for the Monte Carlo runs. (3) The samples can be

supplemented with those obtained by applying those symmetry operations that leave the molecules invariant. In the case of non-chiral ligands, we would also apply a reflection of the ligand in the unbound case, because this will leave the energy unchanged. In this way, the portal moves allow all the symmetric variants of the molecules to be explored so that symmetry is included in a systematic way in the computation of the binding affinity. (4) When forming $\langle \delta \Xi_0 \delta \Xi_0 \rangle$, we should set the intermolecular cross terms to zero because the conformations of the two molecules are independent when $\lambda = 0$.

This method of finding portals depends on the samples “spanning” a volume of phase space. This requires that the dimensionality of phase space be sufficiently small, and this, in turn, implies the use of an implicit solvation model. In addition, the portals can be more reliably found if the “hard” degrees of freedom in the molecules are replaced by constraints (e.g., by fixing the bond lengths and bond angles as described above).

The method of successively subdividing the samples may lead to suboptimal portals in some cases, for example, by dividing a contiguous set of samples. We have recently experimented with fitting a mixture of Gaussians to the samples using the expectation-maximization (EM) algorithm.^{20,21} Because this optimizes the fit to all the samples, it usually results in fewer, better, portals. In this case, we are naturally lead to use Gaussians for the portal functions rather than the more restrictive ellipsoids. By enforcing the symmetries of the ligand when making the fits, ligand symmetry can also be included in a rigorous way. These improvements will be described in a subsequent publication.

The Free Energy Calculation

Prior to the calculation of the free energy, we estimate a suitable value of V_0 , which enters into the definition of the volumes of the $\lambda = 0$ portals, by assigning an estimated statistical weight of $v \exp(-\beta \langle E^* \rangle)$ to each portal and estimating

$$V_0 \sim \frac{\sum_{\lambda=1} v \exp(-\beta \langle E^* \rangle)}{\sum_{\lambda=0} (v/V_0) \exp(-\beta \langle E^* \rangle)},$$

where the sums in numerator (denominator) are over the bound (unbound) portals and v/V_0 for the unbound portals is the volume of the ellipsoid (multiplied by σ) and thus does *not* depend on V_0 . If this estimate for V_0 results in the sampling being too heavily weighted toward $\lambda = 0$ or 1, V_0 may be changed and the current values of the sample sums for K_d can be adjusted to account for this change.

The choice of the starting configuration for the free energy calculation may introduce some bias in the results. We can remove much of this bias by using the portals to select the starting configuration: select a portal w with probability proportional to its estimated statistical weight; select a configuration, $[\Gamma, \lambda]$, uniformly from this portal; and accept the configuration with probability

$$\min[1, \exp(-\beta[E_\lambda^*(\Gamma) - \langle E^* \rangle])],$$




Figure 2. The structure of *p*-amino-benzamidine.

where $\langle E^* \rangle$ is the canonical average energy over the portal. This procedure is repeated until a configuration is accepted. The bias can be further reduced by running the Monte Carlo calculation for several correlation times, defined by eq. (6), prior to gathering the data for eq. (3).

During the course of the free energy calculation, normal and wormhole Monte Carlo moves are mixed. With the normal moves, we only attempt to change Γ and, for this reason, it is possible to have the Monte Carlo step size depend on λ (with, usually, the step size being larger with $\lambda = 0$). In practice, most ($\sim 90\%$) of the attempted moves are wormhole moves because frequently we have $[\Gamma, \lambda] \notin w$ and the attempted wormhole move is inexpensively rejected.

The method is robust in the sense that it does not depend on the particular form of the energy function. In addition, a failure to make transitions between $\lambda = 0$ and 1 can be detected. This may be because the test $[\Gamma, \lambda] \in w$ never succeeds (i.e., the configuration has been trapped in a new well), or because the acceptance criterion is never met, which indicates that there is a deep well within the well of one of the portals. In both cases, the problem can be corrected by adding a new portal based on recent configurations.

Example

The efficacy of the wormhole method depends on how frequently a configuration lies within one of the portals and how often the jump to the new portal is accepted. To assess these questions, we have computed the binding affinity of *p*-amino-benzamidine, whose structure is shown in Figure 2, to the digestive enzyme trypsin, at $T = 290$ K. We emphasize that the primary goal of this exercise is to assess how well wormhole moves allow the free-energy calculation to converge. For this purpose, we are not so interested in comparing the resulting computed binding affinity to the experimental data because this will depend in large measure on the accuracy of the force field and of the protein structure. Nevertheless, because the convergence will depend on the complexity of the energy “landscape,” we use this example to epitomize the binding of a small ligand to a protein.

The coordinates of the atoms in trypsin are taken from a trypsin–benzamidine complex, 1BTY.²² At physiological pH, the amidine group is protonated (net charge of +1) and, in the complex, it is attracted to a negatively charged aspartate residue in trypsin inhibiting its enzymatic action. We employ the Amber 7 force field^{23–25} and the GB/SA solvation model.^{26–28} The protein is taken to be rigid, $n_p = 0$, and two bonds of the ligand are

allowed to rotate, namely, those connecting the amidine and the amino groups to the benzene ring, $n_L = 2$. The published force field²³ does not provide a satisfactory torsion for the bond between the benzene ring and the amidine group, and this term was determined using Gaussian 98²⁹ as $[-14.2 \cos(2\phi) + 3.3 \cos(4\phi) + 0.5 \cos(6\phi)]$ kJ/mol, where ϕ is the dihedral angle.

Five unbound and five bound canonical simulations of 1000 steps each were carried out to find the portals. The resulting configurations were supplemented by those obtained by applying the symmetry operations that leave the ligand invariant. This gave 16 unbound and 8 bound portals. The configurations of the bound portals are all the same “pose” of the ligand on the protein and correspond to the symmetries of *p*-amino-benzamidine given by rotating the amino, benzene, and amidine groups by 180° . The unbound ligand also exhibits eight symmetries given by rotating the amino and amidine groups by 180° relative to the central benzene ring and by including the mirror images. However, because the bond parameters allow partially free rotation of the amidine group, each symmetric set of configurations is represented by two portals.

We estimated a suitable value for V_0 of $0.39 \times 10^{-18} \text{ m}^3$ based on the volumes and the mean energies of the computed portals. During the binding affinity calculation, wormhole moves were attempted on 90% of the steps (the other 10% were standard Monte Carlo moves). Of these attempted wormhole moves, 97% failed (inexpensively) because the configuration was not in the chosen w . Of the remaining 3%, about 60% lead to a successful move, of which about 40% involved switching from $\lambda = 0$ to 1 and vice versa. The major computational cost in the free energy calculation is the evaluation of the bound energies. In this example, the wormhole moves required less than three bound energy evaluations, on average, to effect the transition from a bound configuration to an unbound one and back. This enables an accurate estimate to be made of the ratio of the averages in eq. (3), which after 5×10^6 steps yields $\text{p}K_d = 7.99 \pm 0.01$. The error estimate is derived in the next section and represents a 2% relative error in K_d .

This example illustrates that the method is effective at allowing sufficient transitions between the bound and unbound states to enable the binding affinity of protein–ligand systems to be accurately computed. We note that our computed binding affinity differs from the experimental results of 5.1 to 5.2.^{30,31} This discrepancy may be accounted for by modest, $\sim 20\%$, errors in the force field.

Error Analysis

To assess the errors in the computation of $\text{p}K_d$ in more detail, we carried out 10 independent runs similar to the one described in previous section. Each of these used the same portals and the same value of V_0 . We computed cumulative estimates for $\text{p}K_d$ based on the first s steps of the Markov chains. When forming the averages in eq. (3) we sample every 100th step. (As we shall see, there is a high degree of correlation within 100 steps; thus, there would be little improvement in the estimate of $\text{p}K_d$ by sampling more frequently.) The results for five of these runs are shown in Figure 3. The convergence toward the mean is as $1/\sqrt{s}$; after 5×10^6 steps, the error in $\text{p}K_d$ has been reduced to about ± 0.01 .




Figure 3. Cumulative estimates $pK_d(s)$ obtained by sampling every 100th step from the first s steps of five independent Monte Carlo runs. The dashed lines shows convergence as $1/\sqrt{s}$ to the mean value of 7.99.

For the purposes of further analysis, let us assume that the computation is carried out with $E_\lambda^*(\Gamma) = E_\lambda(\Gamma)$ so that all samples have the same statistical weight. The probability that the system is in the bound (resp. unbound) state is $p = \langle \delta_{\lambda 1} \rangle$ (resp. $q = 1 - p = \langle \delta_{\lambda 0} \rangle$) and eq. (3) becomes

$$K_d = \frac{1}{V_0} \frac{q}{p}. \quad (5)$$

The determination of K_d is then equivalent to estimating q/p by taking the ratio of the number of unbound and bound steps in the Monte Carlo simulation. How well this estimate converges depends, naturally, on the ‘‘correlation time’’ of the system. If wormhole moves rarely cause the system to switch between bound and unbound states, the correlation time is large and the convergence will be slow.

To make these ideas quantitative, we define the λ -correlation function,

$$C_t = \langle (\lambda_s - p)(\lambda_{s+t} - p) \rangle_s,$$

where λ_s is the value of λ at simulation step s and $\langle \dots \rangle_s$ denotes an average over steps. Figure 4 shows C_t for several different values of V_0 . [From eq. (5), we have $q/p \propto V_0$.] In a Markov chain of length s , the expected number of bound states is ps , while, for $s \rightarrow \infty$, the variance in the number of bound states is $2Ds$, where

$$D = \frac{1}{2} C_0 + \sum_{t>0} C_t = \frac{1}{2} C_0 \tau. \quad (6)$$

This provides us with the definition of the correlation time, τ . From Figure 4, we see that C_t decays approximately exponentially so that the sum converges.




Figure 4. The λ -correlation function C_t . The curves show C_t for $V_0/(0.39 \times 10^{-18} \text{ m}^3) =$ (a) $\frac{1}{10}$, (b) $\frac{1}{3}$, (c) 1, (d) 3, and (e) 10. The ratio q/p varies correspondingly (from approximately $\frac{1}{8}$ to $\frac{100}{8}$). The correlation times are (a) 329, (b) 346, (c) 352, (d) 382, and (e) 415. In determining the correlation times, we limit the sum in eq. (6) to $0 < t \leq 2500$ to avoid including the small but noisy terms for large t .

We can compare the Monte Carlo simulation to a simpler Markov process of independent trials, for example, tossing a coin where the probability of heads is p . In this case, we have $C_{t>0} = 0$ and $D = \frac{1}{2} C_0 = \frac{1}{2} pq$. In the limit $s \rightarrow \infty$, the relative errors in estimating p for the Monte Carlo simulation will be the same as that for s/τ tosses of the coin. We can illustrate this by making five independent simulations of a coin tossing experiment to match the data in Figure 3, for which $p = 0.443$ and $\tau = 352$. The results are given in Figure 5, where we plot h_n/t_n against n where h_n




Figure 5. Estimation of the bias of a coin with $p = 0.443$. Five independent runs are shown. In each run, the coin was tossed $5 \times 10^6/\tau = 14,200$ times, where τ is the correlation time for Figure 4c, and the cumulative value of h_n/t_n is recorded. The expected value, p/q , is shown as a dashed line. The axes have been adjusted to allow the figure to be directly compared with Figure 3.

(resp. t_n) is the number of heads (resp. tails) after n throws. As expected, Figures 3 and 5 both exhibit the same convergence behavior.

Because the distribution of outcomes h_n in the coin tossing experiment is the binomial distribution, the mean and variance for $l_n = \ln(h_n/t_n)$ can be calculated, yielding

$$\langle l_n \rangle = \ln\left(\frac{p}{q}\right) + \frac{1}{2n} \left(\frac{p}{q} - \frac{q}{p}\right) + O(n^{-2}), \quad \langle (l_n - \langle l_n \rangle)^2 \rangle = \frac{1}{npq} + O(n^{-2}).$$

The standard error in the estimate of pK_d from a Monte Carlo run with s steps is found by substituting $n = s/\tau$ and scaling by $\ln 10$ (because pK_d is defined in terms of the common logarithm) to give

$$\Delta pK_d \approx \frac{1}{\ln 10} \sqrt{\frac{\tau}{spq}}. \quad (7)$$

In the case of the simulations shown in Figure 3, we find $\Delta pK_d \approx 0.01$ consistent with the data in the figure. From Figure 4, we see that τ is weakly dependent on p . Thus, for a given s , we minimize the error by choosing $p = q = \frac{1}{2}$. Alternatively, we may wish to minimize the error for a given amount of computational effort. If bound steps are f times more expensive to carry out than unbound steps, the error is minimized by taking $q/p = \sqrt{f}$, that is, $p = 1/(1 + \sqrt{f})$.

Discussion

The important aspects of this method that enable us to compute the binding affinity of a ligand to a protein are: (1) formulation in the extended $[\Gamma, \lambda]$ space, which allows the unbound and bound systems to be treated as a single canonical system, and K_d to be expressed as the ratio of canonical averages, eq. (3); (2) the wormhole move that allows transitions between the bound and unbound systems; and (3) the use of an implicit solvation model that reduces the number of degrees of freedom and so allows portals to be identified.

A method similar to ours is “simulated mutational equilibration,”³² which employs an extended phase space, uses an implicit solvation model, and allows jumps between different systems. However, in place of the wormhole move, this work employed a more restrictive “jumping between wells” move, which limited its applicability to computing the difference in the binding affinities for two enantiomers. In contrast, our wormhole method can be used to directly compute the binding affinity of a ligand with several rotatable bonds to a protein target. This allows us to study a wide range of interesting drug-like ligands.

In the special case of binding rigid molecules (for which we have $E_0^*(\Gamma) = 0$), wormhole Monte Carlo is isomorphic to a grand canonical Monte Carlo simulation.³³ Consider a grand canonical system with a single protein molecule and with ligand molecules at a fixed chemical potential μ ; we make the additional restriction that the ligand–ligand interaction energy is infinite, so that the system can only accommodate 0 or 1 ligand molecule. For the

wormhole simulation, we specify the bound portal to include the full simulation volume with arbitrary orientation; the unbound port is degenerate and corresponds merely to specifying $\lambda = 0$. The wormhole moves from $\lambda = 0$ to 1 and vice versa correspond to particle insertion and deletion in the grand canonical simulation, and we can verify that the acceptance probabilities are the same.

The application described here can be extended by allowing a greater degree of flexibility for the protein and the ligand. This permits the treatment of side-chain rotation and ligand-induced loop movement on the part of the protein and the treatment of flexible rings for the ligand. Many standard Monte Carlo techniques can be used with this method, if appropriate—preferential sampling, early rejection, force bias, etc. Wormhole moves could also be used to treat other discrete, or nearly discrete, transitions. Examples are: the treatment of molecules, such as cyclohexane, which can assume distinct conformations; discrete sets of side-chain rotations in the protein; protonation and tautomerization states for either molecule. In each of these cases, the acceptance probability for the transitions should account for the free energy difference between the discrete states.

Appendix: Generalized Lambert Projection

The Lambert azimuthal equal-area projection projects a point on S^2 to a point on a disk in \mathbb{R}^2 . Here, we will generalize this to arbitrary dimensions; that is, we will find a projection from S^n to a ball in \mathbb{R}^n . (Thus, the circle S^1 is projected into a line segment; the surface of a sphere S^2 is projected into a disk; S^3 is projected into a sphere; etc.) The projection is azimuthal, so that directions from the pole are preserved in the projected space.

The sequence S^n can be defined recursively by

$$S^0 = [\pm 1],$$

$$S^n = [\cos \theta, \sin \theta S^{n-1}],$$

where θ is the colatitude and $0 \leq \theta \leq \pi$. The area of S^n lying between θ and $\theta + d\theta$ is, therefore

$$dA = a_{n-1} \sin^{n-1} \theta d\theta,$$

where a_n is the area of S^n . We project that portion of S^n with colatitude in $[0, \theta]$ to a ball in \mathbb{R}^n of radius t . Equating the measures (area on S^n and volume in \mathbb{R}^n), we obtain

$$v_n t^n = a_{n-1} \int_0^\theta \sin^{n-1} \theta' d\theta'$$

where v_n is the volume of a unit ball in \mathbb{R}^n . Using the relation,

$$a_{n-1} = \left. \frac{d(v_n t^n)}{dt} \right|_{t=1} = n v_n,$$

we obtain

$$t = \left(n \int_0^\theta \sin^{n-1} \theta' d\theta' \right)^{1/n}. \quad (8)$$

The general “equal-area” projection is given by the mapping

$$[\cos \theta, \sin \theta S^{n-1}] \rightarrow \mathbf{t},$$

where $\mathbf{t} \in \mathbb{R}^n$, $\hat{\mathbf{t}}$ is given by the point on S^{n-1} , and t is given by eq. (8). Some special cases of eq. (8) are

$$t = \begin{cases} \theta, & \text{for } n = 1, \\ 2 \sin \frac{1}{2} \theta, & \text{for } n = 2, \\ \left[\frac{3}{4} (2\theta - \sin 2\theta) \right]^{1/3}, & \text{for } n = 3, \\ 2 \sin \frac{1}{2} \theta \left[\frac{1}{3} (1 + 2 \cos^2 \frac{1}{2} \theta) \right]^{1/4}, & \text{for } n = 4. \end{cases}$$

The case $n = 1$ corresponds to unwrapping a circle onto a line; and $n = 2$ gives the Lambert azimuthal equal-area projection. We are interested in the case $n = 3$ as a way of mapping orientations from unit quaternion space to \mathbb{R}^3 . A quaternion $q = [\cos \theta, \sin \theta \hat{\mathbf{u}}]$ represents a rotation of $\psi = 2\theta$ about $\hat{\mathbf{u}}$. Noting that q and $-q$ represent the same rotation, we may make the restriction $0 \leq \theta \leq \frac{1}{2} \pi$. We choose, therefore, to map this hemisphere of S^3 onto \mathbf{u} in the unit ball, by making the substitutions $\theta = \frac{1}{2} \psi$ and $t = \left(\frac{3}{4} \pi \right)^{1/3} u$ to give

$$u = \left(\frac{\psi - \sin \psi}{\pi} \right)^{1/3}.$$

(The other hemisphere, corresponding to $\frac{1}{2} \pi < \theta \leq \pi$, maps onto the shell $1 < u \leq 2^{1/3}$.) We call this \mathbf{u} representation of orientations “turn space.”

References

- Mezei, M.; Beveridge, D. L. *Ann NY Acad Sci* 1986, 482, 1.
- Kollman, P. A. *Chem Rev* 1993, 93, 2395.
- Landau, L. D.; Lifshitz, E. M. *Course of Theoretical Physics*; Pergamon Press: New York, 1969, Vol. 5, 2nd ed.
- Bennett, C. H. *J Comp Phys* 1976, 22, 245.
- Luo, H.; Sharp, K. *Proc Nat Acad Sci* 2002, 99, 10399.
- Gilson, M. K.; Given, J. A.; Bush, B. L.; McCammon, J. A. *Biophys J* 1997, 72, 1047.
- Metropolis, N.; Rosenbluth, A. W.; Rosenbluth, M. N.; Teller, A. H.; Teller, E. *J Chem Phys* 1953, 21, 1087.
- Lyubartsev, A. P.; Martsinovski, A. A.; Shevkunov, S. V.; Vorontsov-Velyaminov, P. N. *J Chem Phys* 1992, 96, 1776.
- Tidor, B. *J Phys Chem* 1993, 97, 1069.
- Kong, X.; Brooks, C. L., III. *J Chem Phys* 1996, 105, 2414.
- (a) Morris, M. S.; Thorne, K. S.; Yurtsever, U. *Phys Rev Lett* 1998, 61, 1446; (b) Houdayer, J. *J Chem Phys* 2002, 116, 1783.
- Miller, M. A.; Reinhardt, W. P. *J Chem Phys* 2000, 113, 7035.
- Zhu, Z.; Tuckerman, M. E.; Samuelson, S. O.; Martyna, G. J. *Phys Rev Lett* 2002, 88, 100201.
- Voter, A. F. *J Chem Phys* 1985, 82, 1890.
- Senderowitz, H.; Guarnieri, F.; Still, W. C. *J Am Chem Soc* 1995, 117, 8211.
- Golub, G. H.; Van Loan, C. F. *Matrix Computations*; Johns Hopkins University Press: Baltimore, MD, 1996, 3rd ed.
- Hamilton, W. R. *Proc R Irish Acad* 1847, 3, 1; URL <http://www.maths.tcd.ie/pub/HistMath/People/Hamilton/Quatern2/>.
- Mardia, K. V.; Jupp, P. E. *Directional Statistics*, J. Wiley & Sons: New York, 1999.
- Head, M. S.; Given, J. A.; Gilson, M. K. *J Phys Chem* 1997, 101, 1609.
- Dempster, A. P.; Laird, N. M.; Rubin, D. B. *J R Stat Soc* 1997, 39, 1.
- Verbeek, J. J.; Vlassis, N.; Kröse, B. *Neural Comput* 2003, 15, 469.
- Katz, B. A.; Finer-Moore, J.; Mortezaei, R.; Rich, D. H.; Stroud, R. M. *Biochemistry* 1995, 34, 8264; URL <http://www.rcsb.org/pdb/cgi/expire.cgi?pdbId=1bty>.
- Case, D. A.; Pearlman, D. A.; Caldwell, J. W.; Cheatham, T. E., III; Wang, J.; Ross, W. S.; Simmerling, C. L.; Darden, T. A.; Mertz, K. M.; Stanton, R. V.; et al. *Amber 7*; University of California: San Francisco, 2002; URL <http://amber.scripps.edu/doc7/>.
- Cornell, W. D.; Cieplak, P.; Bayly, C. I.; Gould, I. R.; Merz, K. M., Jr.; Ferguson, D. M.; Spellmeyer, D. C.; Fox, T.; Caldwell, J. W.; Kollman, P. A. *J Am Chem Soc* 1995, 117, 5179.
- Bayly, C. I.; Cieplak, P.; Cornell, W. D.; Kollman, P. A. *J Phys Chem* 1993, 97, 10269.
- Still, W. C.; Tempczyk, A.; Hawley, R. C.; Hendrickson, T. *J Am Chem Soc* 1990, 112, 6127.
- Hawkins, G. D.; Cramer, C. J.; Truhlar, D. G. *Chem Phys Lett* 1995, 246, 122.
- Tsui, V.; Case, D. A. *J Am Chem Soc* 2000, 122, 2489.
- Frisch, M. J.; Trucks, G. W.; Schlegel, H. B.; Scuseria, G. E.; Robb, M. A.; Cheeseman, J. R.; Zakrzewski, V. G.; Montgomery, J. A., Jr.; Stratmann, R. E.; Burant, J. C., et al. *Gaussian 98*, Revision A.9; Gaussian, Inc.: Pittsburg, PA, 1998; URL <http://www.gaussian.com>.
- Mares-Guia, M.; Shaw, E. *J Biol Chem* 1965, 240, 1579.
- Schwarzl, S. M.; Tschopp, T. B.; Smith, J. C.; Fischer, S. *J Comp Chem* 2002, 23, 1143.
- Senderowitz, H.; McDonald, D. Q.; Still, W. C. *J Org Chem* 1997, 62, 9123.
- Adams, D. J. *Mol Phys* 1975, 29, 307.