Free Energy Calculations in Biological Systems. How Useful Are They in Practice?

Christophe Chipot

Equipe de dynamique des assemblages membranaires, UMR CNRS/UHP 7565,
Institut nancéien de chimie moléculaire, Université Henri Poincaré, BP 239,
54506 Vandœuvre–lès–Nancy cedex, France

Abstract. Applications of molecular simulations targeted at the estimation of free energies are reviewed, with a glimpse into their promising future. The methodological milestones paving the road of free energy calculations are summarized, in particular free energy perturbation and thermodynamic integration, in the framework of constrained or unconstrained molecular dynamics. The continuing difficulties encountered when attempting to obtain accurate estimates are discussed with an emphasis on the usefulness of large–scale numerical simulations in non–academic environments, like the world of the pharmaceutical industry. Applications of the free energy arsenal of methods is illustrated through a variety of biologically relevant problems, among which the prediction of protein–ligand binding constants, the determination of membrane–water partition coefficients of small, pharmacologically active compounds — in connection with the blood–brain barrier, the folding of a short hydrophobic peptide, and the association of transmembrane α–helical domains, in line with the "two–stage" model of membrane protein folding. Current strategies for improving the reliability of free energy calculations, while making them somewhat more affordable, and, therefore, more compatible with the constraints of an industrial environment, are outlined.

Key words: Free energy calculations, molecular dynamics simulations, drug design, protein folding, protein recognition and association

1 Introduction

To understand fully the vast majority of chemical and biochemical processes, a close examination of the underlying free energy behavior is often necessary [1]. Such is the case, for instance, of protein–ligand binding constants and membrane–water partition coefficients, that are of paramount important in the emerging field of *de novo*, rational drug design, and cannot be predicted reliably and accurately without the knowledge of the associated free energy changes. The ability to determine *a priori* these physical constants with a reasonable level of accuracy, by means of statistical simulations, is now within reach. Developments on both the software and the hardware fronts have contributed to bring free energy calculations at the level of similarly

robust and well–characterized modeling tools, while widening their field of applications. Yet, in spite of the tremendous progress accomplished since the first published calculations some twenty years ago [2–5], the accurate estimation of free energy changes in large, physically and biologically realistic molecular assemblies still constitutes a challenge for modern theoretical chemistry. Taking advantage of massively parallel architectures, cost–effective, "state–of–the–art" free energy calculations can provide a convincing answer to help rationalizing experimental observations, and, in some instances, play a predictive role in the development of new leads for a specific target.

In the first section of this chapter, the methodological background of free energy calculations is developed, focusing on the methods that are currently utilized to determine free energy differences. Next, four biologically relevant applications are presented, corresponding to distinct facets of free energy simulations. The first one delves into the use of these calculations in *de novo* drug design, through the estimation of protein–ligand binding affinities and water–membrane partition coefficients. The somewhat more challenging application of molecular dynamics (MD) simulations and free energy methods to validate the three–dimensional structure of a G protein–coupled receptor (GPCR) is shown next. Understanding of the intricate physical phenomena that drive protein folding by means of free energy calculations is also reported here, followed by an investigation of the reversible association of transmembrane (TM) α–helices in a membrane mimetic. Conclusions on the role played by free energy calculations in the molecular modeling community are drawn with a prospective look into their future.

2 Methodological Background

In the canonical, (N, V, T), ensemble, the Helmholtz free energy is defined by [6]:

$$A = -\frac{1}{\beta} \ln Q_{\mathrm{NVT}} \tag{1}$$

$\beta = 1/k_B T$, where k_B is the Boltzmann constant and T is the temperature of the N–particle system. Q_{NVT} is its $6N$–dimensional partition function:

$$Q_{\mathrm{NVT}} = \frac{1}{h^{3N} N!} \int\int \exp\left[-\beta \mathcal{H}(\mathbf{x}, \mathbf{p}_x)\right] \, \mathrm{d}\mathbf{x} \, \mathrm{d}\mathbf{p}_x \tag{2}$$

where $\mathcal{H}(\mathbf{x}, \mathbf{p}_x)$ is the classical Hamiltonian describing the system. In (2), integration is carried out over all atomic coordinates, $\{\mathbf{x}\}$, and momenta, $\{\mathbf{p}_x\}$. The normalization factor reflects the measure of the volume of the phase space through the Planck constant, h, and the indistinguishable nature of the particles, embodied in the factorial term, $N!$. In essence, the canonical partition function constitutes the corner stone of the statistical mechanical description of the ensemble of particles. From a phenomenological perspective, it may be viewed as a measure of the thermodynamic states accessible to the system in terms of spatial coordinates and momenta. It can be further restated in terms of energies:

$$Q_{\mathrm{NVT}} = \frac{1}{h^{3N} N!} \int \varrho[\mathcal{H}(\mathbf{x}, \mathbf{p}_x)] \, \exp\left[-\beta \mathcal{H}(\mathbf{x}, \mathbf{p}_x)\right] \, \mathrm{d}\mathcal{H}(\mathbf{x}, \mathbf{p}_x) \tag{3}$$

where $\varrho[\mathcal{H}(\mathbf{x}, \mathbf{p}_x)]$ is the so–called density of states accessible to the system of interest.

The definition of the partition function may be utilized to introduce the concept of probability distribution to find the system in the unique microscopic state characterized by positions $\{\mathbf{x}\}$ and momenta $\{\mathbf{p}_x\}$:

$$\mathcal{P}(\mathbf{x}, \mathbf{p}_x) = \frac{1}{h^{3N} N!} \, \frac{1}{Q_{\mathrm{NVT}}} \, \exp\left[-\beta \mathcal{H}(\mathbf{x}, \mathbf{p}_x)\right] \tag{4}$$

A logical consequence of this expression is that low–energy regions of the phase space will be sampled predominantly, according to their respective Boltzmann weight [7].

2.1 Free Energy Perturbation

Returning to the original definition (1) of the free energy, and using the identity:

$$\int \int \exp\left[+\beta \mathcal{H}(\mathbf{x}, \mathbf{p}_x)\right] \, \exp\left[-\beta \mathcal{H}(\mathbf{x}, \mathbf{p}_x)\right] \, \mathrm{d}\mathbf{x} \, \mathrm{d}\mathbf{p}_x = h^{3N} N! \tag{5}$$

it follows that:

$$\begin{aligned} A &= -\frac{1}{\beta} \, \ln \frac{1}{h^{3N} N!} \int \int \exp\left[-\beta \mathcal{H}(\mathbf{x}, \mathbf{p}_x)\right] \, \mathrm{d}\mathbf{x} \, \mathrm{d}\mathbf{p}_x \\ &= +\frac{1}{\beta} \, \ln \int \int \exp\left[+\beta \mathcal{H}(\mathbf{x}, \mathbf{p}_x)\right] \, \mathcal{P}(\mathbf{x}, \mathbf{p}_x) \, \mathrm{d}\mathbf{x} \, \mathrm{d}\mathbf{p}_x \\ &= +\frac{1}{\beta} \, \ln \left\langle \exp\left[+\beta \mathcal{H}(\mathbf{x}, \mathbf{p}_x)\right] \right\rangle \end{aligned} \tag{6}$$

This expression illuminates the fast growth of $\exp\left[+\beta \mathcal{H}(\mathbf{x}, \mathbf{p}_x)\right]$ with the total energy, $\mathcal{H}(\mathbf{x}, \mathbf{p}_x)$, of the system. It should, therefore, be expected that the weight of the high–energy regions of phase space be significant when evaluating the integral. Yet, as hinted by (4), in simulations of finite length, sampling of these regions is likely to be insufficient to guarantee a correct estimate of A. In most instances, evaluation of accurate absolute free energies from statistical simulations is not possible. The latter may, however, give access to free energy differences between two well–delineated thermodynamic states, provided that a reaction coordinate can be defined to characterize the pathway that connects these two states. In this context, the Hamiltonian, $\mathcal{H}(\mathbf{x}, \mathbf{p}_x)$, describing the transformation is made a function of the reaction coordinate, or "coupling parameter", λ [8]. Conventionally, λ varies between 0 and 1 when the system goes from the initial state, a, to the final state, b,

characterized, respectively, by the Hamiltonians $\mathcal{H}(\mathbf{x}, \mathbf{p}_x; \lambda_a) = \mathcal{H}(\mathbf{x}, \mathbf{p}_x; \lambda = 0)$ and $\mathcal{H}(\mathbf{x}, \mathbf{p}_x; \lambda_b) = \mathcal{H}(\mathbf{x}, \mathbf{p}_x; \lambda = 1)$. In practice, λ can correspond to a variety of reaction coordinates, ranging from a simple distance to determine a potential of mean force (PMF) to non–bonded parameters in the so–called "alchemical transformations" or *in silico* point mutations [4, 5].

Within this framework, the canonical partition function defined in (2) now depends explicitly on the coupling parameter, and so does the free energy:

$$\Delta A_{a\to b} = A(\lambda_b) - A(\lambda_a) = -\frac{1}{\beta} \ln \frac{Q_{\mathrm{NVT}}(\lambda_b)}{Q_{\mathrm{NVT}}(\lambda_a)} \tag{7}$$

Combining the above with the definition of the partition function, and introducing the identity (5), it follows that:

$$\boxed{\Delta A_{a\to b} = -\frac{1}{\beta} \ln \left\langle \exp\left\{-\beta\left[\mathcal{H}(\mathbf{x}, \mathbf{p}_x; \lambda_b) - \mathcal{H}(\mathbf{x}, \mathbf{p}_x; \lambda_a)\right]\right\}\right\rangle_{\lambda_a}} \tag{8}$$

Here, $\langle \cdots \rangle_{\lambda_a}$ denotes an ensemble average over configurations representative of the initial state, a. Validity of perturbation formula (8) only holds for small changes between the initial state, a, and the final state, b, of the transformation [9]. At this stage, the condition of *small changes* should be clarified, as it has often been misconstrued in the past. It does not imply that the free energies characteristic of a and b be sufficiently close, but rather that the corresponding configurational ensembles overlap appropriately to guarantee the desired accuracy [10, 11]. In other words, it is expected that the density of states, $\varrho[\mathcal{H}(\mathbf{x}, \mathbf{p}_x)]$, describing the transformation from a to b be narrow enough — *viz.* typically on the order of $1 \,/\, \beta$, to ascertain that, when multiplied by the exponential term of (3), the resulting distribution be located in a region where ample statistical data have been collected. In most circumstances, however, single–step transformations between rather orthogonal states are unlikely to fulfill this requirement. To circumvent this difficulty, the reaction pathway is broken down into a number of physically meaningless intermediate states connecting a to b, so that between any two contiguous states, the condition of overlapping ensembles is satisfied [12]. The interval separating these intermediate states, which corresponds to selected fixed values of the coupling parameter, λ, is often referred to as "window". It should be reminded that the vocabulary *window* adopted in perturbation theory is distinct from that utilized in "umbrella sampling" (US) simulations [13], where it denotes a range of values taken by the reaction coordinate. For a series of a N intermediate states, the total free energy change for the transformation from a to b is expressed as a sum of $N - 1$ free energy differences [12]:

$$\Delta A_{a\to b} = -\frac{1}{\beta} \sum_{k=1}^{N-1} \ln \left\langle \exp\left\{-\beta\left[\mathcal{H}(\mathbf{x}, \mathbf{p}_x; \lambda_{k+1}) - \mathcal{H}(\mathbf{x}, \mathbf{p}_x; \lambda_k)\right]\right\}\right\rangle_{\lambda_k} \tag{9}$$

Assessing the ideal number of intermediate states, N, between a and b evidently depends upon the nature of the system that undergoes the transformation. The condition of overlapping ensembles should be kept in mind when setting N, remembering that

the choice of $\delta\lambda = \lambda_{k+1} - \lambda_k$ ought to correspond to a perturbation of the system. A natural choice consists in using a number of windows that guarantees a reasonably similar free energy change between contiguous intermediate states. The consequence of this choice is that the width of the consecutive windows connecting a to b may be different.

Performing "alchemical transformations" calls for the definition of topologies that describe the initial and the final states of the mutation. To this end, a single– or a dual–topology paradigm [17] can be employed, as shown in Fig. 1. Each approach has its advantages and inherent drawbacks. To circumvent numerical instabilities in the trajectory caused by the concurrent scaling of the charges, the van der Waals and selected internal parameters, transformation of electrostatic and non–electrostatic terms are usually decoupled in the single–topology approach, thus, giving access to the corresponding free energy contributions. This scheme requires two distinct simulations, when the dual–topology paradigm only involves one. The latter, however, does not provide much information about the contributions that drive the global free energy change — albeit contributions should be interpreted with great care, because, in contrast with the net free energy, which is a state function, the former are path–dependent. The dual–topology approach has been recognized to be generally more sensitive to so–called "end–point catastrophes", when $\lambda \rightarrow 0$ or 1, than the

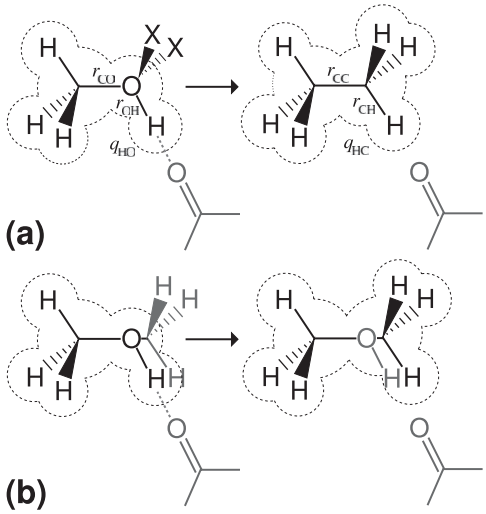

Figure 1. Difference between the single– (**a**) and the dual–topology (**b**) paradigms illustrated in the case of the methanol to ethane "alchemical transformation". In the single–topology approach, a common topology for the initial and the final states of the transformation is designed. Non–bonded parameters are scaled as λ varies from 0 to 1. Modification of the oxygen atom into a carbon one imposes that the chemical bond be shrunk, using, for instance, a PMF–type calculation [14–16]. In the dual–topology paradigm, the initial and the final states coexist , yet without "seeing" each other. The interaction energy of these topologies with their environment is scaled as λ goes from 0 to 1

single–topology paradigm. A number of schemes have been devised to circumvent this problem, among which the use of windows of decreasing width as λ tends towards 0 or 1. Introduction of a soft–core potential [18] to eliminate the singularities at 0 or 1 perhaps constitutes the most elegant method proposed hitherto.

2.2 Thermodynamic Integration

Closely related to the free energy perturbation (FEP) expression, thermodynamic integration (TI) restates the free energy difference between state a and state b as a finite difference [8, 19]:

$$\Delta A_{a\rightarrow b} = A(\lambda_b) - A(\lambda_a) = \int_{\lambda_a}^{\lambda_b} \frac{\mathrm{d}A(\lambda)}{\mathrm{d}\lambda}\,\mathrm{d}\lambda \tag{10}$$

Combining with the definition of the canonical partition function (2), it follows that:

$$\frac{\mathrm{d}A(\lambda)}{\mathrm{d}\lambda} = \frac{\displaystyle\int \frac{\partial\mathcal{H}(\mathbf{x},\mathbf{p}_x;\lambda)}{\partial\lambda}\,\exp -\beta\mathcal{H}(\mathbf{x},\mathbf{p}_x;\lambda)\,\mathrm{d}\mathbf{x}\,\mathrm{d}\mathbf{p}_x}{\displaystyle\int \exp -\beta\mathcal{H}(\mathbf{x},\mathbf{p}_x;\lambda)\,\mathrm{d}\mathbf{x}\,\mathrm{d}\mathbf{p}_x} \tag{11}$$

Consequently, the integrand can be written as an ensemble average:

$$\boxed{\Delta A_{a\rightarrow b} = \int_{\lambda_a}^{\lambda_b} \left\langle \frac{\partial\mathcal{H}(\mathbf{x},\mathbf{p}_x;\lambda)}{\partial\lambda} \right\rangle_{\lambda} \mathrm{d}\lambda} \tag{12}$$

In sharp contrast with the FEP method, the criterion of convergence here is the appropriate smoothness of $A(\lambda)$. Interestingly enough, assuming that the variation of the kinetic energy between states a and b can be neglected, the derivative of $\Delta A_{a\rightarrow b}$ with respect to some reaction coordinate, ξ , is equal to, $-\left\langle F_\xi \right\rangle_\xi$, the average of the force exerted along ξ , hence, the concept of PMF [20].

2.3 Unconstrained Molecular Dynamics and Average Forces

Generalization of the classical definition of a PMF, $w(r)$, based on the pair correlation function, $g(r)$, is not straightforward. For this reason, the free energy as a function of reaction coordinate ξ will be expressed as:

$$A(\xi) = -\frac{1}{\beta}\ln\mathcal{P}(\xi) + A_0 \tag{13}$$

where $\mathcal{P}(\xi)$ is the probability distribution to find the system at a given value, ξ, along that reaction coordinate:

$$\mathcal{P}(\xi) = \int \delta[\xi - \xi(\mathbf{x})]\;\exp[-\beta\mathcal{H}(\mathbf{x},\mathbf{p}_x)]\,\mathrm{d}\mathbf{x}\,\mathrm{d}\mathbf{p}_x \tag{14}$$

Equation (13) corresponds to the classic definition of the free energy in methods like US, in which external biasing potentials are included to ensure a uniform distribution $\mathcal{P}(\xi)$. To improve sampling efficiency, the complete reaction pathway is broken down into "windows", or ranges of ξ, wherein individual free energy profiles are determined. The latter are subsequently pasted together using, for instance, the self–consistent weighted histogram analysis method (WHAM) [21].

For a number of years, the first derivative of the free energy with respect to the reaction coordinate has been written as [22]:

$$\frac{\mathrm{d}A(\xi)}{\mathrm{d}\xi} = \left\langle \frac{\partial \mathcal{V}(\mathbf{x})}{\partial \xi} \right\rangle_{\xi} \tag{15}$$

This description is erroneous because ξ and $\{\mathbf{x}\}$ are evidently not independent variables [23]. Furthermore, it assumes that kinetic contributions can be safely omitted. This may not always be necessarily the case. For these reasons, a transformation of the metric is required, so that:

$$\mathcal{P}(\xi) = \int |J| \, \exp[-\beta \mathcal{V}(\mathbf{q};\xi)] \, \mathrm{d}\mathbf{q} \int \exp[-\beta \mathcal{T}(\mathbf{p}_x)] \mathrm{d}\mathbf{p}_x \tag{16}$$

Introduction in the first derivative of probability $\mathcal{P}(\xi)$, it follows that the kinetic contribution vanishes in $\mathrm{d}A(\xi)/\mathrm{d}\xi$:

$$\frac{\mathrm{d}A(\xi)}{\mathrm{d}\xi} = -\frac{1}{\beta} \frac{1}{\mathcal{P}(\xi)} \int \exp[-\beta \mathcal{V}(\mathbf{q};\xi^*)] \, \delta(\xi^* - \xi) \, \times$$

$$\left\{ -\beta |J| \frac{\partial \mathcal{V}(\mathbf{q};\xi^*)}{\partial \xi} + \frac{\partial |J|}{\partial \xi} \right\} \mathrm{d}\mathbf{q} \, \mathrm{d}\xi^* \tag{17}$$

After back transformation into Cartesian coordinates, the derivative of the free energy with respect to ξ can be expressed as a sum of configurational averages at constant ξ [24]:

$$\boxed{\frac{\mathrm{d}A(\xi)}{\mathrm{d}\xi} = \left\langle \frac{\partial \mathcal{V}(\mathbf{x})}{\partial \xi} \right\rangle_{\xi} - \frac{1}{\beta} \left\langle \frac{\partial \ln |J|}{\partial \xi} \right\rangle_{\xi} = -\langle F_{\xi} \rangle_{\xi}} \tag{18}$$

In this approach, only $\langle F_{\xi} \rangle_{\xi}$ is the physically meaningful quantity, unlike the instantaneous components, F_{ξ}, from which it is evaluated. Moreover, it should be clearly understood that neither F_{ξ} nor $-\partial \mathcal{H}(\mathbf{x}, \mathbf{p}_x)/\partial \xi$ are fully defined by the sole choice of the reaction coordinate.

In practice, F_{ξ} is accumulated in bins of finite size $\delta\xi$ and provides an estimate of $\mathrm{d}A(\xi)/\mathrm{d}\xi$. After a predefined number of observables are accrued, the adaptive biasing force (ABF) [25, 26] is applied along the reaction coordinate:

$$\mathbf{F}^{\mathrm{ABF}} = \boldsymbol{\nabla} \widetilde{A} = -\langle F_{\xi} \rangle_{\xi} \, \boldsymbol{\nabla} \xi \tag{19}$$

As sampling proceeds, $\boldsymbol{\nabla}\widetilde{A}$ is progressively refined. Evolution of the system along ξ is governed mainly by its self–diffusion properties. It is apparent from the present description that this method is significantly more effective than US or its variants, because no *a priori* knowledge of the free energy hypersurface is required to define the necessary biasing potentials that will guarantee uniform sampling along ξ. The latter can easily become intricate in the case of qualitatively new problems, in which variation of the free energy behavior cannot be guessed with the appropriate accuracy. Often misconstrued, it should be emphasized that, whereas ABF undoubtedly improves sampling dramatically along the reaction coordinate, efficiency suffers, like any other free energy method, from orthogonal degrees of freedom in the slow manifolds.

2.4 When is Enough Sampling Really Enough?

When is the trajectory long enough to assume safely that the results are converged? is a recurrent question asked by modelers performing free energy calculations. Assessing the convergence properties and the error associated to a free energy calculation often turns out to be daunting task. Sources of errors likely to be at play are diverse, and, hence, will modulate the results differently. The choice of the force field parameters undoubtedly affects the results of the simulation, but this contribution can be largely concealed by the statistical error arising from insufficient sampling. Paradoxically, exceedingly short free energy calculations employing inadequate non–bonded parameters may, nonetheless, yield the correct answer [14]. Under the hypothetical assumption of an optimally designed potential energy function, quasi non–ergodicity scenarios constitute a common pitfall towards fully converged simulations.

Appreciation of the statistical error has been devised following different schemes. Historically, the free energy changes for the $\lambda \rightarrow \lambda + \delta\lambda$ and the $\lambda \rightarrow \lambda - \delta\lambda$ perturbations were computed simultaneously to provide the hysteresis between the forward and the reverse transformations. In practice, it can be shown that when $\delta\lambda$ is sufficiently small, the hysteresis of such "double–wide sampling" simulation [27] becomes negligible, irrespective of the amount of sampling generated in each window — as would be the case in a "slow–growth" calculation [28]. A somewhat less arguable point of view consists in performing the transformation in the forward, $a \rightarrow b$, and in the reverse, $b \rightarrow a$, directions. Micro–reversibility imposes that, in principle, $\Delta A_{b\rightarrow a} = -\Delta A_{a\rightarrow b}$ — see for instance Fig. 2. Unfortunately, forward and reverse transformations do not necessarily share the same convergence properties. Case in point, the insertion and deletion of a particle [29]: Whereas the former simulation converges rapidly towards the expected excess chemical potential, the latter never does. This shortcoming can be ascribed to the fact that configurations in which a cavity does not exist where a real atom is present are never sampled. In terms of density of states, this scenario would translate into ϱ_a embracing ϱ_b entirely, thereby ensuring a proper convergence of the forward simulation, whereas the same cannot be said for the reciprocal, reverse transformation. Estimation of errors based on forward and reverse simulations should, therefore, be considered with great care.

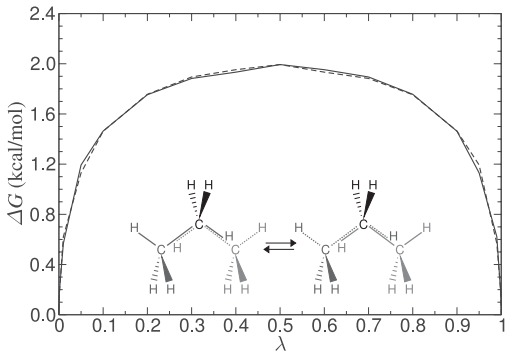

Figure 2. Free energy change characterizing the "zero–sum", for- ce–field–independent mutation of ethane into ethane in water (insert). Forward (*solid line*) and reverse (dashed line) profiles were obtained by dividing the reaction pathway into 40 windows, in which 16 ps of equilibration and 32 ps of data collection were sampled. The dual–topology paradigm was used to define a pseudo–propane molecule

In fact, appropriate combination of the two can be used profitably to improve the accuracy of free energy calculations [10].

In FEP, convergence may be probed by monitoring the time–evolution of the ensemble average (6). This is, however, a necessary, but not sufficient condition for convergence, because apparent plateaus of the ensemble average often conceal anomalous overlap of the density of states characterizing the initial, ϱ_a , and the final, ϱ_b , states [10, 11]. The latter should be the key–criterion to ascertain the local convergence of the simulation for those degrees of freedom that are effectively sampled. Statistical errors in FEP calculations may be estimated by means of a first–order expansion of the free energy:

$$\Delta A = -\frac{1}{\beta}\left\{\ln\left\langle\exp\left[-\beta\Delta\mathcal{V}(\mathbf{x};\lambda)\right]\right\rangle_\lambda \pm \frac{\delta\varepsilon}{\left\langle\exp\left[-\beta\Delta\mathcal{V}(\mathbf{x};\lambda)\right]\right\rangle_\lambda}\right\} \tag{20}$$

Here, $\delta\varepsilon$ is the statistical error on the ensemble average, $\left\langle\exp\left[-\beta\Delta\mathcal{V}(\mathbf{x};\lambda)\right]\right\rangle_\lambda$, defined as:

$$\delta\varepsilon^2 = \frac{1+2\tau}{N}\left\{\left\langle\exp\left[-2\beta\Delta\mathcal{V}(\mathbf{x};\lambda)\right]\right\rangle_\lambda - \left\langle\exp\left[-\beta\Delta\mathcal{V}(\mathbf{x};\lambda)\right]\right\rangle_\lambda^2\right\} \tag{21}$$

N is the number of samples accrued in the FEP calculation, and $(1 + 2\tau)$ is the sampling ratio of the latter [30].

Commonly, in US simulations, convergence is probed by verifying two criteria: (i) convergence of individual windows — the statistical error can be measured through a block–averaging over sub–runs — and (ii) appropriate overlap of the free energy profiles between adjacent windows — which gives rise to a systematic error. The latter can be evaluated by minimizing the difference in the curvature of two consecutive free energy profiles in the region where they overlap.

In the idealistic cases where a thermodynamic cycle can be defined — e.g. investigation of the conformational equilibrium of a short peptide through the $\alpha_R \rightarrow C_{7\mathrm{ax}} \rightarrow \alpha_L \rightarrow \beta' \equiv (\beta, C_5, C_{7eq}) \rightarrow \alpha_R$ successive transformations — closure of the latter imposes that the sum of individual free energy contributions sum up to zero [31]. In principle, any deviation from this target should provide a valuable guidance to improve sampling efficiency. In practice, discrimination of the faulty transformation, or transformations, is cumbersome on account of possible mutual compensation or cancellation of errors.

As has been commented on previously, visual inspection of ϱ_a and ϱ_b indicates whether the free energy calculation has converged [10, 11]. Deficiencies in the overlap of the two distributions is also suggestive of possible errors, but it should be kept in mind that approximations like (20) only reflect the *statistical precision* of the computation, and evidently do not account for fluctuations in the system occurring over long time scales. In sharp contrast, the *statistical accuracy* is expected to yield a more faithful picture of the degrees of freedom that have been actually sampled. The safest route to estimate this quantity consists in performing the same free energy calculation, starting from different regions of the phase space — *viz.* the error is defined as the root mean square deviation over the different simulations [32]. Semantically speaking, the error measured from one individual run yields the statistical precision of the free energy calculation, whereas that derived from the ensemble of simulations provides its statistical accuracy.

3 Free Energy Calculations and Drug Design

One of the grand challenges of free energy calculations is their ability to play a predictive role in ranking ligands of potential pharmaceutical interest, agonist or antagonist, according to their relative affinity towards a given protein. The usefulness of such numerical simulations outside an academic environment can be assessed by answering the following question: Can free energy calculations provide a convincing answer faster than experiments are carried out in an industrial setting? Early encouraging results had triggered much excitement in the community, opening new vistas for *de novo*, rational drug design. They were, however, subsequently shattered when it was realized that accurate free energies would require considerably more computational effort than was appreciated hitherto. Beyond the fundamental need of appropriate sampling to yield converged ensemble averages — today, easily achievable in some favorable cases by means of inexpensive clusters of personal computers — the necessity of well–parameterized potential energy functions, suitable for non–peptide ligands, rapidly turns out to constitute a critical bottleneck for the routine use of free energy calculations in the pharmaceutical industry. Closely related to the parametrization of the force field, setting up free energy calculations — i.e. defining the alternate topology of the mutated moieties and the initial set of coordinates — is sufficiently time–consuming to be incompatible with the high–throughput requirements of industrial environments. In essence, this is where the paradox of free energy calculations lies: They have not yet come of age to be considered as black box

routine jobs, but should evolve in that direction to become part of the arsenal of computational tools available to the pharmaceutical industry. Selection of potent ligand candidates in large data bases containing several millions of real or virtual molecules, employing a screening funnel that involves increasingly complex searching tools, from crude geometrical recognition to more sophisticated flexible molecular docking, offers new prospects for *de novo* drug design. In this pipeline of screening methods, free energy calculations should evidently be positioned at the very end, i.e. at the level of the optimization loop aimed at a limited number of ligands, which also checks adsorption–metabolism/toxicology (ADME/TOX) properties. Ideally, as selection in the funnel proceeds, the computational effort should remain constant — *viz.* the amount of CPU time necessary to perform free energy calculations on a few candidates is equivalent to that involved in the rough geometrical recognition over a number of molecules five to six orders of magnitude larger.

Perhaps the key to the generalized and routine use of free energy calculations for molecular systems of pharmacological relevance is the demonstration that this methodology can be applied fruitfully to problems that clearly go beyond the typical scope of an academic environment. Collaborative projects with the pharmaceutical industry provide such a framework. In the context of the search for therapeutic agents targeted at osteoporosis and other bone–related diseases, free energy calculations have been applied to complexes formed by the multi–domain protein $^{\mathrm{pp60}}$src kinase associated to non–peptide inhibitors. $^{\mathrm{pp60}}$src kinase is involved in signal transduction pathways and is implicated in osteoclast–mediated bone resorption [33]. Of particular interest, its SH2 domain, a common recognition motif of highly conserved protein sequence, binds preferentially phosphotyrosine (pY)–containing peptides. In most circumstances, the latter adopt an extended conformation mimicking a two–pronged plug that interacts at two distinct anchoring sites of the protein — i.e. the hydrophilic phosphotyrosine pocket and the hydrophobic pocket — separated by a flat surface [34, 35]. For instance, the prototypical tetrapeptide pYEEI — a sequence found on the PDGF receptor upon activation, appears to recognize the src SH2 domain with an appropriate specificity.

In silico point mutations have been performed on a series of non–peptide inhibitors — see Fig. 3, using the FEP methodology in conjunction with the dual–topology paradigm. All simulations have been carried in the isothermal–isobaric ensemble, using the program NAMD [36, 37]. The temperature and the pressure were fixed at 300 K and 1 atm, respectively, employing Langevin dynamics and the Langevin piston. To avoid possible end–point catastrophes, 33 windows of uneven width, $\delta\lambda$, were utilized to scale the interaction of the mutated moieties with their environment. The total length of each trajectory is equal to 1 ns, in the free and in the bound states. Forward and reverse simulations were run to estimate the statistical error, with the assumption that the two transformations have identical convergence properties. The structures of the protein–ligand complex were determined by x–ray crystallography [34, 35].

Compared with pYEEI, ligand **I1** adopts a very similar binding mode. The biphenyl moiety occupies the hydrophobic pocket entirely and the interaction of pY with the hydrophilic site is strong. The scaffold of the peptide forms steady van

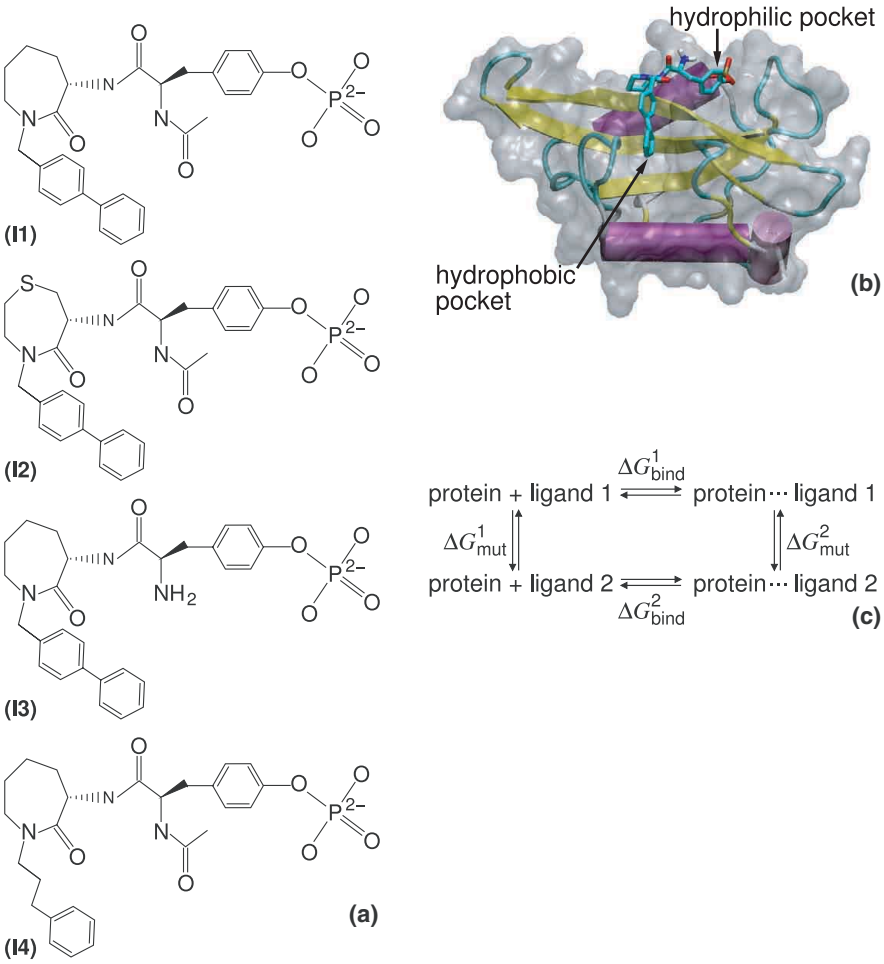

Figure 3. Non–peptide inhibitors of the SH2 domain of the ^{pp60}src kinase (**a**). Complex formed by the src SH2 domain and inhibitor l3 bound to the surface of the latter. Note the diphenyl moiety interacting with the hydrophobic pocket, while the pY motif is buried in the hydrophilic pocket (**b**). Thermodynamic cycle utilized to estimate relative protein–ligand binding free energies. Horizontal transformations are determined experimentally and generally are not amenable to statistical simulations. Vertical transformations correspond to "alchemical transformations" of the ligand in the free state (left) and in the bound state (right), so that $\Delta G^2_{\rm bind} - \Delta G^1_{\rm bind} = \Delta G^2_{\rm mut} - \Delta G^1_{\rm mut}$ (**c**)

der Waals contacts with the surface of the protein. A persistent water molecule is bridged between the carbonyl group of the lactam moiety and the amide –NH group of Lys^{62}. The relative binding free energies corresponding to the "alchemical transformations" of ligand **l1** are gathered in Table 1. Overall, the agreement between the computational and the experimental estimates is good, within chemical accuracy.

Table 1. Point mutations of a series of non–peptide inhibitors of the SH2 domain of the ^{pp60}src kinase. The FEP methodology was employed in association with the dual–topology paradigm. Experimental binding free energies were determined using micro–calorimetry techniques

Transformation	Calculated $\Delta G^2_{\rm mut} - \Delta G^1_{\rm mut}$ (kcal/mol)	Experimental $\Delta G^2_{\rm bind} - \Delta G^1_{\rm bind}$ (kcal/mol)
I1 $\rightarrow$ **I2**	$+0.6\pm0.4$	$+1.3$
I1 $\rightarrow$ **I3**	$+2.9\pm0.4$	$+2.2$
I1 $\rightarrow$ **I4**	$+1.9\pm0.4$	$+1.8$

Replacement of one methylene group in the lactam scaffold by a sulfur atom — see Fig. 3, reduces the binding affinity by about an order of magnitude. No striking structural modification of the scaffold is observed, and the two–pronged plug motif is preserved. A closer look at the protein–ligand interface, however, reveals a loss of van der Waals contacts and unfavorable electrostatic interactions where the point mutation occurred, causing a somewhat increased flexibility in **I2**. Replacement of the amide group of **I1** by an amino group decreases the binding affinity more significantly. Here again, the altered ligand, **I3**, remains perfectly anchored in the hydrophobic and hydrophilic pockets upon mutation. Yet, the interaction of the –NAc moiety with Arg^{14} vanishes as the former is modified into $–NH_2$, resulting in a weaker binding, *ca.* forty times less than for **I1**. Ligands **I1** and **I4** differ in their hydrophobic region, the first phenyl ring of the biphenyl moiety being replaced by two methylene groups. Close examination of the complex formed by the src SH2 domain and **I4** reveals that the second phenyl ring of **I1** and **I4** superimpose nicely. Positioning of the peptide scaffold is, however, modified, so that –NAc no longer interacts with Arg^{14} , but rather with His^{60} . This structural change propagates to the pY ring, which adopts an alternative conformation, corresponding, overall, to a radically different binding mode of the ligand. It is important to underline that the "alchemical transformation" of **I1** into **I4** is able to capture the structural modifications in the protein–ligand association, and, hence, provide an accurate estimate of the relative free energy. This may be ascribed to the rapidly relaxing degrees of freedom of the short peptide, compatible with the time scale of the simulation.

Design of potent leads for a given target is only one element of the complex process that will eventually result in the release of a drug candidate suitable for clinical test phases. Downstream from the ranking of ligands according to their affinity for a protein, bioavailability and toxicity properties should be explored [38]. Of particular interest, the unassisted transport of pharmacologically active molecules across the membrane, in connection with the so–called blood–brain barrier (BBB), is an area where modeling techniques can provide a convincing answer, and, hence, complete the pipeline of *in silico* screening tools. Free energy methods may play a significant role in this effort by offering a detailed picture of the underlying energetics involved in the translocation of a drug from the aqueous medium to the hydrophobic core of the membrane. US calculations were used to simulate the passage of a series of pharmacologically relevant molecules across the interface formed by a lamella of

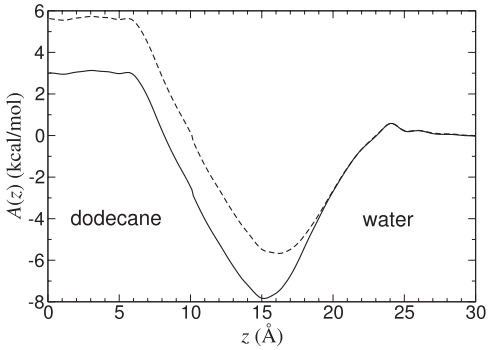

Figure 4. Transfer free energy of antipyrine across the water–dodecane interface. Uncorrected profile (dashed line) and profile incorporating the distortion energy (*solid line*), evaluated from $\Delta A_{\mathrm{distort}} = \left\langle \psi|\hat{\mathcal{H}}_0|\psi \right\rangle - \left\langle \psi_0|\hat{\mathcal{H}}_0|\psi_0 \right\rangle$, where $\hat{\mathcal{H}}_0$ and $|\psi_0\rangle$ are the unperturbed Hamiltonian and wave function. $|\psi\rangle$ is the SCRF wave function that includes the solvent contribution

dodecane in equilibrium with a lamella of water — a pertinent model for estimating partition coefficients [39]. Up to 20 ns were sampled to cover the whole reaction pathway, divided into sequentially overlapping windows to improve efficiency.

An example of such PMFs is depicted in Fig. 4, in the case of antipyrine, for which the net transfer free energy is equal to $+3.0\pm0.5$ kcal/mol, i.e. $\log_{10}P_{\mathrm{wat-dod}} = -2.1$, to be compared to the experimental value of -2.5 [40]. In most circumstances, modification of the environment and the ensuing change in the polarization of the solute is ignored when evaluating differential solvation properties. Traditionally, basis sets that inflate artificially the polarity of the molecule — e.g. 6–31G(d, p), have been utilized to parameterize the electrostatic term of the potential energy function, thereby compensating in an average sense for missing induction effects. Whereas the fitted sets of point charges are generally *ad hoc* for molecular simulations in an aqueous environment, the same cannot be said for non–polar media, in which the polarity of the solute is clearly exaggerated. A convenient framework to circumvent this difficulty is provided by the self–consistent reaction field (SCRF) method [41]. Atomic charges are derived from the electrostatic potential that accounts for the mutual polarization of the molecule by a dielectric continuum surrounding it. Upon translation from one medium to another, less polar one, the positive, reversible work corresponding to the overestimated polarity of the solute in the non–polar environment — i.e. the distortion energy [42], should be accounted for when determining the transfer free energy. In the example of antipyrine, this correction can be as large as *ca.* 2.5 kcal/mol (see Fig. 4), thus, demonstrating the necessity to incorporate this important contribution in solvation free energy calculations.

As indicated above, determination of accurate transfer free energies by means, for instance, of US simulations, is CPU demanding, and, hence, not compatible with the high–throughput requirements of industrial settings. These calculations should, therefore, be performed on the handful of drug candidates discriminated by the

screening process. Faster approaches have been devised, however, based on quantum chemical calculations associated to an SCRF scheme for taking solvent effects into account. Aside from the electrostatic term, van der Waals contributions are usually evaluated from empirical formulation using the solvent accessible surface area (SASA) of the solute. This set of methods is substantially faster than statistical simulations, but, at the same time, only supply simple liquid–liquid partition coefficients, rather than the full free energy behavior characterizing the translocation of the molecule between the two media. This information may, however, turn out to be of paramount importance to rationalize biological phenomena. Such was the case, for instance, of general anesthesia by inhaled anesthetics, an interfacial process shown to result from the accumulation of anesthetics at the water–membrane interface of neuronal tissues [43].

4 Free Energy Calculations and Signal Transduction

The paucity of structural information available for membrane proteins has imparted a new momentum in the *in silico* investigation of these systems. The grand challenge of molecular modeling is to attain the microscopic detail that is often inaccessible to conventional experimental techniques. Of topical interest are seven transmembrane (TM) domain G–protein coupled receptors (GPCRs) [44], which correspond to the third largest family of genes in the human genome, and, therefore, represent privileged targets for *de novo* drug design. Full resolution by x–ray crystallography of the three–dimensional structure of bovine rhodopsin [45], the only GPCR structure to this date, has opened new vistas for the modeling of related membrane proteins. Unfortunately, crystallization of this receptor in its dark, inactive state precludes the use of the structure for homology modeling of GPCR–ligand activated complexes [46]. When neither theory nor experiment can provide atomic–level, three–dimensional structures of GPCRs, their synergistic combination offers an interesting perspective to reach this goal. Such a self–consistent strategy between experimentalists and modelers has been applied successfully to elucidate the structure of the human receptor of cholecystokinin (CCK1R) in the presence of an agonist ligand [47] — *viz.* a nonapeptide (CCK9) [48] of sequence Arg–Asp–S-Tyr–Thr–Gly–Trp–Met–Asp–Phe–NH_2 , where S-Tyr stands for a sulfated tyrosyl amino acid. On the road towards a consistent *in vacuo* construction of the complex, site–directed mutagenesis experiments were designed to pinpoint key receptor–ligand interactions, thereby helping in the placement of TM α–helices and the docking of CCK9.

Whereas *in vacuo* models reflect the geometrical constraints enforced in the course of their construction, it is far from clear whether they will behave as expected when immersed in a realistic membrane environment. Accordingly, the model formed by CCK1R and CC9 was inserted in a fully hydrated palmitoyloleylphosphatidylcholine (POPC) bilayer, resulting in a system of 72,255 atoms, and the complete assembly was probed in a 10.5 ns MD simulation. Analysis of the trajectory reveals no apparent loss of secondary structure in the TM domain, and a distance root mean square deviation for the backbone atoms not exceeding 2 Å. More importantly,

the crucial receptor–ligand interactions are preserved throughout the simulation — e.g. Arg^{336} with Asp^{8} [49], and Met^{195} and Arg^{197} with S-Tyr^{3} [50, 51]. Admittedly, such numerical experiments in essence only supply a qualitative picture of the structural properties of the molecular assembly and its integrity over the time scale explored. Free energy calculations go one step beyond by quantifying intermolecular interactions according to their importance, and, consequently, represent a tangible thermodynamic measure for assessing the accuracy of the model. Furthermore, this information is directly comparable to site–directed mutagenesis experiments utilized in the *in vacuo* construction of the receptor, thereby closing the loop of the modeling process.

To some extent, performing free energy calculations in such large molecular assemblies may be viewed as a bold and perhaps foolish leap of faith, considering the various possible sources of errors likely to affect the final result. Among the latter, attempting to reproduce free energy differences using a three–dimensional model in lieu of a well–resolved, experimentally determined structure casts the greatest doubts on the chances of success of this venture. Of equal concern, "alchemical transformations" involving charged amino acids are driven primarily by the solvation of the ionic moieties, resulting in large free energies, the difference of which, between the free and the bound states, is expected to be small. Assuming a validated model, which appears to be confirmed by the preliminary MD simulation, a key–question, already mentioned in this chapter, remains: When is "enough sampling" really enough? This conundrum should be, in fact, rephrased here as: Are the time scales characteristic of the slowest degrees of freedom in the system crucial for the free energy changes that are being estimated? For instance, is the mutation of the penultimate amino acid of CCK9 — *viz.* Asp^{8} into alanine (see Fig. 5), likely to be affected by the slow collective motions of lipid molecules, or possible vertical and lateral motions of TM α–helices? Nanosecond MD simulations obviously cannot capture these events, which occur over significantly longer times. Yet, under the assumption that the replacement of an agonist ligand by an alternate one does not entail any noticeable rearrangement of the TM domain, current free energy calculations are likely to be appropriate for ranking ligands according to their affinity towards a given GPCR.

The FEP estimate of $+3.1\pm0.7$ kcal/mol for the D8A transformation agrees, indeed, very well with the site–directed mutagenesis experiments that yielded a free energy change equal to $+3.6$ kcal/mol. The *in silico* value was obtained from two runs of 3.4 ns each, in bulk water and in CCK1R, respectively, breaking the reaction path into 114 consecutive windows of uneven width, and using the dual–topology paradigm. The error was estimated from two distinct runs performed at 5.0 and 10.5 ns of the MD simulation. In contrast with an error derived from a first–order expansion of the free energy, which only reflects the statistical precision of the calculation — here, ±0.3 kcal/mol, repeating the simulation from distinct initial conditions accounts for fluctuations of the structure over longer time scales. Put together, while it is difficult to ascertain without ambiguity the correctness of the three–dimensional structure in the sole light of a limited number of numerical experiments, it still remains that the host of observations accrued in theses simulations coincide nicely with the collection of experimental data. *De novo* development of new drug candidates for targets of

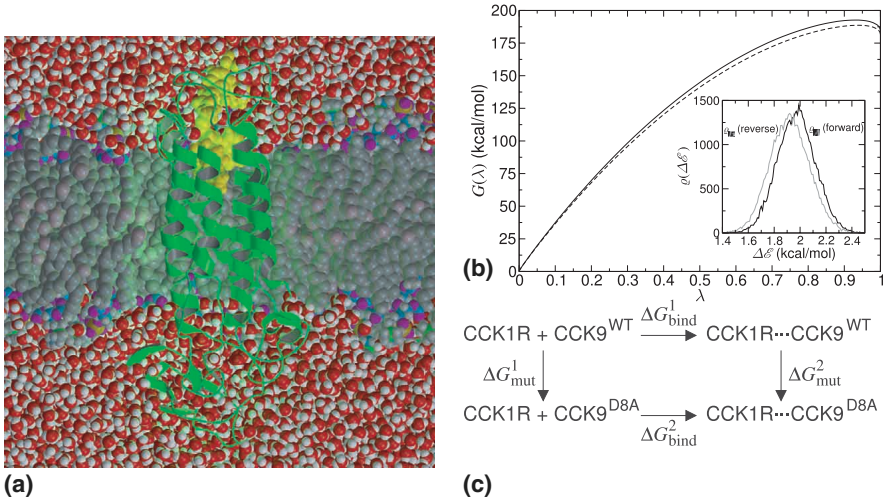

Figure 5. Human receptor of cholecystokinin (CCK1R) embedded in a fully hydrated POPC bilayer. The agonist ligand, nonapeptide CCK9, is highlighted (**a**). Free energy change for the mutation in CCK9 of Asp^8 into alanine: Transformation in the receptor (*solid line*) and in water (*dashed line*). Insert: Overlapping density of states characterizing adjacent states, at $\lambda = 0.5$ (**b**). Thermodynamic cycle utilized to estimate the relative receptor–ligand binding free energy for the D8A point mutation in CCK9 (**c**)

unknown structure constitutes one of the greatest challenges faced today by the pharmaceutical industry. It is envisioned that the encouraging results presented herein for CCK1R will pave the way towards a more self–contained approach to drug design, emancipated from the requirement of well–resolved structures.

5 Free Energy Calculations and Peptide Folding

Capturing the underlying mechanisms that govern protein folding has been one of the holy grails of modern theoretical biophysics. Whereas the sequence of amino acids that forms the protein contains all the necessary information to determine the unique, compact structure of the chain under a given physiological condition, the understanding of the paths that lead to this native, generally biologically active structure remains fragmentary. Two classes of theoretical approaches for tackling the protein folding problem have been employed in recent years. In the first one, an all–atom representation of the protein and its environment and an empirically–based potential energy function are used to evaluate the intra– and intermolecular forces acting on the system. This description is often associated to an MD scheme for exploring the rugged conformational space of the solvated protein. Yet, ab initio folding by means of atomic–level MD simulations featuring both the protein and the solvent remain limited to short peptides and small proteins [52]. The reason for this limitation is rooted in the computational cost of such *in silico* experiments, which does

not allow biologically relevant time scales to be accessed routinely. An alternative to the detailed, all–atom approach consists in turning to somewhat rougher models, that, nonetheless, retain the fundamental characteristics of protein chains. Such is the case of coarse–grained models, in which each amino acid of the protein is represented by a bead located at the vertex of a two– or a three–dimensional lattice [53]. An intermediate description consists of an all–atom representation of the protein in an implicit solvent. It is far from clear, however, whether the delicate interplay of the protein with explicit water molecules is a necessary condition for guaranteeing the correct folding toward the native state.

Whereas MD simulations involving an explicit solvent rarely exceed a few hundreds of ns [54], significantly shorter free energy calculations can be designed advantageously to understand the physical phenomena that drive folding. Among these phenomena, the subtle, temperature–dependent hydrophobic effect [55, 56] remains one of the most investigated to rationalize the collapse of a disordered protein chain into an appropriately folded one. The choice of a pertinent reaction coordinate that characterizes the folding process of a short peptide, let alone a small protein, constitutes a conundrum, unlikely to find a definitive answer in the near future. This intricate problem is rooted in the vast number of degrees of freedom that vary concomitantly as the peptide evolves toward a folded structure. The free energy is, therefore, a function of many variables that cannot be accounted for in a straightforward fashion. Valuable information may, nonetheless, be obtained from simple model systems, for which a non–ambiguous reaction coordinate can be defined. Such is the case of the terminally blocked undecamer of L–leucine organized in an α–helix, the C–terminal residue of which was unfolded from an α–helical conformation to that of a β–strand [57]. Similar calculations have been endeavored with blocked poly–L–alanine of various lengths to examine helix propagation at its N– and C–termini [58]

To highlight the temperature–dependent nature of the hydrophobic effect, the MD simulations were run in the canonical ensemble at 280, 300, 320, 340, 360 and 370 K, using the Nosé–Hoover algorithm implemented in the program COSMOS. Changes in the free energy consecutive to modifications of the last ψ dihedral angle of the homopolypeptide were estimated with the US method. All other torsional angles were restrained softly in a range characteristic of an α–helix. For each temperature, the complete reaction pathway connecting the α–helical state to the β–strand — roughly speaking $-90 \leq \psi \leq +170^\circ$, was broken down in five mutually overlapping windows. The full free energy profiles were subsequently reconstructed employing WHAM. The total simulation length varied from 14 ns at 370 K, to 76 ns at 280 K, on account of the slower relaxation at lower temperatures.

The PMFs shown in Fig. 6 each possess two distinct local minima corresponding to the α–helix and the β–strand, and separated by a maximum of the free energy around 90° . These three conformational states are characterized by different SASAs — *viz.* 134 ± 2, 149 ± 7 and 117 ± 13 Å^2 for the α–helix, the transition state (TS) and the β–strand, respectively. The most striking feature of the PMFs lies in the temperature dependence of the free energy associated to the transition from the α–helix to the TS, $\Delta A_{\alpha_R \to \ddagger}$, and that from the β–strand to the TS, $\Delta A_{\beta \to \ddagger}$. Furthermore, neither $\Delta A_{\alpha_R \to \ddagger}(T)$ nor $\Delta A_{\beta \to \ddagger}(T)$ varies monotonically as the temperature increases. On

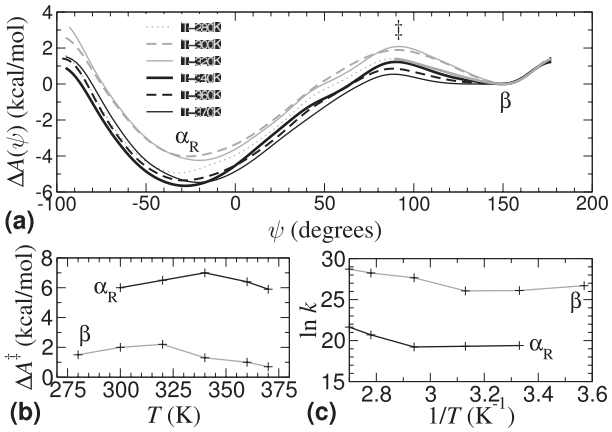

Figure 6. Free energy profile characterizing the unfolding of the undecamer of poly–leucine at different temperatures (**a**). Activation free energies from the α–helix to the transition state, $\ddagger$, and from the β–strand to $\ddagger$ (**b**). Arrhenius plot for the same transitions, inferred from the activated complex theory, *viz.* $\ln k = \ln h/\beta - \beta \Delta A^{\ddagger}$. h is the Planck constant (**c**)

the contrary, the two profiles exhibit a maximum at a characteristic temperature T_i, suggestive of a non–Arrhenius behavior driven by physical principles that are related to the solvation properties of the hydrophobic homopolypeptide. Considering that the variation of entropy between state Ξ — i.e. either the α–helix or the β–strand, and the TS is $\Delta S^{\ddagger}(T) = -\partial \Delta A^{\ddagger}(T)/\partial T$, it follows that at T_i, the entropy is zero. Since $\Delta A^{\ddagger}(T)$ decreases when $T > T_i$, $\Delta S^{\ddagger}(T)$ is positive in this region, hence, implying that the conformational transition is favored entropically. Moreover, as $\Delta A^{\ddagger}(T) > 0$, the change in internal energy, $\Delta U^{\ddagger}(T) = \Delta A^{\ddagger}(T) + T\Delta S^{\ddagger}(T)$, is also necessarily positive, so that Ξ is the thermodynamically favored state. Symmetrically, when $T < T_i$, $\Delta S^{\ddagger}(T)$ is negative and Ξ is the entropically favored state. On account of the greater SASA for the TS, the transition from Ξ to the TS can be viewed as the transfer of a portion of the side chain from a buried state to a solvent–exposed one, similar in spirit to the translocation of a hydrophobic solute from a non–polar medium to water, that is accompanied by a positive free energy change — the thermodynamic fingerprint of the hydrophobic effect [59]. Put together, the free energy barrier can be ascribed at high temperature to the sole internal energy, to which an increasing entropic contribution is added upon lowering the temperature.

A fundamental issue brought to light in this investigation is the non–Arrhenius behavior of the unfolding kinetics of the hydrophobic peptide, and the underlying temperature dependence of the free energy barrier. This result is in blatant contradiction with the hypothesis commonly adopted in lattice model simulations that the activation free energy is a constant, and further suggests that a temperature dependence should be introduced in these models to capture subtle solvent contributions, like the hydrophobic effect [55].

6 Free Energy Calculations and Membrane Protein Association

To a large extent, our knowledge of how membrane protein domains recognize and associate into functional, three–dimensional entities remains fragmentary. Whereas the structure of membrane proteins can be particularly complex, their TM region is often simple, consisting in general of a bundle of α–helices, or barrels of β–strands. An important result brought to light by deletion experiments indicates that some membrane proteins can retain their biological function upon removal of large fractions of the protein. This is suggestive that rudimentary models, like simple α–helices, can be utilized to understand the recognition and association processes of TM segments into complex membrane proteins. On the road to reach this goal, the "two–stage" model [60] provides an interesting view for rationalizing the folding of membrane proteins. According to this model, elements of the secondary structure — *viz.* in most cases, α–helices, are first formed and inserted into the lipid bilayer, prior to specific inter–helical interactions that drive the TM segments towards well–ordered, native structures. Capturing the atomic detail of the underlying mechanisms responsible for α–helix recognition and association requires model systems that are supported by robust experimental data to appraise the accuracy of the computations endeavored. Glycophorin A (GpA), a glycoprotein ubiquitous to the human erythrocyte membrane, represents one such system. It forms non–covalent dimers through the reversible association of its membrane–spanning domain — i.e. residues 62 to 101, albeit only residues 73 to 96 actually adopt an α–helical conformation [61–63]. Inter–helical association has been shown to result from specific interactions involving a heptad of residues, essentially located on one face of each TM segment, as may be seen in Fig. 7.

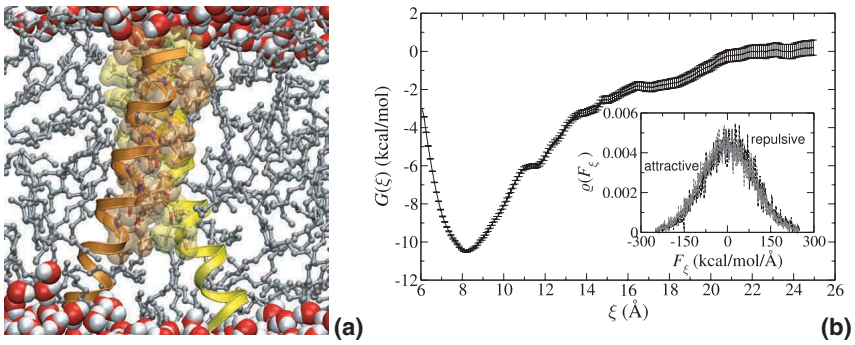

Figure 7. TM domain of GpA formed by a homodimer of α–helices, embedded in a lipid membrane mimetic. The heptad of residues involved in the association of the TM segments are shown as transparent van der Waals spheres. Note the crossing angle between the two α–helices, equal to *ca.* 40° (**a**). Free energy profile delineating the reversible association of the TM segments. Insert: Gaussian–like distributions of the force acting along the reaction coordinate, ξ , in the repulsive region — *viz.* $\xi < 8$ Å, and in the attractive region — *viz.* $\xi > 8$ Å (**b**)

The reversible association of GpA in a lipid bilayer was modelled using its dimeric, α–helical TM segments immersed in a membrane mimetic formed by a lamella of dodecane placed between two lamellae of water. The ABF method introduced in the program NAMD [26] was employed to allow the TM segments to diffuse freely along the reaction coordinate, ξ , chosen to be the distance separating the centers of mass of the two α–helices. Such a free energy calculation is not only challenging methodologically, but it is also of paramount importance from a biophysical standpoint, because it bridges structural data obtained from nuclear magnetic resonance (NMR) [61–63] to thermodynamic data obtained from analytical ultracentrifugation [64, 65] and fluorescence resonance energy transfer (FRET) [66, 67], providing a dynamic view of the recognition and association stages.

Visual inspection of the PMF derived from a 125 ns simulation and describing the reversible association of the α–helices reveals a qualitatively simple profile, featuring a single minimum characteristic of the native, dimeric state. As ξ increases, so does the free energy, progressing by steps that correspond to the successive breaking of inter–helical contacts. Beyond 21 Å, the TM segments are sufficiently separated to assume that they no longer interact. Integration of the PMF in the limit of α–helix association yields the association constant and, hence, the free energy of dimerization, equal to $+11.5\pm0.4$ kcal/mol. Direct and precise comparison of this value with experiment is not possible, because measurements were carried in different environments, namely hydrocarbon *vs.*detergent micelles. It can, nonetheless, be inferred that the value in dodecane probably constitutes an upper bound to the experimental estimates determined in micelles, on account of (i) the greater order imposed by the detergent chains, and, (ii) the hydrophobic fraction of the system that increases with the length of the chain [67].

Deconvolution of the PMF into free energy components illuminates two distinct regimes controlling recognition and association. At large separations, as inter–helical contacts vanish, the helix–helix term becomes progressively negligible, resulting essentially from the interaction of two macro–dipoles. The TM segments are stabilized by favorable helix–solvent contributions. In contrast, at short separations, helix–helix interactions are prominent and govern the change in the free energy near the global minimum. Association proceeds through the transient formation of early, non–native contacts involving residues that act as recognition sites. These contacts are subsequently replaced by contacts in the heptad of residues responsible for association, concomitantly with the tilt of the two α–helices from an upright position to that characteristic of the native dimer.

7 Conclusion

Free energy calculations constitute a tangible link between theory and experiment, by quantifying at the thermodynamic level the physical phenomena modelled by statistical simulations. With twenty years of hindsight gained from methodological development and characterization, a variety of problems of both chemical and biological relevance can now be tackled with confidence. Among the progresses achieved in

recent years, a significant step forward has been made in the calculation of free energies along a reaction coordinate, in particular employing the concept of an average force acting along this coordinate [25, 26]. "Alchemical transformations" utilized to mimic site–directed mutagenesis experiments have also benefited from advances in the understanding of the methodology and how the latter should be applied [68]. Efforts to characterize and estimate the error affecting the simulations [10, 11, 69] have equally played an active role in turning free energy calculations into another tool in the arsenal of computational methods available to the modeler. Put together, free energy calculations have come of age to become a predictive approach, instead of remaining at the stage of a mere proof of concept [70]. As has been illustrated in this chapter, they can be applied to numerous problems, ranging from *de novo* drug design to the understanding of biophysical processes in lipid membranes. Free energy calculations, however, cannot yet be considered as "black box", routine jobs. A robust, reliable methodology does not necessarily imply that it can be used blindly. The nature of the problem dictates the choice of the method and the associated protocol — e.g. US *vs.*FEP *vs.*TI *vs.*average force, the number of intermediate λ–states along a reaction coordinate, the pertinent choice of this reaction coordinate, the amount of sampling per individual λ–state, a single *vs.*a dual topology paradigm for "alchemical transformations", or constrained *vs.*unconstrained MD. Furthermore, little effort has been hitherto devoted to the automatization of free energy calculations, through, for instance, a user–friendly definition of the topologies representative of the initial and the final states of a transformation. With the increased access to massively parallel architectures as the price/performance ratio of computer chips continues to fall inexorably, the bottleneck of free energy calculations has shifted from a purely computational aspect to a human one, due to the need of qualified modelers to set these calculations up. This explains why their use in industry, and particularly in the pharmaceutical world, over the past years has remained scarce. Cutting–edge applications of free energy calculations emanate essentially from academic environments, where the focus is not so much on high throughput, but rather on well–delineated, specific problems that often require more human attention than computational power. Yet, it is envisioned that in a reasonably near future, free energy methods will become an unavoidable element of screening pipelines, discriminating between candidates selected from cruder approaches, to retain only the best leads towards a given target.

Acknowledgments

Jérôme Hénin, Surjit Dixit, Olivier Collet, Eric Darve, Andrew Pohorille and Alan E. Mark are gratefully acknowledged for fruitful and inspiring discussions. The author thank the Centre Informatique National de l'Enseignement Supérieur (CINES) and the centre de Calcul Réseaux et Visualisation Haute Performance (CRVHP) for generous provision of CPU time on their SGI Origin 3000 architectures.

References

[1] Kollman, P. A., Free energy calculations: Applications to chemical and biochemical phenomena, *Chem. Rev.* 93, 2395–2417, 1993.

[2] Postma, J. P. M.; Berendsen, H. J. C.; Haak, J. R., Thermodynamics of cavity formation in water: A molecular dynamics study, *Faraday Symp. Chem. Soc.* 17, 55–67, 1982.

[3] Warshel, A., Dynamics of reactions in polar solvents. Semiclassical trajectory studies of electron transfer and proton transfer reactions, *J. Phys. Chem.* 86, 2218–2224, 1982.

[4] Bash, P. A.; Singh, U. C.; Langridge, R.; Kollman, P. A., Free energy calculations by computer simulation, *Science* 236, 564–568, 1987.

[5] Bash, P. A.; Singh, U. C.; Brown, F. K.; Langridge, R.; Kollman, P. A., Calculation of the relative change in binding free energy of a protein–inhibitor complex, *Science* 235, 574–576, 1987.

[6] McQuarrie, D. A., *Statistical mechanics*, Harper and Row: New York, 1976.

[7] Allen, M. P.; Tildesley, D. J., *Computer Simulation of Liquids*, Clarendon Press: Oxford, 1987.

[8] Kirkwood, J. G., Statistical mechanics of fluid mixtures, *J. Chem. Phys.* 3, 300–313, 1935.

[9] Zwanzig, R. W., High–temperature equation of state by a perturbation method. I. Nonpolar gases, *J. Chem. Phys.* 22, 1420–1426, 1954.

[10] Lu, N.; Singh, J. K.; Kofke, D. A.; Woolf, T. B., Appropriate methods to combine forward and reverse free–energy perturbation averages, *J. Chem. Phys.* 118, 2977–2984, 2003.

[11] Lu, N.; Kofke, D. A.; Woolf, T. B., Improving the efficiency and reliability of free energy perturbation calculations using overlap sampling methods, *J. Comput. Chem.* 25, 28–39, 2004.

[12] Mark, A. E. Free Energy Perturbation Calculations. in *Encyclopedia of computational chemistry*, Schleyer, P. v. R.; Allinger, N. L.; Clark, T.; Gasteiger, J.; Kollman, P. A.; Schaefer III, H. F.; Schreiner, P. R., Eds., vol. 2. Wiley and Sons, Chichester, 1998, pp. 1070–1083.

[13] Torrie, G. M.; Valleau, J. P., Nonphysical sampling distributions in Monte Carlo free energy estimation: Umbrella sampling, *J. Comput. Phys.* 23, 187–199, 1977.

[14] Pearlman, D. A.; Kollman, P. A., The overlooked bond–stretching contribution in free energy perturbation calculations, *J. Chem. Phys.* 94, 4532–4545, 1991.

[15] Boresch, S.; Karplus, M., The role of bonded terms in free energy simulations: I. Theoretical analysis, *J. Phys. Chem. A* 103, 103–118, 1999.

[16] Boresch, S.; Karplus, M., The role of bonded terms in free energy simulations: II. Calculation of their influence on free energy differences of solvation, *J. Phys. Chem. A* 103, 119–136, 1999.

[17] Pearlman, D. A., A comparison of alternative approaches to free energy calculations, *J. Phys. Chem.* 98, 1487–1493, 1994.

[18] Beutler, T. C.; Mark, A. E.; van Schaik, R. C.; Gerber, P. R.; van Gunsteren, W. F., Avoiding singularities and neumerical instabilities in free energy calculations based on molecular simulations, *Chem. Phys. Lett.* 222, 529–539, 1994.
[19] Straatsma, T. P.; Berendsen, H. J. C., Free energy of ionic hydration: Analysis of a thermodynamic integration technique to evaluate free energy differences by molecular dynamics simulations, *J. Chem. Phys.* 89, 5876–5886, 1988.
[20] Chandler, D., *Introduction to modern statistical mechanics*, Oxford University Press, 1987.
[21] Kumar, S.; Bouzida, D.; Swendsen, R. H.; Kollman, P. A.; Rosenberg, J. M., The weighted histogram analysis method for free energy calculations on biomolecules. I. The method, *J. Comput. Chem.* 13, 1011–1021, 1992.
[22] Pearlman, D. A., Determining the contributions of constraints in free energy calculations: Development, characterization, amnd recommendations, *J. Chem. Phys.* 98, 8946–8957, 1993.
[23] den Otter, W. K.; Briels, W. J., The calculation of free–energy differences by constrained molecular dynamics simulations, *J. Chem. Phys.* 109, 4139–4146, 1998.
[24] den Otter, W. K., Thermodynamic integration of the free energy along a reaction coordinate in Cartesian coordinates, *J. Chem. Phys.* 112, 7283–7292, 2000.
[25] Darve, E.; Pohorille, A., Calculating free energies using average force, *J. Chem. Phys.* 115, 9169–9183, 2001.
[26] Hénin, J.; Chipot, C., Overcoming free energy barriers using unconstrained molecular dynamics simulations, *J. Chem. Phys.* 121, 2904–2914, 2004.
[27] Jorgensen, W. L.; Ravimohan, C., Monte Carlo simulation of differences in free energies of hydration, *J. Chem. Phys.* 83, 3050–3054, 1985.
[28] Chipot, C.; Kollman, P. A.; Pearlman, D. A., Alternative approaches to potential of mean force calculations: Free energy perturbation versus thermodynamic integration. Case study of some representative nonpolar interactions, *J. Comput. Chem.* 17, 1112–1131, 1996.
[29] Widom, B., Some topics in the theory of fluids, *J. Chem. Phys.* 39, 2808–2812, 1963.
[30] Straatsma, T. P.; Berendsen, H. J. C.; Stam, A. J., Estimation of statistical errors in molecular simulation calculations, *Mol. Phys.* 57, 89–95, 1986.
[31] Chipot, C.; Pohorille, A., Conformational equilibria of terminally blocked single amino acids at the water–hexane interface. A molecular dynamics study, *J. Phys. Chem. B* 102, 281–290, 1998.
[32] Chipot, C.; Millot, C.; Maigret, B.; Kollman, P. A., Molecular dynamics free energy perturbation calculations. Influence of nonbonded parameters on the free energy of hydration of charged and neutral species, *J. Phys. Chem.* 98, 11362–11372, 1994.
[33] Soriano, P.; Montgomery, C.; Geske, R.; Bradley, A., Targeted disruption of the c–src proto–oncogene leads to osteopetrosis in mice., *Cell* 64, 693–702, 1991.
[34] Lange, G.; Lesuisse, D.; Deprez, P.; Schoot, B.; Loenze, P.; Benard, D.; Marquette, J. P.; Broto, P.; Sarubbi, E.; Mandine, E., Principles governing the

binding of a class of non–peptidic inhibitors to the SH2 domain of src studied by X-ray analysis, *J. Med. Chem.* 45, 2915–2922, 2002.

[35] Lange, G.; Lesuisse, D.; Deprez, P.; Schoot, B.; Loenze, P.; Benard, D.; Marquette, J. P.; Broto, P.; Sarubbi, E.; Mandine, E., Requirements for specific binding of low affinity inhibitor fragments to the SH2 domain of ^{pp60}Src are identical to those for high affinity binding of full length inhibitors, *J. Med. Chem.* 46, 5184–5195, 2003.

[36] Kale, L.; Skeel, R.; Bhandarkar, M.; Brunner, R.; Gursoy, A.; Krawetz, N.; Phillips, J.; Shinozaki, A.; Varadarajan, K.; Schulten, K., NAMD2: Greater scalability for parallel molecular dynamics, *J. Comput. Phys.* 151, 283–312, 1999.

[37] Bhandarkar, M.; Brunner, R.; Chipot, C.; Dalke, A.; Dixit, S.; Grayson, P.; Gullingsrud, J.; Gursoy, A.; Humphrey, W.; Hurwitz, D. et al. NAMD *users guide, version 2.5*. Theoretical biophysics group, University of Illinois and Beckman Institute, 405 North Mathews, Urbana, Illinois 61801, September 2003.

[38] Carrupt, P.; Testa, B.; Gaillard, P. Computational approaches to lipophilicity: Methods and applications. in *Reviews in Computational Chemistry*, Lipkowitz, K.; Boyd, D. B., Eds., vol. 11. VCH, New York, 1997, pp. 241–345.

[39] Wohnsland, F.; Faller, B., High–throughput permeability pH profile and high–throughput alkane–water log P with artificial membranes, *J. Med. Chem.* 44, 923–930, 2001.

[40] Bas, D.; Dorison-Duval, D.; Moreau, S.; Bruneau, P.; Chipot, C., Rational determination of transfer free energies of small drugs across the water–oil interface, *J. Med. Chem.* 45, 151–159, 2002.

[41] Rivail, J. L.; Rinaldi, D., A quantum chemical approach to dielectric solvent effects in molecular liquids, *Chem. Phys.* 18, 233–242, 1976.

[42] Chipot, C., Rational determination of charge distributions for free energy calculations, *J. Comput. Chem.* 24, 409–415, 2003.

[43] Pohorille, A.; Wilson, M.A.; New, M.H.; Chipot, C., Concentrations of anesthetics across the water–membrane interface; The Meyer–Overton hypothesis revisited, *Toxicology Lett.* 100, 421–430, 1998.

[44] Takeda, S.; Haga, T.; Takaesu, H.; Mitaku, S., Identification of G protein–coupled receptor genes from the human genome sequence, *FEBS Lett.* 520, 97–101, 2002.

[45] Palczewski, K.; Kumasaka, T.; Hori, T.; Behnke, C. A.; Motoshima, H.; Fox, B. A.; Le Trong, I.; Teller, D. C.; Okada, T.; Stenkamp, R. E.; Yamamoto, M.; Miyano, M., Crystal structure of rhodopsin: A G protein–coupled receptor, *Science* 289, 739–745, 2000.

[46] Archer, E.; Maigret, B.; Escrieut, C.; Pradayrol, L.; Fourmy, D., Rhodopsin crystal: New template yielding realistic models of G–protein–coupled receptors ?, *Trends Pharmacol. Sci.* 24, 36–40, 2003.

[47] Talkad, V. D.; Fortune, K. P.; Pollo, D. A.; Shah, G. N.; Wank, S. A.; Gardner, J. D., Direct demonstration of three different states of the pancreatic cholecystokinin receptor, *Proc. Natl. Acad. Sci. USA* 91, 1868–1872, 1994.

[48] Moroder, L.; Wilschowitz, L.; Gemeiner, M.; Göhring, W.; Knof, S.; Scharf, R.; Thamm, P.; Gardner, J. D.; Solomon, T. E.; Wünsch, E., Zur Synthese von Cholecystokinin–Pankreozymin. Darstellung von [28–Threonin, 31–Norleucin]– und [28–Threonin, 31–Leucin]– Cholecystokinin–Pankreozymin–(25–33)–Nonapeptid, *Z. Physiol. Chem.* 362, 929–942, 1981.

[49] Gigoux, V.; Escrieut, C.; Fehrentz, J. A.; Poirot, S.; Maigret, B.; Moroder, L.; Gully, D.; Martinez, J.; Vaysse, N.; Fourmy, D., Arginine 336 and Asparagine 333 of the human cholecystokinin–A receptor binding site interact with the penultimate aspartic acid and the C–terminal amide of cholecystokinin, *J. Biol. Chem.* 274, 20457–20464, 1999.

[50] Gigoux, V.; Escrieut, C.; Silvente-Poirot, S.; Maigret, B.; Gouilleux, L.; Fehrentz, J. A.; Gully, D.; Moroder, L.; Vaysse, N.; Fourmy, D., Met–195 of the cholecystokinin–A interacts with the sulfated tyrosine of cholecystokinin and is crucial for receptor transition to high affinity state, *J. Biol. Chem.* 273, 14380–14386, 1998.

[51] Gigoux, V.; Maigret, B.; Escrieut, C.; Silvente-Poirot, S.; Bouisson, M.; Fehrentz, J. A.; Moroder, L.; Gully, D.; Martinez, J.; Vaysse, N.; Fourmy, D., Arginine 197 of the cholecystokinin–A receptor binding site interacts with the sulfate of the peptide agonist cholecystokinin, *Protein Sci.* 8, 2347–2354, 1999.

[52] Daggett, V., Long timescale simulations, *Curr. Opin. Struct. Biol.* 10, 160–164, 2000.

[53] Taketomi, H.; Ueda, Y.; Gō, N., Studies on protein folding, unfolding and fluctuations by computer simulation. 1. The effect of specific amino acid sequence represented by specific inter–unit interactions, *Int. J. Pept. Protein Res.* 7, 445–459, 1975.

[54] Duan, Y.; Kollman, P. A., Pathways to a protein folding intermediate observed in a 1–microsecond simulation in aqueous solution, *Science* 282, 740–744, 1998.

[55] Pratt, L. R., Molecular theory of hydrophobic effects: "She is too mean to have her name repeated", *Annu. Rev. Phys. Chem.* 53, 409–436, 2002.

[56] Pratt, L. R.; Pohorille, A., Hydrophobic effects and modeling of biophysical aqueous solution interfaces, *Chem. Rev.* 102, 2671–2692, 2002.

[57] Collet, O.; Chipot, C., Non–Arrhenius behavior in the unfolding of a short, hydrophobic α–helix. Complementarity of molecular dynamics and lattice model simulations, *J. Am. Chem. Soc.* 125, 6573–6580, 2003.

[58] Young, W. S.; Brooks III, C. L., A microscopic view of helix propagation: N and C–terminal helix growth in alanine helices, *J. Mol. Biol.* 259, 560–572, 1996.

[59] Shimizu, S.; Chan, H. S., Temperature dependence of hydrophobic interactions: A mean force perspective, effects of water density, and non–additivity of thermodynamics signature, *J. Am. Chem. Soc.* 113, 4683–4700, 2000.

[60] Popot, J. L.; Engelman, D. M., Membrane protein folding and oligomerization: The two–stage model, *Biochemistry* 29, 4031–4037,1990.

[61] MacKenzie, K. R.; Prestegard, J. H.; Engelman, D. M., A transmembrane helix dimer: Structure and implications, *Science* 276, 131–133, 1997.

[62] MacKenzie, K. R.; Engelman, D. M., Structure–based prediction of the stability of transmembrane helix–helix interactions: The sequence dependence of glycophorin A dimerization, *Proc. Natl. Acad. Sci. USA* 95, 3583–3590, 1998.

[63] Smith, S. O.; Song, D.; Shekar, S.; Groesbeek, M.; Ziliox, M.; Aimoto, S., Structure of the transmembrane dimer interface of glycophorin A in membrane bilayers, *Biochemistry* 40, 6553–6558, 2001.

[64] Fleming, K. G.; Ackerman, A. L.; Engelman, D. M., The effect of point mutations on the free energy of transmembrane α–helix dimerization, *J. Mol. Biol.* 272, 266–275, 1997.

[65] Fleming, K. G., Standardizing the free energy change of transmembrane helix–helix interactions, *J. Mol. Biol.* 323, 2002, 563–571.

[66] Fisher, L. E.; Engelman, D. M.; Sturgis, J. N., Detergents modulate dimerization, but not helicity, of the glycophorin A transmembrane domain, *J. Mol. Biol.* 293, 639–651, 1999.

[67] Fisher, L. E.; Engelman, D. M.; Sturgis, J. N., Effects of detergents on the association of the glycophorin A transmembrane helix, *Biophys. J.* 85, 3097–3105, 2003.

[68] Dixit, S. B.; Chipot, C., Can absolute free energies of association be estimated from molecular mechanical simulations ? The biotin–streptavidin system revisited, *J. Phys. Chem. A* 105, 9795–9799, 2001.

[69] Rodriguez-Gomez, D.; Darve, E.; Pohorille, A., Assessing the efficiency of free energy calculation methods, *J. Chem. Phys.* 120, 3563–3570, 2004.

[70] Simonson, T.; Archontis, G.; Karplus, M., Free energy simulations come of age: Protein–ligand recognition, *Acc. Chem. Res.* 35, 430–437, 2002.