Published in final edited form as:

J Comput Chem. 2011 April 30; 32(6): 1135–1143. doi:10.1002/jcc.21695.

Extending fragment-based free energy calculations with library Monte Carlo simulation: Annealing in interaction space

Steven Lettieri, Artem B. Mamonov, and Daniel M. Zuckerman

Department of Computational and Systems Biology University of Pittsburgh, Pittsburgh PA, 15206

Abstract

Pre-calculated libraries of molecular fragment configurations have previously been used as a basis for both equilibrium sampling (via library-based Monte Carlo) and for obtaining absolute free energies using a polymer-growth formalism. Here, we combine the two approaches to extend the size of systems for which free energies can be calculated. We study a series of all-atom poly-alanine systems in a simple dielectric solvent and find that precise free energies can be obtained rapidly. For instance, for 12 residues, less than an hour of single-processor is required. The combined approach is formally equivalent to the annealed importance sampling algorithm; instead of annealing by decreasing temperature, however, interactions among fragments are gradually added as the molecule is grown. We discuss implications for future binding affinity calculations in which a ligand is grown into a binding site.

1 Introduction

Free energy differences, ΔF , are fundamental to physical chemistry. In the context of biomacromolecules, ΔF values can quantify folding stability, relative populations, and binding affinity [1]. Although computer simulations have been used to estimate ΔF values for biomolecules, success has been hampered by well-appreciated sampling problems [1, 2].

A large number of numerical techniques have been used to calculate molecular free energy differences, but a smaller subset is capable of estimating absolute free energies values

$$F = -k_B T \ln Z \quad (1)$$

where Z is the dimensionless configurational partition function. The most straightforward way to estimate an absolute free energy is using a reference system with an exactly calculable free energy (F_{ref}), so that $F = F_{\text{ref}} + \Delta F$. ΔF is then obtained using a standard free energy difference technique, yielding the absolute free energy of the system. This long-established strategy (e.g., [3]) was first suggested for molecular systems by Stoessel and Nowak using a harmonic reference system [4]. Other strategies for calculating absolute free energies are also possible, as demonstrated by the work of Gilson and coworkers [1, 5, 6, 7, 8, 9], as well as by Meirovitch and coworkers [10, 11, 12, 13, 14, 15, 16] and Brooks and coworkers [17, 18, 19].

The present study builds on earlier work in our group using reference systems to calculate absolute free energies for molecular fragments which subsequently are combined into a full molecule [20, 21]. This polymer-growth strategy employs a reference system of non-interacting fragments (e.g. amino acids), which are combined to yield a ΔF correction accounting for all interactions in the full molecule. Our earlier study yielded the absolute free energy for tetra-alanine (Ace-Ala₄-Nme), but could not easily be applied to significantly larger systems [20]. Clark et al. developed a closely related fragment-based approach for estimating binding affinities without however, accounting for fragment flexibility [22]

Here we employ a rigorous “annealing” strategy [23, 24, 25, 26] which integrates the polymer-growth approach with our recently developed library-based Monte Carlo (LBMC) method [27, 28]. We anneal not by lowering temperature, but by adding interactions between previously non-interacting fragments. The addition of interactions is formally equivalent to “growing” the polymer [20, 27]. Interactions among fragments are added gradually over several stages, permitting the calculation of incremental free energy differences until the full molecule and final free energy is produced. A weighted configurational ensemble is generated at each stage, which is then “relaxed” by canonical simulation based on the stage-specific interactions. This alternation of adding interactions and relaxing is formally equivalent to “annealed importance sampling” (AIS) [23, 24, 25, 26]. We also employ resampling as a variance-reduction technique, following our previous temperature-based annealing [26].

The annealing strategy is enhanced by our use of LBMC [27] for the relaxation phases. Like other free energy methods, annealed importance sampling employs canonical sampling (here termed “relaxation”) at each stage - which can be performed by any algorithm that correctly samples the stage-specific distributions. In the present study, LBMC is a natural choice because it is based on fragments and was shown to be highly efficient for sampling flexible peptides in implicit solvent [28]. The addition of arbitrary interactions for staging is also straightforward in LBMC. Nevertheless, the sampling method is “orthogonal” to the free energy calculation, and other canonical sampling algorithms (e.g. [29, 30, 31, 32, 33, 34, 35, 36, 21]) possibly combined with hardware-based improvements [37, 38, 39, 40, 41] could be used instead of LBMC.

Aside from the issue of calculating absolute free energies, the annealing approach for obtaining ΔF values can be set in the context of other methods which directly couple equilibrium sampling and free energy estimation [18, 17, 23, 24, 25, 26, 42, 43, 44, 45, 46]. That is, once the “staging” of the ΔF calculation is established –adding interactions in our case, and incrementing a parameter λ in many others [3] – equilibrium sampling at each stage can be performed independently or in a coupled way. Thermodynamic integration [3] performs independent simulations for each stage, but in many more recent approaches, ensembles at a given stage are used to aid sampling at other stages [18, 17, 23, 24, 25, 26, 42, 43, 44, 45] with “lambda dynamics” being a good example [47, 19, 17, 18]. The annealing strategy couples sampling to staging in a unidirectional way, presumably starting from the stage which is easiest to sample. In our case, it is trivial to generate fully independent configurations for the reference stage of non-interacting fragments; interactions are then “annealed in.”

The annealing approach combining library-based polymer-growth and LBMC yields good results in this initial study, significantly extending our previous growth work, which was limited to peptides of 5 residues [20]. Using this method, we can compute precise free energies for Ace-Ala₁₂-Nme in several hours and for peptides up to 22 residues in about two weeks of computing on a single 3.0 GHz processor core. Free energies and equilibrium ensembles of smaller peptides can be obtained in seconds or minutes depending on the desired precision. We validate our results by comparing the equilibrium ensembles produced during annealing to independent Langevin simulations with a collective simulation time of 1 μ s.

2 Methods

We describe how our earlier fragment-based growth method [20] can be improved using library-based Monte Carlo [27, 28]. Formally, however, our procedure is not new, but is a special case of “annealed importance sampling” (AIS) [23, 25, 26]. We therefore describe

our procedure in terms of AIS, which indeed provides a very natural formal framework. The basic idea of AIS, with resampling, is to start with a well sampled ensemble for an “easy” condition (e.g. high temperature) and then progressively “relax” that ensemble through a series of states (e.g. lower temperature) until the target ensemble is reached. Resampling is used to ensure the correct distribution is obtained at each stage. In our approach however, instead of lowering temperature, we incrementally *add interactions* among molecular fragments. Said another way, interactions between fragments are incrementally “annealed in” - i.e. simply turned on - between successive growth stages.

2.1 Polymer growth with relaxation: annealed importance sampling

The formalism introduced by Neal as AIS [23] will be applied to generalize standard polymer growth algorithms; see also [24]. It can be described in a straightforward way based on an arbitrary set of un-normalized distributions $\pi_i(\mathbf{x})$, with $i = 1, 2, \dots, N$ representing the index of the (growth) stage. In our case, these distributions are standard Boltzmann factors:

$$\pi_i(\mathbf{x}) \propto e^{-U_i(\mathbf{x})/k_B T_i} \quad (2)$$

where U_i is the potential energy at growth stage i at temperature T_i . In our case, annealing is performed only in interaction space so that $T_i = 298\text{K}$ for all i . Following the previous convention in our growth study [20], the initial distribution is π_1 , and π_N is the targeted distribution. In physical terms, π_1 will represent the distribution of all atoms in non-interacting fragments and π_N will be the fully interacting molecule. Full details of the stages are given below in sec. 2.2.

Our AIS procedure has only a few simple steps and follows our earlier work [26]. The process starts with a well-sampled ensemble of M configurations at stage $i = 1$ in the initial distribution π_1 . AIS does not specify a procedure for sampling π_1 , but assumes it can be accomplished. In our case, the noninteracting fragment ensemble can be sampled almost perfectly using internal coordinate Monte Carlo, as described in refs [20, 27, 28].

The ensemble progresses to the next growth stage by “annealing in” interactions – i.e. “turning on” interactions – between fragments according to the growth pathway shown in Fig. 1 and equations 5. The annealing process shown schematically in Fig. 1 corresponds to the case where we are growing a target molecule composed of smaller non-overlapping fragments $\mathbf{x} = (x_A, x_B, \dots, x_Y)$.

Formally, the M configurations from the current stage i are resampled into the next distribution π_{i+1} based on the weights:

$$w(\mathbf{x}) = \pi_{i+1}(\mathbf{x}) / \pi_i(\mathbf{x}) \quad (3)$$

There are numerous procedures for resampling [48], but here we use the simplest approach of generating M new configurations for ensemble $i + 1$ proportional to the weights from Eq. 3. This approach leads to some higher-weight configurations being duplicated – a fact which is exploited in AIS. After the simple resampling procedure, all weights become equal to one.

Although the resampled set of M configurations for stage $i + 1$ is a statistically valid ensemble for π_{i+1} , it has suffered some “attrition” in quality after growth. Specifically, the uncertainty in calculated observables will be larger than if we had M truly independent configurations – the non-independence is explicit in the duplicated configurations. In AIS, one therefore performs some “relaxation” simulation on each configuration in the ensemble. This can be done using any canonical sampling algorithm, thus preserving the π_{i+1} ensemble

but improving the statistical quality. Our library-based procedure for canonical sampling is described in sec. 2.4. The degree of improvement in ensemble quality depends on the amount of relaxation, a point which we return to later. Nevertheless, after relaxation, a valid ensemble of M configurations in the π_{i+1} ensemble remains. Reweighting and relaxation are repeated through the growth stages until the targeted distribution, π_N , has been sampled.

We summarize our AIS procedure as follows:

- i. Generate an initial distribution of the π_1 ensemble for stage $i = 1$. This is performed by drawing a random configuration from the pre-calculated library for each fragment: see sec. 2.4.
- ii. Resample to the next stage, $i + 1$, by “annealing in” interactions via the weight $w(\mathbf{x}) = \pi_{i+1}(\mathbf{x})/\pi_i(\mathbf{x})$. Our stages are specified in sec. 2.2.
- iii. Relax each configuration via any canonical sampling algorithm, e.g. LBMC, which maintains the π_{i+1} distribution.
- iv. Repeat steps (ii) and (iii) until the target distribution $\pi_N(\mathbf{x})$ has been reached.

2.2 Choice of stages: Progressive addition of interactions

To establish notation, we first divide the full set of coordinates $\mathbf{x}$ into N non-overlapping fragments

$$\mathbf{X} = x_A, x_B, x_C, \dots, x_Y \quad (4)$$

The total energy, $U(\mathbf{x})$, of any fragment-based configuration can be decomposed into two parts. The first contribution is a sum over the energies internal to each fragment (see U_1 below); the second is a sum over energies between interacting fragment pairs (see U_N below).

For a target molecule consisting of N fragments, we employ N intermediate models (stages) such that interactions between fragments are gradually turned on along the growth pathway shown in Fig. 1. The first stage, i.e. the reference state U_1 corresponding to the distribution π_1 , is sampled at the library generation stage and only includes interactions internal to each fragment. Subsequent intermediate stages “anneal in” the indicated interactions among fragments $A, B, C \dots Y$. The energies of the intermediate models can be written recursively as:

$$\begin{aligned} U_1(\mathbf{x}) &= U_A(x_A) + U_B(x_B) + U_C(x_C) + \dots + U_Y(x_Y) \\ U_2(\mathbf{x}) &= U_1(\mathbf{x}) + U_{AB}(x_A, x_B) \\ U_3(\mathbf{x}) &= U_2(\mathbf{x}) + U_{AC}(x_A, x_C) + U_{BC}(x_B, x_C) \\ &\vdots \\ U_N(\mathbf{x}) &= U_{N-1}(\mathbf{x}) + \sum_{\alpha \neq Y} U_{\alpha Y}(x_\alpha, x_Y) \end{aligned} \quad (5)$$

The energy of the last stage, $U_N(\mathbf{x})$, is the full energy of the desired target molecule. The sum extends over interactions between the last fragment Y , with all previous fragments in the molecule.

2.3 Free energy calculation in annealed importance sampling

The free energy of the fully interacting target ensemble relative to the reference state can be expressed in terms of free energy differences between neighboring levels of the annealing ladder:

$$F_N - F_1 = (F_2 - F_1) + (F_3 - F_2) + \dots + (F_N - F_{N-1}) \quad (6)$$

or

$$\Delta F_{1,N} = \Delta F_{1,2} + \Delta F_{2,3} + \dots + \Delta F_{N-1,N} \quad (7)$$

Possession of the ensembles at each level of the annealing ladder directly permits the calculation of free energy differences between levels. If the configuration space is progressively narrowed through the stages and the temperature is constant, as in our case [20], the required values can be obtained simply using

$$\exp(-\beta \Delta F_{i,i+1}) = \langle \exp(-\beta \Delta U_{i,i+1}) \rangle_i \quad (8)$$

where $\Delta U_{i,i+1} = U_{i+1} - U_i$ and the ensemble average is over the configurations from stage i .

Although this relation is sufficient for our studies, in more difficult cases, “two-sided” calculations could be performed - e.g., using the Bennett method [49].

We also point out from Eq. 7 that if the absolute free energy of stage 1, F_1 , is known then the absolute free energy of stage N can be simply found via

$$F_N = F_1 + \Delta F_{1,N} \quad (9)$$

Absolute free energies of the molecular fragments have already been determined in our previous work [20], therefore it is straightforward to convert all free energy differences reported in this paper into their absolute values.

2.4 Library-based Monte Carlo for relaxation

AIS requires a canonical sampling procedure for the “relaxation” process, and we employ library-based Monte Carlo (LBMC) [27, 50]. LBMC is a natural choice because it can employ the same fragments used in the staging choices, and is also highly efficient for sampling implicitly solvated peptides [28]. LBMC uses pre-generated libraries of fragment configurations, echoing extensive work with the Rosetta folding program [51]. LBMC is a canonical sampling procedure which can be used with an arbitrary forcefield and solvent model. Full details regarding LBMC have been given in previous work [27, 50, 28] but we summarize the essentials here.

In simplest terms, LBMC is an ordinary MC procedure which can employ a special fragment-swap trial move: exchange of the configuration of a fragment with a pre-calculated configuration chosen from a “library” or ensemble of pre-calculated configurations. When the library is distributed according to the Boltzmann factor of the target forcefield for all interactions internal to the fragment, the Metropolis criterion [52] is particularly simple:

$$p_{\text{acc}} = \min \left(1, \exp \left[-\beta \Delta U^{\text{rest}} \right] \right) \quad (10)$$

where ΔU^{rest} is the change in fragment-fragment interaction energy due to the trial fragment swap.

More precisely, if one is performing a trial swap move by changing a single fragment configuration $x_J \rightarrow x'_J$, then

$$\Delta U^{\text{rest}} = \left[U(x_A, \dots, x'_j, \dots, x_Y) - U(x_A, \dots, x_j, \dots, x_Y) \right] - \left[U_j(x'_j) - U_j(x_j) \right] \quad (11)$$

The first two energy terms are calculated per trial move. The second two energy terms are simply the energies of the single fragment configurations x_j and x'_j — they are extracted from the pre-calculated libraries.

Many variants of LBMC are possible, but this simple scheme has shown to be successful for flexible all-atom peptides [28]. In particular, all degrees of freedom are included in the libraries – so an amino acid fragment consists of all atomic coordinates plus the six connector degrees of freedom which exactly specify the position and orientation of the next fragment. Swap moves are attempted on configurations drawn uniformly from the Boltzmann distributed libraries.

Libraries for Ace, Ala and Nme are generated by internal-coordinate Monte Carlo, as described in our earlier work [27]. Libraries are distributed according to the $T = 298\text{ K}$ Boltzmann factor of all OPLSAA energy terms internal to each fragment - both bonded and non-bonded. The libraries include additional dummy atoms which encode the six degrees of freedom necessary for positioning a fragment with respect to the previous fragment [27]. The fragments Ace, Ala and Nme contain 6, 10, and 6 atoms respectively. Each fragment library contained 10^5 such distinct configurations and their corresponding energies. Collectively, these libraries occupy approximately 300MB of computer memory. Although smaller libraries probably can be effective, we are still investigating optimal sizes.

For the π_1 distribution of non-interacting fragments, LBMC is not necessary. Rather, the distribution is sampled by drawing a random configuration from the pre-calculated library for each fragment: see U_1 in Eq. 2.

2.5 System and simulation details

We use library-based AIS to calculate free energy changes along the growth pathway in Fig. 1 for four polypeptide systems: Ace-Ala₄-Nme, Ace-Ala₁₂-Nme, Ace-Ala₁₆-Nme, Ace-Ala₂₀-Nme. For all systems under investigation, our libraries implemented the OPLS-AA forcefield [53] with uniform and constant dielectric at constant temperature 298K. The uniform dielectric constant was chosen to be $\epsilon = 60$ and no potential cutoffs are used in the calculation of the energy terms. While other implicit and explicit solvent models are within scope of library-based methods, the goal of this work was to extend library-based free energy calculations[20] and demonstrate that complete sampling and accurate free energy measurements are easily attainable for larger systems.

During each AIS simulation, the free energy change between growth stages i and $i + 1$ is measured according to equation (8) using configurations obtained throughout the relaxation procedure. For the purpose of examining our data, we also calculate intermediate $\Delta F_{i,i+1}$ values as relaxation proceeds; these values are obtained using Eq. (8) for the set of M partially relaxed configurations at various “time” points. The final free energy difference between growth stages i and $i + 1$ is then determined by exponentially averaging all intermediate ΔF values. The total free energy difference between the target and reference system is calculated by summing the exponentially averaged free energy differences for each growth stage, i.e. via equation (7).

For the systems under investigation, we repeat simulations with various amount of relaxation and ensemble sizes to observe the effects on sampling quality and ΔF fluctuations. In principle, these parameters may be adjusted automatically until a desired threshold for free energy accuracy is achieved, however this automation is not implemented

in our current work; see sec. 4.3. The total number of relaxation steps is between $10^7 - 10^9$ LBMC trial moves for each simulation and the steps are distributed evenly over each growth stage. Note that relaxation in the early stages of growth is considerably faster than later stages because there are fewer terms to calculate in Eq.(10). In particular, in order to grow a to grow a single polyalanine chain Ace-Ala_{*n*}-Nme containing *n* alanine residues, only one call to the energy function is made. However, as the peptide chain increases in length, the energy call becomes more computationally expensive due to the fact that the number of pairwise energy terms between atoms increases as 10^{n+1} . Finally, for each simulation, at least 10 repeats have been performed in order to obtain accurate statistics on variations in sampling quality and free energies $\Delta F_{1,N}$.

2.6 Validation method

To validate proper sampling of our systems, we compare the target ensemble of configurations with those obtained from ten independent Langevin Dynamics trajectories. Such a comparison is possible by choosing a strongly discriminating representation of phase space based on a Voronoi construction described below. The ten independent Langevin simulations were performed using TINKER for a collective run time of $1\mu s$. All Langevin simulations were run at $T = 298K$ with a friction constant of 5 psec^{-1} .

We compare the target systems' phase-space distributions with those obtained from well-sampled Langevin Dynamics simulations. Briefly, to obtain a representation of the phase space distribution, we choose five independent and dissimilar reference structures (similarity metric is based on RMSD) from the Langevin Dynamics trajectory of the target system. Configuration space is then partitioned into 5 distinct regions or "bins" based on a Voronoi procedure so that each bin contains all configurations closest to one reference structure. This representation of the phase-space distribution provides an extremely sensitive test which is not always "passed" as seen in the next section. Full details for this procedure can be found in ref. [54, 55, 56, 20].

3 Validation and Results

3.1 Sampling Validation

First, to check whether sufficient sampling has been performed, Fig. 2 plots the configuration-space distributions mentioned in sec. 2.6 for the four systems examined. In each plot, we compare the distributions resulting library-based AIS and Langevin Dynamics simulations. Error bars have a width of two standard deviations and are based on results obtained from at least 10 independent simulations of each method. In all cases, there is good agreement between both methods, validating the free energy measurements and sampling capacity of this method. Note that although the sampling error bars are large for the Ace-Ala₂₀-Nme system, the free energy standard deviation is still reasonably small (0.39 kcal/mol) as described below.

For reference, in Fig. 3 we plot the configuration-space distributions for the peptide Ace-Ala₁₂-Nme as obtained from three methods: pure LBMC of the full system (no staging), pure growth as in ref. [20] (no relaxation), and Langevin dynamics simulations. The data underscores the fact that pure library-based growth simulations are unable to sample these larger systems. However, by implementing relaxation combined with growth (i.e. AIS), we are able to recover the correct equilibrium distribution – compare Fig. 3 and Fig. 2b.

3.2 Free Energy Measurements and Statistics

To assess the free energy estimates, we report the mean and standard deviation of the free energy for each polypeptide relative to the non-interacting reference state for varying

amounts of relaxation and ensemble sizes as shown in Table 1. The statistics are based on at least 10 independent simulations for each set of parameters. The table indicates the computing time required for different levels of precision, although further optimization may be possible (sec. 4.3).

The principal result embodied in Table 1 is that for polyaniline systems up to 16 residues, only a couple of hours is required to reach a level of precision comparable to the accuracy of forcefields (~ 0.5 kcal/mol) [2]. It can be seen from Table 1 that free energy variances can be decreased by increasing the overall amount of relaxation. Importantly, however, increasing the ensemble size implies that each configuration will receive less relaxation if the simulations are to be run in equal amounts of time. Although the sampling error bars are large for the Ace-Ala₂₀-Nme system (e.g Fig. 2(d)), the free energy estimate is still reasonably precise, with a standard deviation of 0.39 kcal/mol.

The series of constant-time simulations for Ace-Ala₁₂-Nme in Table 1 indicates that decreasing the ensemble size by a factor of 10 is roughly equivalent to increasing the number of relaxation steps by the same factor insofar as free energy variances are concerned.

A representative plot of intermediate ΔF measurements during the relaxation of each growth stage for the Ace-Ala₁₂-Nme system is shown in Fig. 4 along with the exponentially averaged results shown in the inset for each growth stage. In Fig. 4, the sharp decrease in $\Delta F_{i \rightarrow i+1}$ at growth stage $i = 6$ is attributed to the fact that, on average, configurations in the ensemble become long enough so that more contacts are formed, e.g. H-bonds and steric clashes. The LBMC acceptance rate (not shown) also follows a similar trend since there is a larger chance of steric overlap when more atoms are interacting.

We investigated the issue of the choice of ensemble size (M) for a given computing investment. Thus, for Ace-Ala₁₂-Nme, four separate runs were performed, each using a total of 10^8 relaxation steps as shown in Table 1. The standard deviations, σ , suggest $M \sim 10^3$ is close to optimal for this system.

4 Discussion

4.1 Sampling quality and free energy precision

What kind of sampling quality is required for reasonable precision (0.5 kcal/mol) in free energy estimates? To address this issue, it is interesting to compare results obtained in cases of good and poor sampling. We have measured the free energy $\Delta F_{1,N}$ using pure growth simulations (no relaxation) for the four polypeptide systems examined, see Table 1. Pure growth simulations, as mentioned previously, will under-sample configuration space for these larger systems – see Fig. 3 for example. In the case of the smallest system we examined, Ace-Ala₄-Nme, relaxation makes very little difference in sampling quality since pure growth simulations are able to accurately predict the correct equilibrium distributions and free energies[20]. However, for the larger systems Ace-Ala₁₂-Nme and Ace-Ala₁₆-Nme, the difference in mean $\Delta F_{1,N}$ values as obtained from poorly-sampled (e.g., Fig. 3) and well-sampled simulations (e.g, Fig. 2b) are 1.18 and 1.62 kcal/mol respectively. This underscores the fact that seemingly small differences in free energies could mask poorly-sampled ensembles. Such differences would be expected to increase in more complex systems.

The coupling of sampling and free energy calculation suggests one can examine such approaches purely in terms of their ability to provide equilibrium sampling. To set this in context, note that we recently showed that a fragment-based polymer growth strategy without annealing/relaxation could sample equilibrium ensembles of all-atom peptides but

was limited to about five residues [20]. The present study, by adding a relaxation phase, greatly extends the size of peptides which can be sampled.

4.2 Limitations of the relaxed growth method

The essential limitation of our approach, not surprisingly, is sampling. It is not a coincidence that our present implementation becomes dramatically more expensive when studying systems beyond the efficient range of current LBMC simulation, roughly 16–20 residues for polyalanine. Other sampling methods, or improved versions of LBMC, likely could extend the system sizes amenable for free energy estimation. The basic requirements, as is implicit in Eq. (8) and the subsequent discussion, is to be able to sample the ensembles at every stage to a sufficient degree to permit calculation of the required free energies. Improved staging, possibly more incremental, could also be useful for larger systems.

4.3 Possible improvements

Insofar as annealed importance sampling is a formal method based on (i) arbitrary stages and (ii) an arbitrary (correct) sampling method, improvements in either of those components will improve free energy calculations. Sampling itself could be improved in a number of ways: with LBMC variants better optimized for structured (e.g., folded) systems; with alternative sampling algorithms; or with other hardware optimizations, such as based on graphics processing units (GPUs), for LBMC or for more traditional algorithms [37]. The relaxation of M configurations at every stage may be amenable to CPU or GPU parallelization.

Staging could also be improved, as is generally the case in free energy calculations [57, 58]. Our current procedure adds all interactions for a newly added fragment in a single stage - but this becomes a significant perturbation as system size increases. The interaction-space staging scheme naturally permits “sub-staging” of added interactions, and this strategy will be explored in future work.

However, even for a fixed set of stages and a specified sampling algorithm, significant further optimization may be possible. For instance, as Fig. 4 illustrates graphically, some stages are much easier than others - i.e., have greater overlap. This suggests that the amount of relaxation per stage could be adjusted on the fly based on a convergence criterion, perhaps for a block-averaged variance [59] of free energy estimates. The size of the ensemble, M , could also be optimized in a system and stage-specific way. These issues will be taken up in future studies.

4.4 Future applications to protein affinity estimates

The resources expended in this study were quite modest, and suggest significantly larger systems could be addressed with more computing. Yet we do not believe that the present methodology would permit the growth and sampling of all-atom proteins. What are the prospects for estimating binding affinities?

Importantly, the estimation of binding affinities would not require the growth of a full protein. Rather, the affinity a ligand for a receptor would be calculated from the difference of free energy values calculated by growing the ligand into solvent and into the receptor binding site. In other words, only the ligand needs to be grown, and typical ligands are smaller than the peptides studied in this report. Such a growth procedure adds interactions, and can be seen as the inverse of a decoupling procedure [60, 61]. At the same time, free energy values require good sampling of the full system, which is never easy for proteins. In this regard, LBMC was designed to handle hybrid models in a natural way - with an atomistic binding site and a reduced representation elsewhere [50]. Good sampling of such

models via LBMC may permit rapid, statistically based affinity estimates within the context of hybrid models. Such a strategy echoes the work of Roux and coworkers [62, 63, 64, 65], but hybrid models would permit potentially important allosteric coupling to regions distant from the binding site. Another strategy based on fragments has been proposed [22], but it does not account for flexibility internal to fragments.

4.5 Improvements in implementation compared to previous work

By comparison with previous a previous study by our group [20], the reported computing times in Table 1 may seem incongruously fast. Our earlier study examined a series of peptides, with the largest being tetra-alanine, noting that 50 minutes of computing time were required to obtain a precision of 0.29 kcal/mol for tetra-alanine. By contrast, in our present work, the same “pure growth” calculation (no relaxation) required 14 s. The improvement results primarily from the fact that our previous work was a scripting based post-analysis of data generated by LBMC code, whereas we now calculate free energy values directly within the LBMC code. Additionally, the growth pathway implemented in this study differs slightly from that in ref. [20]. Our current growth pathway (Fig. 1) adds new fragments in a more efficient way, increasing overlap between neighboring growth stages. However the number of stages and energy calls is the same, with no relaxation.

5 Summary and Conclusions

We reported free energy calculations for implicitly solvated polyalanine peptides, ranging in size from four to 20 residues. The calculations combine two previously developed techniques, a fragment-based polymer growth strategy [20] and library-based Monte Carlo simulation [27, 28]. Our new implementation greatly extends the system sizes amenable to free energy estimation compared to a previous study by our group [20]. Because the calculations for the peptides are so inexpensive, we hope the approach can be useful for protein-ligand affinity estimation in the future, as described in sec. 4.4. The present calculations required seconds to days of single-CPU computing, depending on system size and required precision.

The results here are another application of the memory-intensive strategy of using pre-calculated libraries of molecular-fragment configurations [51, 27, 28]. That strategy has been useful for rapid sampling of semi-atomistic protein models [27] and of implicitly solvated peptides [28]. Libraries of configurations have previously been applied extensively in the Rosetta protein folding software [51] albeit not for canonical sampling. The potential for ongoing improvements in memory size and access speed, orthogonal to CPU speed, suggests the value of continued pursuit of memory-intensive computations.

From a formal point of view, we have shown that one can perform annealing in interaction space. That is, our approach is formally equivalent to the annealed importance sampling strategy described by Neal [23], except that instead of lowering temperature, we add interactions among molecular fragments.

Acknowledgments

We greatly appreciate insightful discussions with Ying Ding, Divesh Bhatt and Xin Zhang, as well as financial support from the NIH (Grants GM076569 and GM070987) and NSF(Grant MCB-0643456).

References

1. Gilson MK, Zhou HX. Calculation of protein-ligand binding affinities. Annual review of biophysics and biomolecular structure. 2007; 36:21.

2. Shirts MR, Pitera JW, Swope WC, Pande VS. Extremely precise free energy calculations of amino acid side chain analogs: Comparison of common molecular mechanics force fields for proteins. *The Journal of Chemical Physics*. 2003; 119:5740.
3. Frenkel D, Ladd AJC. New Monte Carlo method to compute the free energy of arbitrary solids. Application to the fcc and hcp phases of hard spheres. *The Journal of Chemical Physics*. 1984; 81:3188.
4. Stoessel JP, Nowak P. Absolute free energies in biomolecular systems. *Macromolecules*. 1990; 23(7):1961–1965.
5. Head MS, Given JA, Gilson MK. Mining Minima: Direct Computation of Conformational Free Energy. *J. Phys. Chem. A*. 1997; 101(8):1609–1618.
6. Chang CE, Gilson MK. Free Energy, Entropy, and Induced Fit in Host- Guest Recognition: Calculations with the Second-Generation Mining Minima Algorithm. *J. Am. Chem. Soc.* 2004; 126(40):13156–13164. [PubMed: 15469315]
7. Moghaddam S, Inoue Y, Gilson MK. Host- Guest Complexes with Protein- Ligand-like Affinities: Computational Analysis and Design. *J. Am. Chem. Soc.* 2009; 131(11):4012–4021. [PubMed: 19133781]
8. Zhou HX, Gilson MK. Theory of free energy and entropy in noncovalent binding. *Chem. Rev.* 2009; 109(9):4092–4107. [PubMed: 19588959]
9. Hnizdo V, Tan J, Killian BJ, Gilson MK. Efficient calculation of configurational entropy from molecular simulations by combining the mutual-information expansion and nearest-neighbor methods. *Journal of computational chemistry*. 2008; 29(10):1605–1614. [PubMed: 18293293]
10. Meirovitch H. Simulation of a free energy upper bound, based on the anticorrelation between an approximate free energy functional and its fluctuation. *The Journal of Chemical Physics*. 1999; 111:7215.
11. White RP, Meirovitch H. A simulation method for calculating the absolute entropy and free energy of fluids: Application to liquid argon and water. *Proceedings of the National Academy of Sciences of the United States of America*. 2004; 101(25):9235. [PubMed: 15197270]
12. Cheluvaraja S, Meirovitch H. Simulation method for calculating the entropy and free energy of peptides and proteins. *Proceedings of the National Academy of Sciences of the United States of America*. 2004; 101(25):9241. [PubMed: 15197271]
13. Meirovitch H. Entropy, pressure, and chemical potential of multiple chain systems from computer simulation. I. Application of the scanning method. *The Journal of Chemical Physics*. 1992; 97:5803.
14. Meirovitch H, Cheluvaraja S, White RP. Methods for calculating the entropy and free energy and their application to problems involving protein flexibility and ligand binding. *Current protein & peptide science*. 2009; 10(3):229. [PubMed: 19519453]
15. Meirovitch H. Methods for calculating the absolute entropy and free energy of biological systems based on ideas from polymer physics. *Journal of Molecular Recognition*. 2010; 23(2):153–172. [PubMed: 19650071]
16. Cheluvaraja S, Meirovitch H. Calculation of the entropy and free energy of peptides by molecular dynamics simulations using the hypothetical scanning molecular dynamics method. *The Journal of chemical physics*. 2006; 125:024905.
17. Banba S, Guo Z, Brooks CL III. Efficient Sampling of Ligand Orientations and Conformations in Free Energy Calculations Using the [lambda]-Dynamics Method. *J. Phys. Chem. B*. 2000; 104(29):6903–6910.
18. Guo Z, Brooks CL III, Kong X. Efficient and Flexible Algorithm for Free Energy Calculations Using the [lambda]-Dynamics Approach. *J. Phys. Chem. B*. 1998; 102(11):2032–2036.
19. Kong XJ, Brooks CL. Lambda-dynamics: A new approach to free energy calculations. *J. Chem. Phys.* 1996; 105(6):24142423.
20. Zhang X, Mamonov AB, Zuckerman DM. Absolute free energies estimated by combining precalculated molecular fragment libraries. *Journal of Computational Chemistry*. 2009; 30(11): 1680–1691. [PubMed: 19504588]
21. Ytreberg FM, Zuckerman DM. Simple estimation of absolute free energies for biomolecules. *The Journal of Chemical Physics*. 2006; 124:104105. [PubMed: 16542066]

22. Clark M, Meshkat S, Talbot GT, Carnevali P, Wiseman JS. Fragment-Based Computation of Binding Free Energies by Systematic Sampling. *J. Chem. Inf. Model.* 2009; 49(8):1901–1913. [PubMed: 19610599]
23. Neal RM. Annealed importance sampling. *Statistics and Computing.* 2001; 11(2):125–139.
24. Gary A. Huber and J. Andrew McCammon. Weighted-ensemble simulated annealing: Faster optimization on hierarchical energy surfaces. *Phys. Rev. E.* Apr; 1997 55(4):4822–4825.
25. Lyman E, Zuckerman DM. Annealed importance sampling of peptides. *The Journal of chemical physics.* 2007; 127:065101. [PubMed: 17705625]
26. Lyman E, Zuckerman DM. Resampling improves the efficiency of a fast-switch equilibrium sampling protocol. *The Journal of chemical physics.* 2009; 130:081102. [PubMed: 19256587]
27. Mamonov AB, Bhatt D, Cashman DJ, Ding Y, Zuckerman DM. General Library-Based Monte Carlo Technique Enables Equilibrium Sampling of Semi-atomistic Protein Models. *J. Phys. Chem. B.* 2009; 113(31):10891–10904. [PubMed: 19594147]
28. Ding Y, Mamonov AB, Zuckerman DM. Efficient equilibrium sampling of all-atom peptides using library-based Monte Carlo. *J. Phys. Chem. B.* 2010; 114(17):5870–5877. [PubMed: 20380366]
29. Swendsen RH, Wang JS. Replica Monte Carlo simulation of spin-glasses. *Physical Review Letters.* 1986; 57(21):2607–2609. [PubMed: 10033814]
30. Geyer CJ. Markov Chain Monte Carlo Maximum Likelihood. 1991:156.
31. Berg BA, Neuhaus T. Multicanonical ensemble: A new approach to simulate first-order phase transitions. *Physical Review Letters.* 1992; 68(1):9–12. [PubMed: 10045099]
32. Okamoto Y. Generalized-ensemble algorithms: enhanced sampling techniques for Monte Carlo and molecular dynamics simulations. *Journal of Molecular Graphics and Modelling.* 2004; 22(5): 425–439. [PubMed: 15099838]
33. Ifitimie R, Salahub D, Wei D, Schofield J. Using a classical potential as an efficient importance function for sampling from an ab initio potential. *The Journal of Chemical Physics.* 2000; 113:4852.
34. Gelb LD. Monte Carlo simulations using sampling from an approximate potential. *The Journal of Chemical Physics.* 2003; 118:7747.
35. Hetenyi B, Bernacki K, Berne BJ. Multiple time step Monte Carlo. *The Journal of Chemical Physics.* 2002; 117:8203.
36. Lyman E, Ytreberg FM, Zuckerman DM. Resolution exchange simulation. *Physical review letters.* 2006; 96(2):28105.
37. Friedrichs MS, Eastman P, Vaidyanathan V, Houston M, Legrand S, Beberg AL, Ensign DL, Bruns CM, Pande VS. Accelerating molecular dynamic simulation on graphics processing units. *Journal of computational chemistry.* 2009; 30(6):864–872. [PubMed: 19191337]
38. Phillips JC, Braun R, Wang W, Gumbart J, Tajkhorshid E, Villa E, Chipot C, Skeel RD, Kale L, Schulten K. Scalable molecular dynamics with NAMD. *Journal of computational chemistry.* 2005; 26(16):1781–1802. [PubMed: 16222654]
39. Freddolino PL, Liu F, Gruebele M, Schulten K. Ten-microsecond molecular dynamics simulation of a fast-folding WW domain. *Biophysical journal.* 2008; 94(10):L75–L77. [PubMed: 18339748]
40. Kalé 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* 1. *Journal of Computational Physics.* 1999; 151(1):283–312.
41. Voelz VA, Bowman GR, Beauchamp K, Pande VS. Molecular Simulation of ab Initio Protein Folding for a Millisecond Folder NTL9 (1–39). *J. Am. Chem. Soc.* 2010; 132(5):1526–1528. [PubMed: 20070076]
42. Bitetti-Putzer R, Yang W, Karplus M. Generalized ensembles serve to improve the convergence of free energy simulations. *Chemical Physics Letters.* 2003; 377(5-6):633–641.
43. Tidor B. Simulated annealing on free energy surfaces by a combined molecular dynamics and Monte Carlo approach. *The Journal of Physical Chemistry.* 1993; 97(5):1069–1073.
44. Zheng L, Chen M, Yang W. Random walk in orthogonal space to achieve efficient free-energy simulation of complex systems. *Proceedings of the National Academy of Sciences.* 2008; 105(51): 20227.

45. Marc Fasnacht, Robert H. Swendsen; Rosenberg, John M. Adaptive integration method for monte carlo simulations. *Phys. Rev. E.* May.2004 69(5):056704.
46. Li H, Fajer M, Yang W. Simulated scaling method for localized enhanced sampling and simultaneous alchemical free energy simulations: A general method for molecular mechanical, quantum mechanical, and quantum mechanical/molecular mechanical simulations. *The Journal of chemical physics.* 2007; 126:024106. [PubMed: 17228942]
47. Knight JL, Brooks CL. Lambda-dynamics free energy simulation method. *J. Comput. Chem.* 2009; 30:1692–1700. [PubMed: 19421993]
48. Liu, JS. Monte Carlo strategies in scientific computing. Springer Verlag; 2008.
49. Bennett CH. Efficient estimation of free energy differences from Monte Carlo data. *Journal of Computational Physics.* 1976; 22(2):245–268.
50. Mamonov AB, Zhang X, Zuckerman DM. Rapid equilibrium sampling of all-atom peptides using a library-based polymer-growth approach. Arxiv preprint arXiv:0910.2495. 2010 Recently accepted to *J. Comput. Chem.*
51. Rohl CA, Strauss CEM, Misura K, Baker D. Protein structure prediction using Rosetta. *Methods in enzymology.* 2004; 383:66–93. [PubMed: 15063647]
52. Meteopolis N, Ulam S. The monte carlo method. *Journal of the American Statistical Association.* 1949; 44(247):335–341. [PubMed: 18139350]
53. Jorgensen WL, Maxwell DS, Tirado-Rives J. Development and testing of the OPLS all-atom force field on conformational energetics and properties of organic liquids. *J. Am. Chem. Soc.* 1996; 118(45):11225–11236.
54. Lyman E, Zuckerman DM. Ensemble-based convergence analysis of biomolecular trajectories. *Biophysical journal.* 2006; 91(1):164–172. [PubMed: 16617086]
55. Lyman E, Zuckerman DM. On the structural convergence of biomolecular simulations by determination of the effective sample size. *J. Phys. Chem. B.* 2007; 111(44):12876–12882. [PubMed: 17935314]
56. Voronoi G. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Premier mémoire. Sur quelques propriétés des formes quadratiques positives parfaites. *Journal für die reine und angewandte Mathematik (Crelle's Journal).* 1908; 1908(133):97–102.
57. Kofke DA, Cummings PT. Quantitative comparison and optimization of methods for evaluating the chemical potential by molecular simulation. *Molecular Physics.* 1997; 92(6):973–996.
58. Kofke DA, Cummings PT. Precision and accuracy of staged free-energy perturbation methods for computing the chemical potential by molecular simulation. *Fluid Phase Equilibria.* 1998; 150:41–49.
59. Grossfield A, Zuckerman DM. Quantifying Uncertainty and Sampling Quality in Biomolecular Simulations. *Annual reports in computational chemistry.* 2009; 5:23–48. [PubMed: 20454547]
60. Hamelberg D, McCammon JA. Standard free energy of releasing a localized water molecule from the binding pockets of proteins: double-decoupling method. *J. Am. Chem. Soc.* 2004; 126(24): 7683–7689. [PubMed: 15198616]
61. Boresch S, Tettinger F, Leitgeb M, Karplus M. Absolute binding free energies: a quantitative approach for their calculation. *J. Phys. Chem. B.* 2003; 107(35):9535–9551.
62. Beglov D, Roux B. Finite representation of an infinite bulk system: solvent boundary potential for computer simulations. *Journal of Chemical Physics.* 1994; 100(12):9050–9063.
63. Wang J, Deng Y, Roux B. Absolute binding free energy calculations using molecular dynamics simulations with restraining potentials. *Biophysical journal.* 2006; 91(8):2798–2814. [PubMed: 16844742]
64. Deng Y, Roux B. Computation of binding free energy with molecular dynamics and grand canonical Monte Carlo simulations. *The Journal of chemical physics.* 2008; 128:115103. [PubMed: 18361618]
65. Deng Y, Roux B. Computations of standard binding free energies with molecular dynamics simulations. *J. Phys. Chem. B.* 2009; 113(8):2234–2246. [PubMed: 19146384]

Figure 1.

A general schematic of the growth and relaxation simulation procedure. Each box represents an ensemble of fragments. The process starts with non-interacting fragments (e.g. the amino acids) which are then used to “grow” the full molecule. At each intermediate stage, the next fragment in the sequence is added by “turning on” the additional interactions due to that fragment. Free energy changes, ΔF , are computed and accumulated at each growth stage. This procedure is repeated until the fully interacting molecule is obtained at stage N . The grey colored blocks represent ensembles which have been relaxed according to library-based Monte Carlo.

Figure 2.

Comparison of the equilibrium distributions obtained using both AIS and Langevin Dynamics simulations for the four systems under investigation. Correct equilibrium sampling is confirmed for systems a), b) and c) using the new method. System d) has not reached the precision of the 100 nsec Langevin simulations. Error bars for both of these methods have a width of two standard deviations and are a result of statistics obtained from at least 10 independent simulations.

Figure 3.

Comparison of equilibrium distributions obtained by growth, LBM and standard Langevin Dynamics simulations. The plot shows the fractional population in different “bins” of phase space for the peptide Ace-Ala₁₂-Nme. The error bars have a width of 2 standard deviations and are based on data from 10 independent simulations. The lines drawn between data points are a guide to the eye. For a system of this size, a simple growth procedure without relaxation fails to produce the correct distribution.

Figure 4. Intermediate free energy changes measured during the relaxation process for each growth stage for Ace-Ala₁₂-Nme. Each data point represents Eq. (8) applied to a subset of the data – different colors are displayed to more clearly distinguish each growth stage. Shown in the inset is the exponentially averaged final free energy difference that is obtained from averaging all data via. Eq. (8) at each growth stage. Lines between data points are drawn to guide the eye. These are representative results from a single simulation.

Table 1

Free energy statistics. The mean values (μ) and standard deviations (σ) in kcal/mol for the four peptide systems are tabulated as a function of ensemble size and amount of relaxation. Statistics are based on at least ten independent simulations.

System	Ensemble Size (M)	Tot no. Steps	$\mu(\Delta F_{L,N})$	$\sigma(\Delta F_{L,N})$	CPU time
Ace-Ala ₄ -Nme	10 ³	no relaxation/pure growth	-5.581	0.043	14 sec.
	10 ³	10 ⁷	-5.562	0.001	7.8 min.
Ace-Ala ₁₂ -Nme	10 ⁴	no relaxation/pure growth	-20.19	0.985	20 sec.
	10 ¹	10 ⁷	-21.22	0.299	1 hr.
	10 ¹	10 ⁸	-21.21	0.235	8.6 hrs.
	10 ²	10 ⁸	-21.38	0.223	8.6 hrs.
	10 ³	10 ⁸	-21.24	0.169	8.6 hrs.
	10 ⁴	10 ⁸	-21.28	0.191	8.6 hrs.
	10 ⁴	10 ⁹	-21.37	0.057	3.5 days
Ace-Ala ₁₆ -Nme	10 ⁴	no relaxation/pure growth	-29.78	1.191	35 sec.
	10 ¹	10 ⁷	-30.78	0.459	1.9 hrs.
	10 ²	10 ⁸	-31.02	0.396	16.6 hrs.
	10 ⁴	10 ⁹	-31.40	0.179	7 days
Ace-Ala ₂₀ -Nme	10 ³	no relaxation/pure growth	-38.99	2.101	13 sec.
	10 ¹	10 ⁸	-41.66	0.870	1.3 days
	10 ³	10 ⁹	-42.45	0.394	14 days