

Published in final edited form as:

J Comput Chem. 2009 August ; 30(11): 1680–1691. doi:10.1002/jcc.21337.

Absolute free energies estimated by combining pre-calculated molecular fragment libraries

Xin Zhang^{*},

Department of Physics & Astronomy, University of Pittsburgh, Pittsburgh, Pennsylvania 15260

Artem B. Mamonov^{*}, and

Department of Computational Biology, School of Medicine, University of Pittsburgh, Pittsburgh, Pennsylvania 15213

Daniel M. Zuckerman[†]

Department of Computational Biology, School of Medicine, University of Pittsburgh, Pittsburgh, Pennsylvania 15213

Abstract

The absolute free energy — or partition function, equivalently — of a molecule can be estimated computationally using a suitable reference system. Here, we demonstrate a practical method for staging such calculations by growing a molecule based on a series of fragments. Significant computer time is saved by pre-calculating fragment configurations and interactions for re-use in a variety of molecules. We employ such fragment libraries and interaction tables for amino acids and capping groups to estimate free energies for small peptides. Equilibrium ensembles for the molecules are generated at no additional computational cost, and are used to check our results by comparison to standard dynamics simulation. We explain how our work can be extended to estimate relative binding affinities.

I. INTRODUCTION

The use of a reference system for free energy calculations has a long history in physics and chemistry.¹ The basic idea is to employ a reference system (“ref”) for which the absolute free energy is available, and which is as similar as possible to the physical system of interest (“phys”):

$$F^{\text{phys}} = F^{\text{ref}} + \Delta F_{\text{ref} \rightarrow \text{phys}}, \quad (1)$$

where F^x is the absolute free energy of model “x,” and $\Delta F_{\text{ref} \rightarrow \text{phys}}$ is the free energy difference between the systems. In essence, this paper is about practical choices for the both the reference system and strategies for calculating $\Delta F_{\text{ref} \rightarrow \text{phys}}$ when the physical system is a large molecule.

Historically, Stoessel and Nowak applied the reference-system strategy to a molecular system for the first time, using a solid harmonic reference system in Cartesian coordinates.² Zuckerman and Ytreberg extended that work in two ways designed to improve overlap between the reference and physical systems:³ (i) by using internal coordinates; and (ii) by using a more flexible, numerically exact reference system based on histograms from a short

[†]ddmmzz+@pitt.edu.

^{*}These authors contributed equally.

dynamics simulation, rather than an artificial analytically tractable reference state. Huang and Makarov also employed the reference-system approach embodied in (1), but in a different way.⁴

Other efforts to calculate absolute free energies for molecular systems have been ongoing for years in the groups of Meirovitch^{5–8} and Gilson^{9–12} and more approximately using harmonic and quasi-harmonic methods.¹³ The work of Meirovitch builds on long-standing polymer-growth methodologies for estimating partition functions which date to the work of Rosenbluth and Rosenbluth.¹⁴ The original Rosenbluth work was generalized for higher efficiency and more realistic models by many workers.^{15–27} Ideas from these polymer-growth sampling methods also inform the present work, in the sense that each of our fragments is a kind of “monomer”.

The present paper significantly extends the previous work by Ytreberg and Zuckerman³ by estimating absolute free energies for molecules built up gradually from molecular fragments. Larger molecules can be treated, compared to our previous paper³, because a series of staged intermediate systems are adopted. In essence the free energy difference of Eq. (1) is sub-divided (staged) into a sum of easy-to-calculate terms. Staging increments are highly tunable, based on the choice of fragment sizes and even by selection of subsets of interactions as detailed below. The use of fragments in other types of molecular mechanics calculations has a long history.^{28–30}

A key contribution of this work is a practical strategy of pre-calculation which minimizes the number of energy terms which need to be computed at each stage. Specifically, for each fragment, a statistical library — i.e., an ensemble of configurations and their energies — is stored; we have also used such libraries in Monte Carlo sampling³¹. Additionally, for each covalently bonded fragment pair, we store the full interaction energy (based on all atoms) for every possible pair of configurations. Such storage is quite practical on typical modern computers with > 1 GB of RAM. During production simulations it is only necessary to compute interactions between fragments separated by one or more other fragments. Needless to say, the stored libraries and interaction tables can be re-used in future simulations of the same or different molecules. The pre-calculation strategy, which has early conceptual roots^{32–34}, appears to represent a significant practical advance over earlier polymer-growth calculations. The use of non-statistical libraries has been popularized in the Rosetta protein folding program.³⁵

There is substantial flexibility in the division of a molecule into fragments. We have used single amino acids as fragments in this study, but larger segments and even different interaction subsets as detailed below – may also be practical. The fragment-based approach could also be used to study protein-ligand binding, by growing small molecules into receptor binding pockets and estimating the free energies. This can be seen as a statistical mechanical generalization of fragment-based ideas developed earlier.^{29,30}

Our results, which employ single amino-acid fragments, are extremely encouraging. The data indicate that absolute free energies for small peptides can be calculated rapidly and reliably. Specifically, high-precision free energy estimates, with fluctuations of ~ 0.3 kcal/mole, are obtained for 52-atom tetra-alanine in less than an hour of single-processor computing time, with a simple dielectric “solvent”. We check our data by comparing the equilibrium ensembles (obtained simultaneously with the free energy estimates) with independent Langevin simulations. As a further check, in one case, the free energy results are verified by an independent calculation using different fragments.

The remaining sections of the paper describe the methods, results, and our conclusions. Our methods section provides full details for performing the calculations, including the

generation of fragment libraries and interaction tables. We also correct a technical error in our earlier study.³ Our results describe both the free energy values and the analysis of the equilibrium ensembles. Our discussion section describes possible improvements to the method and extension to the estimation of relative binding affinities using absolute free energies.

II. METHODS

Our basic approach is to calculate the free energy of the physical system of interest based on the difference from a known reference system, as in Eq. (1), and also to stage the calculation using molecular fragments. The fragments not only permit the gradual staging of the calculation but also a tremendous savings of computer time based on the storage of (i) fragment configurations, (ii) energies internal to each fragment configuration, and (iii) interaction energies between covalently bonded fragments. The low cost and high precision of the resulting estimates suggests we are far from the practical limit of the approach in the present implementation. However, a number of improvements to the implementation appear to be within easy reach, as described in our Discussion. All fragment libraries used in the present calculations are available at our website (www.ccbb.pitt.edu/Zuckerman).

A. Model and systems

All calculations employ a standard atomistic forcefield, OPLS-AA36 at $T = 298K$. In the present report, our fragments are individual amino acids and capping groups. For simplicity in this initial investigation, we model the solvent by a simple uniform dielectric constant $\epsilon = 60$. We compute free energy estimates for alanine dipeptide (Ace-Ala-Nme), di-alanine (Ace-(Ala)₂-Nme), and tetra-alanine (Ace-(Ala)₄-Nme). Following standard conventions, Ace is Acetyl (CH₃-CO), Ala is Alanine (NH-CH(CH₃)-CO), and Nme is N-methylamide (NH-CH₃).

B. A simple example

Consider the calculation of the configurational free energy of alanine dipeptide based on a division into three fragments (Ace, Ala, Nme) which can be denoted (A, B, C) respectively (see Fig. 1). We will discuss this example qualitatively, to give the reader a sense of what to expect. Full details and all equations will be presented in succeeding subsections.

In advance, we calculate statistical libraries of configurations for each fragment, which are constant-temperature OPLS-AA ensembles based only on the atoms within the given fragment. The libraries additionally include the six degrees of freedom necessary for joining the fragments, based on the use of “dummy atoms” as described below. During the library generation process, the absolute free energy for each fragment is also calculated using a reference system based on independent internal coordinates, as described previously.³ A typical library will contain 10,000 configurations. We also pre-calculate every possible interaction energy between covalently bound fragments — i.e., a table of 10^8 interaction energies for the A-B and B-C fragment pairs.

The calculation proceeds as schematized in Fig. 1, where the presence of a line connecting two fragments indicates that all interactions between the fragments is included. The reference system (not shown) consists of fully independent coordinates, so that the fragments are not yet constructed. The first intermediate consists of the three non-interacting fragments, which include, however, all interactions *within* each fragment. Thus the fragment free energies, which are calculated and stored in advance, properly include the interactions among all degrees of freedom *internal* to each fragment. Other interactions are added in three stages: A-B interactions first, followed by B-C, and completed by A-C couplings.

In the first intermediate stage, the absolute free energies for the individual fragments are retrieved from disk. (They are initially calculated following reference³ as detailed below.) Next, A-B interactions are added by a standard free energy difference calculation. Specifically, an ensemble of non-interacting A-B configurations is generated by random combination of fragments from the A and B libraries, and the resulting energy change is exponentially averaged in the usual way — via Eq. (6) below. The energy changes due to the combination do not need to be calculated in our scheme, however, because they have been tabulated in advance. Additionally, the now interacting A-B fragments are “resampled”³⁷ to correspond to the full potential for all degrees of freedom in *both* fragments. The details of resampling are given below — see Eq. (8) — but the bottom line is that one obtains 10,000-configuration ensemble of the partially grown molecule consisting of the A-B fragments.

The calculations then proceeds as if there are two fragments, A-B and C. The two libraries are joined combinatorially *but only accounting for the B-C interactions at this stage*. The A-C interactions will be handled at a later stage. Once again, the free energy change is calculated and the ensemble is resampled to reflect B-C interactions. The resulting ensemble contains 10,000 configurations of the full molecule reflecting all interactions except those between fragments A and C.

In the final stage sketched in Fig. 1, the A-C interactions are added in a standard free energy difference calculation based on the the ensemble of the previous stage. However, a standard energy call (for the whole molecule) is not required to save CPU time. Rather, our code only computes energy terms specific to the A and C fragments — i.e., electrostatic and van der Waals interactions between atoms of A and those of C. Once the energy changes have been obtained, the full free energy is rapidly estimated. Resampling into the fully interacting ensemble can also be performed rapidly without additional energy calculations.

It is not difficult to imagine generalizing this example to systems with more fragments. It is also worth noting that, strictly speaking, the final stage was not necessary. That is, we could have added the A-C interactions simultaneously with the B-C combination since the full molecular configurations were constructed at that point. These choices illustrate the flexibility intrinsic to staging the calculation with fragments, as we detail further in our Discussion. Additional staging flexibility results, of course, from the initial choice of the fragments — i.e., smaller fragments lead to staging in finer increments.

C. Basic formalism

Our calculation of the absolute free energy F^{phys} for a molecule divided into fragments is based on standard, straightforward equations. The only novel aspect of the formalism is our particular choice of stages based on the addition of fragments and/or inter-fragment interactions. Although our heavy reliance on pre-calculated information has very significant practical implications, it does not affect the formalism.

Definition of the free energy—The fundamental object of interest is the absolute classical free energy F^{phys} for an implicitly solvated molecule. The molecule is taken to consist of N atoms, and its *internal*-coordinate configurations are denoted by $\mathbf{x} = (x_1, x_2, \dots, x_{3N-6})$. The potential energy function will be a standard forcefield (here, OPLS-AA36), possibly augmented by an implicit solvation model; the full potential energy including any solvation will be denoted by $U^{\text{phys}}(\mathbf{x})$. The free energy, which is a functional of U^{phys} , is defined by the dimensionless *configurational* partition function at temperature $T = 1/k_B\beta$ via

$$F^{\text{phys}}[U^{\text{phys}}] = -k_B T \ln \left\{ \frac{1}{(1 \text{ \AA})^{3N-3}} \int d\mathbf{x} J(\mathbf{x}) e^{-\beta U^{\text{phys}}(\mathbf{x})} \right\}, \quad (2)$$

where the Jacobian is given by

$$J(\mathbf{x}) = \prod_{i=1}^{3N-6} J_i(x_i), \quad (3)$$

with $J_i = x_i^2$ for bond lengths, $J_i = \sin(x_i)$ for bond angles, and $J_i = 1$ for dihedrals. Kinetic energy terms have already been integrated out. Both the dimensionless character of the partition function in Eq. (2) and the angstrom-based normalization result from a particular choice for the standard concentration C° (defined in references^{3,10}) — or equivalently, for the volume containing our implicitly solvated molecule. In particular, we have chosen $C^\circ \equiv$

$8\pi^2(1 \text{ \AA})^{3N-3} Q_p / \sigma$, where $Q_p = \prod_{i=1}^N (2\pi m_i k_B T / h^2)^{3/2}$ results from the momentum integrals, h is Planck's constant, and σ is the molecule's symmetry number. We note that our chosen standard concentration varies based on the molecule (i.e., based on the number and masses of its atoms), and also eliminates the temperature dependence of Q_p . However, in almost every application of interest (see Discussion, below), the absolute free energy calculated here ultimately will be used to estimate a free energy difference and eliminate any artifacts due to C° .

Our single-molecule formulation, as noted, allows for “implicit solvation” using an effective solvent term in U^{phys} that is solely a function of the internal molecular coordinates \mathbf{x} . The present calculations employ a simple uniform dielectric constant ($\epsilon = 60$). In our Discussion, we address the minor technical issues involved with using a more realistic implicit solvent model.

One issue of dimensionality is worth emphasizing. Although there are $3N - 6$ internal coordinates for a molecule consisting of N atoms, the integral of Eq. (2) has dimensionality of length to the power $3N - 3$. This is because $N - 1$ bond lengths remain in the full set of internal coordinates \mathbf{x} , each of which contributes three powers of length. Put another way, of the six excluded rigid-body/center-of-mass coordinates, the three orientation angles are dimensionless; more specifically, the angles integrate to the factor of $8\pi^2$ included in the definition of C° .

Staging the free energy calculation—As illustrated in the example of Sec. IIB, we will calculate the free energy in a series of stages. These can be understood most easily by adding and subtracting the free energies corresponding to k intermediate models,

$$F^{\text{phys}} = (F^{\text{phys}} - F_k) + (F_k - F_{k-1}) + \dots + (F_1 - F^{\text{ref}}) + F^{\text{ref}} \quad (4)$$

$$= \Delta F_{k \rightarrow \text{phys}} + \Delta F_k + \dots + \Delta F_1 + F^{\text{ref}} \quad (5)$$

where $\Delta F_j = F_j - F_{j-1}$ and $F_j[U_j]$ is defined in analogy to Eq. (2) for the intermediate models defined by U_j . The U_j potentials will be specified below.

All free energy difference calculations will be performed here using the “perturbation” formulation.³⁸ Explicitly, for two arbitrary potential energy functions U_a and U_b , one has

$$\Delta F_{a \rightarrow b} = F_b[U_b] - F_a[U_a] = -k_B T \ln \langle \exp[-\beta(U_b - U_a)] \rangle_a \quad (6)$$

$$\simeq -k_B T \ln \left\{ \frac{1}{N_a} \sum_{i=1}^{N_a} \exp[-\beta(U_b(\mathbf{x}_i) - U_a(\mathbf{x}_i))] \right\} \quad (7)$$

where the subscript a denotes an average performed over configurations distributed according to the U_a ensemble and N_a is the number of configurations in that ensemble. Eq. (7) is used to estimate the free energy differences required in Eq. (5), and it is exact in the limit $N_a \rightarrow \infty$.

We emphasize that succeeding intermediates are constructed to have progressively narrower distributions as more interactions are added, as in the alanine dipeptide example. In other words, we ensure good overlap and reliable ΔF estimates by proceeding in the generalized “insertion” direction.^{39,40}

Resampling to obtain staged equilibrium ensembles—As our calculation proceeds through the various stages, we will require the corresponding equilibrium ensembles for each stage, primarily for use in Eq. (7). These are obtained by “resampling,” the process of converting an ensemble for one distribution into another by eliminating, duplicating, and/or adjusting the weights of configurations in the original distribution.³⁷ In our case, we primarily use elimination of configurations from a larger ensemble (e.g., all combinations of fragments A and B) to create a smaller one (e.g., the interacting A-B ensemble); we do not adjust weights. More specifically, to resample an ensemble of configurations \mathbf{x}_a generated according to U_a into a U_b ensemble, the original configurations are randomly selected with probability proportional to the ratio of Boltzmann factors,

$$e^{-\beta[U_b(\mathbf{x}_a) - U_a(\mathbf{x}_a)]}. \quad (8)$$

Operationally, we select configurations by forming a cumulative distribution function (cdf) based on the normalized set of ratios (8), and then choosing from this cdf as many times as desired.

D. Choice of intermediate models

As already noted, the set of intermediate models $\{U_j\}$ can be chosen in a variety of ways. In the present study, we employ k intermediates for a molecule divided into k fragments. This was exemplified for alanine dipeptide, which is divided into three fragments.

The present study uses a uniform staging strategy for all molecules examined, as exemplified in Fig. 1 and Fig. 2. The reference system consists of *all* coordinates fully independent — both within and between fragments — as in our previous work.³ The first intermediate stage adds interactions *within fragments*, so that one has true molecular fragments but no interactions between fragments. We then add interactions between neighboring, covalently bound fragments — i.e., among all the atoms in the neighboring fragment pair — one fragment pair at a time. The final stage of this simple scheme involves the addition of all remaining interactions, which occur solely between non-adjacent

fragments. The result is a molecule with atoms interacting fully according to a standard forcefield and possibly continuum solvent model.

Because interactions among previously non-interacting components are added at every stage, it is expected that the configuration space will become increasingly narrow. Such a progressive narrowing justifies the use of the perturbation expression (7).

To explicitly illustrate the staging scheme employed here, consider the case of a molecule divided into the three fragments A, B, and C, as in Fig. 1. We denote by u_i^{ref} the reference potential for internal coordinate x_i , where the full set is $\mathbf{x} = (x_1, x_2, \dots)$. For the fragments, we let U_y be the potential energy for all interactions internal to fragment y , and U_{yz} is the potential energy for all interactions between the y and z fragments. A three-fragment molecule would be staged as follows:

$$\begin{aligned} U^{\text{ref}} &= \sum_{i=1}^{3N-6} u_i^{\text{ref}}(x_i) \\ U_1 &= U_A + U_B + U_C \\ U_2 &= U_1 + U_{AB} \\ U_3 &= U_2 + U_{BC} \\ U^{\text{phys}} &= U_3 + U_{AC} \end{aligned} \quad (9)$$

The choice of the reference potentials $\{u_i^{\text{ref}}\}$ is guided by the forcefield, as detailed below in Sec. IIG. Building on our earlier work,³ we will use independent internal coordinates with simulation-guided distributions.

A four-fragment molecule, such as di-alanine (Ace-(Ala)₂-Nme) schematized in Fig. 2, would be staged according to:

$$\begin{aligned} U^{\text{ref}} &= \sum_{i=1}^{3N-6} u_i^{\text{ref}}(x_i) \\ U_1 &= U_A + U_B + U_C + U_D \\ U_2 &= U_1 + U_{AB} \\ U_3 &= U_2 + U_{BC} \\ U_4 &= U_3 + U_{CD} \\ U^{\text{phys}} &= U_4 + U_{AC} + U_{BD} + U_{AD} \end{aligned} \quad (10)$$

As described in the Discussion, it is also possible to stage the final (“non-bonded”) pairwise interactions separately.

We anticipate that significant optimization can be obtained by adjusting fragmentation and staging schemes. While our Discussion, below, describes more gradual staging strategies, the present initial report is limited to the single staging strategy given above.

E. The non-interacting reference system

The computation of the (absolute) reference free energy F^{ref} is perhaps the most technically involved step of the calculation. The remaining free energy differences in the decomposition of F^{phys} in Eq. (5) are estimated using a simple, standard method. For the reference free energy, however, great care must be taken with the normalization and Jacobian factors of the

chosen probability distributions. Indeed, our previous report³ includes an error in this regard, as explained at the end of this subsection.

As described in our discussion of staging (Sec. IID), the reference system for all molecules studied here consists of the set of non-interacting internal coordinates. The reference potential energy function will be constructed, following our previous work³, so that the reference partition function is normalized to one. That is, we will *construct* our reference model U^{ref} so that

$$Z^{\text{ref}}[U^{\text{ref}}] = e^{-\beta F^{\text{ref}}} = \frac{1}{(1 \text{ \AA})^{3N-3}} \int d\mathbf{x} J(\mathbf{x}) e^{-\beta U^{\text{ref}}(\mathbf{x})} \equiv 1, \quad (11)$$

where the same standard concentration as in Eq. (2) has been used implicitly. (From this point forward, we will omit writing the length units, but they should be implicitly associated with every bond-length integration.) The motivation for the unit normalization of Z^{ref} is that application of a logarithm leads to the simplifying value,

$$F^{\text{ref}} \equiv 0, \quad (12)$$

for every system.

While there are many ways to construct U^{ref} to satisfy the required normalization of Eq. (11), we primarily use the strategy of employing independent internal coordinates as in our earlier work.³ As usual, the full set of $3N - 6$ internal coordinates, $\mathbf{x} = (x_1, x_2, \dots, x_{3N-6})$ consists of $N - 1$ bond lengths, $N - 2$ bond angles, and $N - 3$ dihedrals. So long as the distribution of each individual coordinate is normalized when integrated with the appropriate Jacobian factor J , the full distribution will be normalized.

Because total reference energy is given by a simple sum of independent terms,

$$U^{\text{ref}}(\mathbf{x}) = \sum_{i=1}^{3N-6} u_i^{\text{ref}}(x_i) \quad (13)$$

the desired normalization (11) is ensured by enforcing

$$\int dx_i J_i(x_i) e^{-\beta u_i^{\text{ref}}(x_i)} = 1, \quad (14)$$

where the inclusion of inverse length units is understood for bond-angle integrals. In words, then, each individual potential u_i^{ref} must include suitable normalization as described in Sec. IIG.

(As detailed in Sec. IIG, peptide ϕ and ψ angles were, in fact, sampled together from a single distribution based on a pairwise energy function $u_{\phi\psi}^{\text{ref}}$. These angles are independent from all other coordinates, however. We emphasize that this exception does not alter the basic formalism, which has been simplified very slightly for clarity.)

It is very useful to observe that normalization of the coordinate distributions via Eq. (14) can be achieved either using standard analytic forms — e.g., Gaussians — or via numerical

histogramming procedures.³ Thus, there is great flexibility in the choice reference distributions embodied in the reference potentials. In addition to forcefield terms, prior knowledge, such as from a simulation, can be used in constructing the set $\{u_i^{\text{ref}}\}$. The reference potentials chosen for the present study are described below in Sec. IIG on library construction.

One word of warning is appropriate here. Although it is possible to describe the internal configuration of a molecule using additional bond angles to substitute for dihedrals in some cases, the Jacobian for such a description appears *not* to be well-defined. Therefore, it is necessary to use the standard description with $N - 2$ bond angles and $N - 3$ dihedrals. Unfortunately, we were unaware of this point during our original study,³ and therefore an erratum will be prepared correcting the resulting numerical errors.

F. First intermediate: non-interacting fragments

The first intermediate stage adds only localized interactions to the non-interacting reference model, as illustrated in Fig. 1 and Fig. 2. Specifically, once a molecule is divided into fragments (A, B, C, ...), the first intermediate includes only interactions occurring *within fragments*. The fragments exactly divide all coordinates so that we can write $\mathbf{x} = (\mathbf{x}_A, \mathbf{x}_B, \dots)$ and the potential energy function for this stage is given by

$$U_1(\mathbf{x}) = U_A(\mathbf{x}_A) + U_B(\mathbf{x}_B) + \dots, \quad (15)$$

where U_y includes all interactions from the *full forcefield* (OPLS-AA, in our case) among the fragment coordinates \mathbf{x}_y for $y = A, B, \dots$. Importantly, the fragment potential U_y includes all non-bonded interactions — electrostatic, van der Waals — among the atoms of the fragment. (Sec. IIG on our libraries describes the treatment of connecting “dummy” atoms.)

The free energy for this stage — i.e., $F_1[U_1]$ for use in the key equation (5), recalling $F^{\text{ref}} = 0$ — can be calculated by using the standard perturbation relation (7). For such a computation, one would use $U_a = U^{\text{ref}}$ and $U_b = U_1$ along with an ensemble distributed according to the Boltzmann factor of U^{ref} .

In practice, *once the libraries are generated, no calculation of energies needs to be done*. As detailed below the libraries are generated (just once, for repeated use in many systems) based on the U^{ref} distribution. Thus, during the library generation process, it is a trivial matter to calculate the absolute free energy for each fragment using Eq. (7). Thus, individual fragment free energies F_y are calculated in advance that exactly sum to the first-stage free energy:

$$F_1 = F_A + F_B + \dots. \quad (16)$$

Further, the independent-coordinate distributions are subsequently resampled based on Eq. (8) to generate the library distributions — i.e., ensembles for the U_y Boltzmann factors — for use in subsequent stages.

G. Construction of fragment libraries

As just described, fragment libraries are critical to the calculation of the free energy of the first intermediate stage, F_1 . The libraries also greatly facilitate computations for the rest of the intermediates.

In general terms, fragment configurations are generated by sampling internal coordinates according to the independent probability distributions which constitute the reference system. The generated configurations are then used to calculate fragment free energies, F_y for $y = A, B, \dots$. The configurations are also reweighted into an ensemble distributed according to the full forcefield for \mathbf{x}_y , the degrees of freedom internal to fragment y . Typically, such a procedure can be applied only to systems with a sufficiently small number of degrees of freedom. For large systems with enough correlated degrees of freedom, there tends to be insufficient overlap with the reference system of independent coordinates. That is, only a tiny fraction of the reference-distributed configurations will be important in the interacting ensemble. Therefore, the choice of the generating probability is essential for the efficient generation of libraries.

We found that for fragments the size of alanine residues, rather simple probability functions were sufficient for generating tens of thousands of (statistically independent) configurations in weeks of single-CPU time. This is a negligible cost because once a library is generated, it can be used in multiple simulations.

Different coordinate types are best sampled with different distributions, as is suggested by the forcefield terms. In our initial tests we employed the same functions and parameters as in the OPLS-AA forcefield — i.e. bonds and angles were described by harmonic functions and dihedrals by a linear combination of several cosine functions. However, this choice of energy functions and parameters did not work well because of strong correlations between coordinates. Better results were obtained when we extracted parameters from a short Langevin simulation of alanine dipeptide, using harmonic functions for bonds, angles and “stiff” dihedrals such as in the relatively planar peptide groups. So for bonds, angles and “stiff” dihedrals the reference potential used in Eq. (14) is

$$u_i^{\text{ref}}(x_i) = \widehat{u}_i(x_i) + k_B T \ln \widehat{z}_i, \quad (17)$$

where \widehat{u}_i is simply a harmonic function and \widehat{z}_i is a constant that enforces the normalization of u_i^{ref} . The harmonic function is described by

$$\widehat{u}_i(x_i) = \frac{(x_i - \bar{x}_i)^2}{2\sigma_i^2} k_B T, \quad (18)$$

where \bar{x}_i is the equilibrium position and σ_i^2 is the variance of a corresponding coordinate. The normalization constant is a single-coordinate partition function

$$\widehat{Z}_i = \int dx_i J_i(x_i) e^{-\beta \widehat{u}_i(x_i)} \quad (19)$$

where inverse length units are implied for all bond-angle integrals.

For “soft”, rotatable dihedrals — such as ϕ , ψ and χ angles in amino acids — we extracted histograms from a Langevin simulation of alanine dipeptide, as described in Ref.3 Thus, a two-dimensional (correlated) probability function was used for the (ϕ, ψ) dihedral pair. In particular, the joint histogram for the (ϕ, ψ) dihedrals consisted of $360 \times 360 = 129,600$ bins with a width of one degree. Histograms for “soft” dihedrals employed 6,000 bins with a width of 0.06 degree.

Coordinates were sampled according to their probability distributions by constructing a numerical cumulative distribution function (cdf) for each:

$$\text{cdf}(x_i) = \int_{-\infty}^{x_i} dx_i J_i(x_i) e^{-\beta u_i^{\text{ef}}(x_i)}. \quad (20)$$

For the histogrammed variables (ϕ , ψ and χ), the cdf will exactly reproduce the distribution in the histogram. For variables with harmonic energy functions, we used a finely discretized numerical cdf to generate a close approximation to the desired distribution. Specifically, bond-length and bond-angle cdf's employed 3,000 bins with a bin-width of 0.001 Å and 0.06 degree respectively. Histograms for dihedrals used 6,000 bins with a width of 0.06 degree. To sample from a given cdf, a uniform random number (0,1] is generated and equated to the cdf to find the corresponding x_i value. Within each bin, the x_i value is chosen uniformly. Parameters for all fragment distributions are available in the Supplementary Material.

Based on the distributions just described, internal coordinates were sampled independently (except for pairwise sampling of ϕ and ψ dihedrals) using an in-house program written in C. Generated configurations were converted to Cartesian coordinates and the potential energy was evaluated using the “analyze” module of the Tinker software package.⁴¹ Based on these values and the known reference energies, the individual fragment free energies were calculated using Eq. (7). A simple resampling procedure³⁷ was used to generate a fragment ensemble distributed according to the forcefield; see Eq. (8). Only a small fraction ($\approx 10^{-4}$) of reference-ensemble configurations remain after resampling, requiring extensive sampling of the reference ensemble and weeks of CPU cost, as mentioned earlier.

For this study, we generated libraries consisting of 10,000 configurations. All fragment libraries were sampled according to OPLS-AA forcefield at $T = 298\text{K}$, with a simple dielectric constant ($\epsilon = 60$) modeling solvent. The choice of dielectric constant was motivated by the reasonable behavior observed in separate Langevin simulations of polyaniline systems (data not shown).

As noted earlier, for all possible 10^8 *covalent* (neighboring) pairings of fragments, we also tabulated the interaction energies from the forcefield, accounting for all atoms in the fragment pair. Suitable corrections for dummy atoms (see below) were made. In other words, for a simple two-fragment system, all interactions are stored.

Use of dummy atoms—Because fragments are sampled independently from each other, the six degrees of freedom that specify the relative orientation of neighboring fragments are included with the fragments. For this purpose “dummy” atoms are used to provide the extra coordinates. When fragments are joined, the dummy atoms are simply removed because they are replaced with the “real” fragment atoms. Although fragment combination is simplest when dummy atoms do not interact with real atoms, we chose to have dummy atoms interact with the fragment atoms for better overlap with subsequent ensembles. Interacting dummy atoms can mimic the neighboring fragment atoms and improve the efficiency of the method. When fragments are joined, however, the interaction energy of the dummy atoms must be subtracted from the full fragment energy because they are replaced with the neighboring fragment atoms. When interacting dummy atoms are employed care must be taken to avoid introducing extra degrees of freedom (e.g., certain bond lengths and angles) because they can affect the ensemble distributions.

The dummy atoms used at the N-terminus of a fragment are carbonyl C, carbonyl O and terminal alpha-C with valence set to one. The dummy atoms used at the C-terminus are

amide N, amide H, and terminal alpha-C with valence also set to one. The dummy atoms were assigned the same forcefield parameters as used in the corresponding fragment atoms.

H. The second and subsequent intermediates: adding neighboring fragment interactions

Returning again to the scheme embodied in Eq. (5), as well as in Fig. 1 and Fig. 2, the next intermediates add interactions between neighboring fragments. These can be considered the “bonded” interactions in the space of fragments, but non-bonded interactions among all atoms in the neighboring pair are included. Explicitly, the models for the remaining intermediates are described by

$$U_2(\mathbf{x})=U_A(\mathbf{x}_A)+U_B(\mathbf{x}_B)+\cdots+U_{AB}(\mathbf{x}_A,\mathbf{x}_B) \quad (21)$$

$$U_3(\mathbf{x})=U_2(\mathbf{x})+U_{BC}(\mathbf{x}_B,\mathbf{x}_C)+\cdots, \quad (22)$$

where U_{yz} is the full interaction energy — based on the forcefield and solvent model — between fragments y and z .

Formally, it is clear what needs to be done. The ensemble of the previous stage $j - 1$ should be used to calculate ΔF_j using the perturbation relation (7) with U_{j-1} and U_j .

Again, however, possession of the libraries and interaction tables leads to dramatic practical implications. For instance, by construction, the energy $U_2 - U_1$ is simply the pre-stored energy U_{AB} ; similarly $U_3 - U_2 = U_{BC}$. These tabulated energies are used directly in Eq. (7) without the need for additional energy calls. The required ensembles for each stage are generated by the rapid resampling procedure of Eq. (8). In this way, one readily generates the free energy differences $\Delta F_2, \Delta F_3, \dots$ required for the evaluation of F^{phys} via Eq. (5).

Caution is required when the molecule of interest contains repeated fragment pairs. While the same libraries can be used for the repeats, say at intermediate stages j and m , the corresponding values of ΔF_j and ΔF_m will be different in general. To see the reason, consider the case of the tetra-alanine peptide studied below. The term ΔF_2 corresponds to including the interaction of the already combined fragments Ace-Ala with the next Ala. Note that the free energy difference ΔF_2 is calculated via Eq. (7) using the Ace-Ala ensemble as the “a” system. By contrast, consider the calculation of ΔF_3 for the addition of the next Ala — now to the Ace-(Ala)₂ ensemble. Although the free energy change will be based upon the identical (tabulated) interactions, the associated Boltzmann factors in Eq. (7) will be weighted differently — i.e., occur with different frequencies — due to the differing initial “a” ensembles. In turn, this will lead to different free energy changes, so that $\Delta F_3 \neq \Delta F_2$.

I. The final free energy difference: non-neighboring interactions

As described in the master scheme of Eq. (5) and illustrated in Fig. 1 and Fig. 2, the final calculation needed to obtain F^{phys} entails the inclusion of all remaining interactions in the forcefield and solvent model. These interactions, excluded until now, occur between atoms in non-neighboring pairs. As described in Sec. II D, for a molecule of k fragments, the full physical potential energy function (i.e., the forcefield) can be written as the difference from the final (k th) intermediate:

$$U^{\text{phys}}(\mathbf{x}) = U_k(\mathbf{x}) + \sum_{y..z} U_{yz}(\mathbf{x}_y, \mathbf{x}_z), \quad (23)$$

where the sum is over non-neighboring pairs of fragments — i.e., AC, AD, BD,

In this case, the necessary energy terms for use in the calculation of $\Delta F_{k \rightarrow \text{phys}}$ via Eq. (7) must be calculated. They cannot readily be stored in advance, due to the combinatorial explosion of possible configurations. For instance, with libraries of 10^4 configurations, there are 10^{12} possible configurations for three fragments, which is beyond the range of current commercial machines.

J. Generating an equilibrium ensemble without additional energy calls

The physical ensemble, distributed according to the Boltzmann factor of the forcefield, can be generated by resampling the U_k ensemble — the last intermediate — using Eq. (8). In this case, the “a” ensemble corresponds to U_k and the “b” ensemble to the full forcefield and (implicit) solvent model. Because all energy terms have already been calculated, no additional energy calls need to be made. The necessary resampling computation is extremely fast compared with preceding stages of the protocol.

K. Checking the code and estimating uncertainty

Although the formalism governing the present study is mostly straightforward, our in-house computer program not only needs to reproduce standard forcefield results, but also requires complicated “dissections” of various subsets of forcefield terms. We therefore performed three types of checks on our code. (i) We checked that the forcefield energy for full molecular configurations exactly reproduces the results reported in Tinker (data not shown). This verifies that we have correctly accounted for our dummy-atom energy terms, (ii) Using our previously developed “structural histograms” for analyzing configuration-space distributions,^{42,43} we checked that the equilibrium ensembles produced during our free energy calculations agree with independent Langevin simulations. This data is shown in the Results section, and generated as explained below, (iii) Finally, we performed a check to ensure that our final free energy values are independent of the choice of fragments. This data, for two- and three-fragment decompositions of alanine dipeptides is also shown in the Results section.

Statistical error—Statistical uncertainties were calculated by running 20 independent computations for every free energy value reported. Twice the standard deviation among these 20 values is reported, which quantifies the scale of expected statistical error for a *single* simulation. The repeated simulations were run using 20 independent sets of libraries for the various fragments — i.e., the calculation was started all the way at the beginning in each repeat. However, because the overlap between various stages is the limiting factor in the quality of the free energy results, rather than our fairly large libraries, we anticipate similar error estimates would be obtained for one set of libraries.

Analyzing equilibrium ensembles/distributions—In two previous studies,^{42,43} we have developed methods for comparing equilibrium distributions for molecular systems of arbitrary complexity. The central idea is to employ a “structural histogram” which simply classifies (divides) configuration space into a number of “bins” (regions). Two correct simulations should yield the same results for the fractional populations of the bins, within statistical error, regardless of whether the bins correspond to physical states/free energy basins. (Furthermore, the statistical uncertainty in the population estimates can be used to

quantify the “effective sample size”).⁴³ In the present work, we compare equilibrium distributions from fragment combination and from standard Langevin simulations based on structural histograms. The particular histograms employed in the present study have five bins derived from a Voronoi construction;⁴⁴ the reference structures for the Voronoi procedure are derived from the equi-probability scheme described in reference.⁴³ Although the resulting bins are not exactly equally probable, each is guaranteed to represent a contiguous region in configuration space due to the Voronoi construction.

III. RESULTS

The absolute configurational free energy F^{phys} was calculated for the monomer, dimer, and tetramer alanine peptides: alanine dipeptide (Ace-Ala-Nme), di-alanine (Ace-(Ala)₂-Nme), and tetra-alanine (Ace-(Ala)₄-Nme). For alanine dipeptide, the free energy was estimated based on two different fragment sets as a check on our code. Twenty independent calculations for every F^{phys} estimate were performed to quantify uncertainty, as described above. Additionally, every free energy calculation also yields an equilibrium ensemble, which is compared to independent Langevin simulations.

The results are very encouraging in every regard, and rather rapid as reported at the end of the Results section. The amount of memory used, which is a key to the present calculations, is also reported.

The results reflect the uniform protocol adopted here. First, absolute free energies for non-interacting fragments are calculated. Then free energy changes resulting from interactions among covalently bound fragments are added (“bonded” terms, in the space of fragments), one at a time in sequence. Finally, all remaining interactions are added, which account to (“non-bonded”) interactions among all non-sequential fragment atoms. The final free energy values reflect the full OPLS-AA forcefield³⁶ as implemented in Tinker.⁴¹

A. Alanine dipeptide using two different fragmentations

Because of the complexity of the fragmentation procedure and the lack of reference standards for absolute free energy values, we wanted to ensure our code was introducing no artifacts. We were particularly concerned about the interacting dummy atoms which introduce “temporary” energy terms, that must be corrected for properly at every combination stage. We find good agreement between free energy estimates based on two- and three-fragment decompositions.

“Standard” three-fragment decomposition—In our standard decomposition for the present study, we separate peptide and amino acid groups. For alanine dipeptide (AD), then, the three standard fragments are Ace, Ala, and Nme, and the corresponding stages for the free energy calculation are given in Eq. (9). Recalling our convention that $F^{\text{ref}} \equiv 0$, the free energy terms from Eq. (5) can be written as

$$\begin{aligned}
 F^{\text{ref}} &\equiv 0 \\
 \Delta F_1 &= F_{\text{Ace}} + F_{\text{Ala}} + F_{\text{Nme}} \\
 \Delta F_2 &= \Delta F_{\text{Ace} \rightarrow \text{Ala}} \\
 \Delta F_3 &= \Delta F_{\text{Ala} \rightarrow \text{Nme}} \\
 \Delta F_{3 \rightarrow \text{phys}} &= \Delta F_{\text{nonbonded}}
 \end{aligned}
 \tag{24}$$

where F_y is the absolute free energy (including dummy atoms) for fragment y and $\Delta F_{x \rightarrow y}$ indicates the free energy change of combining fragments x and y (which includes all bonded and non-bonded terms, as well as the correction of dummy terms). Finally, $\Delta F_{\text{nonbonded}}$

denotes the free energy change in going from an ensemble where sequentially separated fragments do not interact to a fully interacting ensemble (in this case, the Ace-Nme interactions are added).

Two-fragment decomposition—As an alternative decomposition, we used Ace-Ala as one fragment and Nme as the other. Importantly, the Ace-Ala library and absolute free energy were *not* generated from a combination of the two smaller libraries, but instead from a ground-up calculation based on independent coordinates as described in the Methods section.

In this case, then, the free energy terms from Eq. (5) become

$$\begin{aligned} F^{\text{ref}} &\equiv 0 \\ \Delta F_1 &= F_{\text{Ace-Ala}} + F_{\text{Nme}} \\ \Delta F_{1 \rightarrow \text{phys}} &= \Delta F_{\text{Ace-Ala} \rightarrow \text{Nme}}, \end{aligned} \quad (25)$$

where it is notable that in the two-fragment case, *all* interactions are included in the libraries and interaction tables. In other words, no energy calls at all are needed.

Comparison of free energies—There is essentially perfect agreement between free energies estimate via the two independent decompositions, which provides a reassuring check on our computer program. The full results are given in Table I. Notably, the two-fragment decomposition has a higher variance, which probably results from a decreased “precision” in the pre-generated Ace-Ala ensemble. In the composite pre-generated Ace-Ala ensemble, the whole configuration space is represented by 10^4 configurations, whereas when Ace and Ala from separate 10^4 -member libraries are combined, there is a much denser coverage of configuration space.

Equilibrium ensemble compared to standard simulation—Our free energy computation produces an equilibrium ensemble through repeated resampling procedures at each stage, as explained in the Methods section (Sec. IIJ). As a further check on our data, we compare the equilibrium ensembles generated from our fragment combination procedure to those produced by long Langevin simulations performed in Tinker.⁴¹ The results, shown in Fig. 3(a), indicate that our computation is indeed producing correct equilibrium ensembles. The graph shows the populations of different regions of configuration space, which was divided up using a Voronoi procedure explained above (Sec. IIK). The alanine dipeptide equilibrium distribution was generated from the three-fragment protocol, and the 1 μ sec. Langevin simulation (20 \times 50 nsec) was performed in Tinker using a friction constant of 10.0 ps⁻¹ at $T = 298K$.

B. Di-alanine

Using the same libraries as for the alanine monomer above, we now calculate the absolute configurational free energy for the di-alanine peptide (Ace-(Ala)₂-Nme). The staging used is described in Eq. (10), which corresponds to the following free energy terms for use in Eq. (5):

$$\begin{aligned}
 F^{\text{ref}} &\equiv 0 \\
 \Delta F_1 &= F_{\text{Ace}} + 2 \cdot F_{\text{Ala}} + F_{\text{Nme}} \\
 \Delta F_2 &= \Delta F_{\text{Ace} \rightarrow \text{Ala}} \\
 \Delta F_3 &= \Delta F_{\text{Ala} \rightarrow \text{Ala}} \\
 \Delta F_4 &= \Delta F_{\text{Ala} \rightarrow \text{Nme}} \\
 \Delta F_{4 \rightarrow \text{phys}} &= \Delta F_{\text{nonbonded}}.
 \end{aligned}
 \tag{26}$$

It is important to note that $\Delta F_{\text{Ala} \rightarrow \text{Ala}}$ for di-alanine differs in principle from the term with the same name in Eq. (24) for alanine-dipeptide because the prior ensemble is different. This point was explained in Sec. III. However, in practice, the difference among the three systems is not statistically significant.

The free energy values are once again calculated with high precision: fluctuations are a fraction of one kcal/mole. The data for all free energy terms is given Table II, where we see a significant change in the $\Delta F_{\text{nonbonded}}$ term, reflecting the increased number of attractive interactions in this larger molecule (compared to alanine dipeptide).

Similarly, the agreement among bin populations for di-alanine in Fig. 3(b) is good, which provides an independent reason for having confidence in the free energy results. The Langevin simulation for di-alanine was performed with exactly the same parameters as for alanine-dipeptide.

C. Tetra-alanine

Our results are of high precision (~ 0.1 kcal/mole standard deviation) even for tetra-alanine. The staging follows our standard procedure, with the only subtlety in the present case is that the addition of every Ala residue is different, because the “growing” ensemble is different in every case. Thus we consider the first alanine (Ala1) separate from the second (Ala2), and so on.

$$\begin{aligned}
 F^{\text{ref}} &\equiv 0 \\
 \Delta F_1 &= F_{\text{Ace}} + 4 \cdot F_{\text{Ala}} + F_{\text{Nme}} \\
 \Delta F_2 &= \Delta F_{\text{Ace} \rightarrow \text{Ala1}} \\
 \Delta F_3 &= \Delta F_{\text{Ala1} \rightarrow \text{Ala2}} \\
 \Delta F_4 &= \Delta F_{\text{Ala2} \rightarrow \text{Ala3}} \\
 \Delta F_5 &= \Delta F_{\text{Ala3} \rightarrow \text{Ala4}} \\
 \Delta F_6 &= \Delta F_{\text{Ala4} \rightarrow \text{Nme}} \\
 \Delta F_{6 \rightarrow \text{phys}} &= \Delta F_{\text{nonbonded}}.
 \end{aligned}
 \tag{27}$$

Our data for each of these terms is given in Table II. Although the different alanine additions are based on different ensembles, the results show they are statistically indistinguishable in this case. However, the “non-bonded” term $\Delta F_{\text{nonbonded}}$ again is significantly different from the previous molecules, as expected.

In comparing the equilibrium distributions from fragment combination and Langevin simulation, once again there is good statistical agreement. For Langevin simulation of tetra-alanine, all parameters were set as before, except for a friction constant of 5.0 ps^{-1} , which does not alter the equilibrium distribution. The contrast between the large fluctuations in the bin populations p_i and the high precision of F^{phys} in Table II reveals an important lesson: sampling is harder than free energy calculation. Future work will describe efforts to improve efficiency for fragment-based equilibrium sampling.

D. Timing and memory usage

The calculations were reasonably inexpensive, taking 20 minutes for alanine dipeptide, 30 minutes for di-alanine, and 50 minutes for tetra-alanine using one processor of an Intel Xeon 3.20 GHz machine. Concerning memory, a single library containing 10,000 configurations requires 11 MB for Ala, 12 MB for Ace-Ala complex and 5.7 MB for Ace and Nme. An interaction table containing 10^8 pair-wise interactions uses 1.3 GB.

IV. DISCUSSION

A. The overall strategy and results

Overall, the precision of our free energy estimates was very high, which we attribute to two related factors. First, our ensembles in the reference and intermediate stages were of good statistical quality — i.e., characterized by a large effective sample size (data not shown).⁴³ Second, there was good overlap between the stages, which indeed contributed to maintaining the sample size throughout the stages. The overlap is present by design, as interactions were always *added* between stages. The addition of interactions or, equivalently, correlations among degrees of freedom is guaranteed to reduce the entropy.⁴⁵ This progressive narrowing of configuration space is consistent with Kofke's proposal to calculate free energy differences in the "insertion" direction.^{39,40} For larger systems, however, one expects limitations to maintaining the sample size using the present protocol, as explained below.

We note that our present approach differs from our previous work³ primarily because the earlier study did not employ staging. Effectively, it was a single-fragment calculation from the perspective of the present study.

B. Application of fragment combination for estimating relative protein-ligand affinities

Because the fragment combination procedure can be applied to fragments of small molecules, and not just to peptides as in the present report, the approach can be applied to calculate approximate relative affinities. That is, one can grow a ligand into the binding pocket of a protein receptor and calculate its free energy. A number of different approximations can be imagined. Most simply, the receptor can be held rigid and the ligand grown in the fields (van der Waals and electrostatic) of the receptor. In a better approximation, the binding-site side-chains can be grown along with the ligand. One can expect affinities based on such free energy calculations to be superior to their docking counterparts because entropy is included. To produce a relative affinity estimate between two ligands, the respective solvation terms would need to be included as usual.¹²

C. Efficiency of fragment combination for equilibrium sampling

As we have already noted, the fragment combination protocol we have described produces equilibrium ensembles simultaneously with free energy estimates. It is natural to wonder whether such ensembles are produced more efficiently than by standard dynamics simulation especially given that small peptides have been shown to have multi-nanosecond relaxation times.⁴³ In fact, as we will document in a separate publication, fragment combination can lead to sampling that is faster by several orders of magnitude. However, somewhat more sophisticated resampling schemes and different fragment sizes are useful in reaching the highest levels of efficiency, as will be reported.

D. Use of implicit solvent models

It is interesting and important to consider the additional costs which would be entailed by using a standard implicit solvent model, such as GBSA.⁴⁶ First, both the libraries and the

interaction tables would need to be regenerated using the implicit solvent model. Although this could take several weeks of single-CPU time, it needs to be done only a single time for a given model. The second cost is for additional solvent calculations not included in the libraries and tables. We hope to report on the staging and computational expense in a forthcoming publication.

E. Relaxation simulations for large systems

The protocol described in this report will encounter challenges with larger systems, as our results hint. In the present study, the equilibrium ensemble generated at one stage, say j , is used to calculate the incremental free energy difference to the next stage, $j + 1$. To continue the process, the ensemble at stage $j + 1$ is produced by resampling ensemble j as described in the Methods section. However, it is possible the resampled ensemble will contain a small number of distinct configurations in an important part of configuration space. Such configurations will have high weight prior to resampling and thus the problem can be diagnosed by noting whether any configurations are resampled multiple times. Clearly, duplicated configurations will not be statistically independent and lead to increased statistical error in free energy estimates.

One solution to this problem would be to relax duplicated configurations — i.e. to perform short equilibrium simulations to create distinct configurations. The statistical justification of such an approach is somewhat technical⁴⁷ and will be described in future work as required.

F. Alternative staging using partial interactions

Additional incremental stages can be added by considering only subsets of interactions. For instance, in the case of di-alanine (Ace-(Ala)₂-Nme) which is composed of four fragments (A, B, C, D), there are three sets of non-bonded interactions: AC, AD, and BD. Our present implementation adds all three in a single stage, but they could be added one at a time. Undoubtedly, in larger systems, such finer staging will be necessary and probably should be required by relaxation of duplicated configurations as just described.

V. SUMMARY AND CONCLUSIONS

We have extended our earlier work on computation of absolute free energies for molecular systems **3** in two ways: (i) by staging the calculation using molecular fragments; and (ii) by pre-calculating “statistical libraries” of fragment configurations, intra-fragment energies, and inter-fragment interactions. For a series of test systems – the alanine monomer, dimer, and tetramer — we were able to compute extremely precise free energies, with fluctuations $\ll 1$ kcal/mole. The calculations were quite fast, furthermore, with the slowest requiring less than an hour of single-processor computer time. The speed results from employing (infinitely re-usable) pre-calculated library configurations and interactions, which pre-empt expensive energy calculations. Our statistical libraries of amino-acid and capping-group fragments are available on our website (www.cccb.pitt.edu/Zuckerman).

A future application of potential importance in computational biochemistry is the estimation of binding affinities of small molecules to proteins. We hope to develop fragment libraries suitable for small molecules — following ideas developed long ago^{29,30} — to be used in computing free energies within a potentially flexible protein binding site. Protein flexibility could be included using the side-chains of the binding site as fragments in the calculations.

Another application closely related to the present report is the use of library-based fragment combination for equilibrium sampling. While the calculations reported here already yield equilibrium ensembles, we are actively studying more efficient schemes based on advanced

resampling approaches^{37,48} and a range of fragment sizes. This work will be reported in a separate publication.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We greatly appreciate the advice of Divesh Bhatt, Ying Ding, Marty Ytreberg, and Bin Zhang, as well as the financial support of the NIH (Grants GM076569 and GM070987) and NSF (Grant MCB-0643456).

References

1. Frenkel, D.; Smit, B. *Understanding Molecular Simulation*. San Diego: Academic Press; 1996.
2. Stoessel JP, Nowak P. *Macromol*. 1990; 23:1961.
3. Ytreberg FM, Zuckerman DM. *J. Chem. Phys.* 2006; 124:104105. [PubMed: 16542066]
4. Huang L, Makarov DE. *J. Chem. Phys.* 2006; 124:64108. [PubMed: 16483197]
5. Meirovitch H. *J. Chem. Phys.* 1992; 97:5803.
6. Meirovitch H. *J. Chem. Phys.* 1999; 111:7215.
7. Chelvaraja S, Meirovitch H. *Proc. Natl. Acad. Sci. U. S. A.* 2004; 101:9241. [PubMed: 15197271]
8. White RP, Meirovitch H. *Proc. Natl. Acad. Sci. U. S. A.* 2004; 101:9235. [PubMed: 15197270]
9. Head MS, Given JA, Gilson MK. *J. Phys. Chem. B.* 1997; 101:1609.
10. Chang C-E, Potter MJ, Gilson MK. *J. Phys. Chem. B.* 2003; 107:1048.
11. Chang C, Gilson M. *J. Am. Chem. Soc.* 2004; 126:13156. [PubMed: 15469315]
12. Chen CC, W, Gilson M. *Biophys. J.* 2004; 87:3035. [PubMed: 15339804]
13. Go N, Scheraga H. *J. Chem. Phys.* 1969; 51:4751.
14. Rosenbluth MN, Rosenbluth AW. *J. Chem. Phys.* 1955; 23:356.
15. Wall FT, Erpenbeck JJ. *J. Chem. Phys.* 1959; 30:634.
16. Meirovich H. *J. Phys. A: Math. Gen.* 1982; 15:L735.
17. Meirovich H. *Phys. Rev. A.* 1985; 32:3699. [PubMed: 9896539]
18. Garel T, Orland H. *J. Phys. A: Math. Gen.* 1990; 23:L621.
19. Garel T, Niel JC, Orland H, Smith J, Velikson B. *J. Chem. Phys.* 1991; 88:2479.
20. Velikson B, Garel T, Niel JC, Orland H, Smith JC. *J. Comput. Chem.* 1992; 13:1216.
21. Bascle J, Garel T, Orland H, Velikson B. *Biopolymers.* 1993; 33:1843. [PubMed: 8268410]
22. Grassberger P. *J. Phys. A: Math. Gen.* 1993; 26:2769.
23. Grassberger P, Hegger R. *J. Phys.- Condens. Mat.* 1995; 7:3089.
24. Grassberger P. *Phys. Rev. E.* 1997; 56:3682.
25. Bastolla U, Frauenkron H, Gerstner E, Grassberger P, Nadler W. *Proteins.* 1998; 32:52. [PubMed: 9672042]
26. Zhang JL, Liu JS. *J. Chem. Phys.* 2002; 117:3492.
27. Zhang J, Lin M, Chen R, Liang J, Liu JS. *Proteins.* 2007; 66:61. [PubMed: 17039507]
28. Gibson KD, Scheraga HA. *J. Comput. Chem.* 1987; 8:826.
29. Miranker A, Karplus M. *Proteins.* 1991; 11:29. [PubMed: 1961699]
30. Leach, AR. *Molecular Modelling: Principles and Applications*. Pearson Education Limited; 2001.
31. Mamonov AB, Bhatt D, Cashman DJ, Zuckerman DM. submitted.
32. Wall FT, Rubin RJ, Isaacson LM. *J. Chem. Phys.* 1957; 27:186.
33. Alexandrowicz Z. *J. Chem. Phys.* 1969; 51:561.
34. Macedonia MD, Maginn EJ. *Mol. Phys.* 1999; 96:1375.
35. Rohl CA, Strauss CEM, Misura KMS, Baker D. *Method Enzymol.* 2004; 383:66.
36. Jorgensen WL, Maxwell DS, Tirado-Rives J. *J. Am. Chem. Soc.* 1996; 118:11225.

37. Liu, JS. Monte Carlo strategies in scientific computing. Springer; 2004.
38. Zwanzig RW. J. Chem. Phys. 1954; 22:1420.
39. Peter DAK, Cummings T. Mol. Phys. 1997; 92:973.
40. Kofke DA, Cummings PT. Fluid Phase Equilibr. 1998; 150–151:41.
41. Ponder JW, Richard FM. J. Comput. Chem. 1987; 8:1016.
42. Lyman E, Zuckerman DM. Biophys. J. 2006; 91:164. [PubMed: 16617086]
43. Lyman E, Zuckerman DM. J. Phys. Chem. B. 2007; 111:12876. [PubMed: 17935314]
44. Voronoi G. Journal für die Reine und Angewandte Mathematik. 1907; 133:97.
45. Cover, TM.; Thomas, JA. Elements of Information Theory. 2nd Edition. Wiley; 2006.
46. Qiu D, Shenkin PS, Hollinger FP, Still WC. J Phys. Chem. A. 1997; 101:3005.
47. Neal RM. Stat. comput. 2001; 11:125.
48. Fearnhead P, Clifford P. J. R. Stat. Soc. B. 2003; 65:887.




Figure 1.

Stages for calculating the absolute free energy of a molecule by combining three fragments, based on Eq. (9). Connecting lines schematize full interactions between fragments, including both bonded and non-bonded atomistic terms, (a) The first intermediate stage comprises non-interacting fragments, but include all interactions *internal* to each fragment, (b) The second stage adds interactions among the atoms of fragments A and B, while (c) the third stage does the same for fragments B and C. (d) In the final stage, representing the desired free energy F^{phys} , all interactions are added, including among non-sequential fragments and possibly including an implicit solvent model.




Figure 2.

Stages used in the free energy calculation of a four-fragment molecule, corresponding to Eq. (10). The initial stages proceed in analogy to Fig. 1, with pair-wise interactions added one at a time for neighboring (“bonded”) fragments. In the final stage, *all* remaining interactions are added. Other, more incremental staging schemes are possible, but were not necessary in the present study.




Figure 3. Comparison of equilibrium distributions from fragment combination and Langevin simulation. The graphs show the fractional population in different regions of configuration space, as described in Sec. II K. Three peptides are considered: (a) alanine dipeptide, (b) di-alanine, and (c) tetra-alanine. The error bars for both the fragment combination and Langevin results reflect twice the standard deviations among 20 independent simulations, roughly a 95% confidence interval. Each Langevin simulation was 50 nsec long. The statistical agreement is good in every case.

Table I

Comparison between the absolute free energy for alanine dipeptide estimate using two different fragmentation schemes. Our “standard” three-fragment decomposition (Ace, Ala, Nme) is compared to a two-fragment grouping (Ace-Ala, Nme). The table gives free energy values in kcal/mole, as well as two standard deviations (in parentheses) based on 20 independent calculations.

Alanine dipeptide free energy terms from Eq. (24) and Eq. (25)			
Three Fragments		Two Fragments	
Term	Estimate [kcal/mol]	Term	Estimate [kcal/mol]
F_{Ace}	14.783(0.003)	$F_{\text{Ace-Ala}}$	47.311(0.027)
F_{Ala}	33.326(0.015)	F_{Nme}	16.574(0.003)
F_{Nme}	16.574(0.003)	$\Delta F_{\text{Ace-Ala} \rightarrow \text{Nme}}$	-0.792(0.002)
$\Delta F_{\text{Ace} \rightarrow \text{Ala}}$	-0.801(0.002)		
$\Delta F_{\text{Ala} \rightarrow \text{Nme}}$	-0.499(0.007)		
$\Delta F_{\text{nonbonded}}$	-0.285(0.008)		
F^{phys}	63.098(0.015)	F^{phys}	63.093(0.028)

Table II

Free energy terms used in calculating the absolute free energy for di-alanine and tetra-alanine. The table gives free energy values in kcal/mole, as well as two standard deviations (in parentheses) based on 20 independent calculations.

Free energy terms for di-alanine from Eq. (26) and for tetra-alanine from Eq. (27)			
Di-alanine		Tetra-alanine	
Term	Estimate [kcal/mol]	Term	Estimate [kcal/mol]
F_{Ace}	14.783(0.003)	F_{Ace}	14.783(0.003)
F_{Ala}	33.326(0.015)	F_{Ala}	33.326(0.015)
F_{Nme}	16.574(0.003)	F_{Nme}	16.574(0.003)
$\Delta F_{\text{Ace} \rightarrow \text{Ala}}$	-0.801(0.002)	$\Delta F_{\text{Ace} \rightarrow \text{Ala1}}$	-0.801(0.002)
$\Delta F_{\text{Ala} \rightarrow \text{Ala}}$	-0.771(0.014)	$\Delta F_{\text{Ala1} \rightarrow \text{Ala2}}$	-0.774(0.013)
$\Delta F_{\text{Ala} \rightarrow \text{Nme}}$	-0.499(0.013)	$\Delta F_{\text{Ala2} \rightarrow \text{Ala3}}$	-0.774(0.012)
$\Delta F_{\text{nonbonded}}$	-0.809(0.031)	$\Delta F_{\text{Ala3} \rightarrow \text{Ala4}}$	-0.771(0.014)
		$\Delta F_{\text{Ala4} \rightarrow \text{Nme}}$	-0.498(0.009)
		$\Delta F_{\text{nonbonded}}$	-1.986(0.284)
F^{phys}	95.128(0.057)	F^{phys}	159.057(0.293)