

Protein Thermostability Calculations Using Alchemical Free Energy Simulations

Daniel Seeliger and Bert L. de Groot*

Computational Biomolecular Dynamics Group, Max-Planck-Institute for Biophysical Chemistry, Göttingen, Germany

ABSTRACT Thermal stability of proteins is crucial for both biotechnological and therapeutic applications. Rational protein engineering therefore frequently aims at increasing thermal stability by introducing stabilizing mutations. The accurate prediction of the thermodynamic consequences caused by mutations, however, is highly challenging as thermal stability changes are caused by alterations in the free energy of folding. Growing computational power, however, increasingly allows us to use alchemical free energy simulations, such as free energy perturbation or thermodynamic integration, to calculate free energy differences with relatively high accuracy. In this article, we present an automated protocol for setting up alchemical free energy calculations for mutations of naturally occurring amino acids (except for proline) that allows an unprecedented, automated screening of large mutant libraries. To validate the developed protocol, we calculated thermodynamic stability differences for 109 mutations in the microbial Ribonuclease Barnase. The obtained quantitative agreement with experimental data illustrates the potential of the approach in protein engineering and design.

INTRODUCTION

Rational engineering of proteins (1) to optimize a natural protein for a specific task (e.g., to achieve higher thermal stability, altered substrate specificity, or solubility) is one of the most exciting tasks in biotechnology. This is particularly true for enzymes. Two recently published pioneering articles describe the design of novel enzymes which catalyze chemical reactions that are not known to be catalyzed by any naturally occurring enzyme (2,3). Additionally of great industrial importance is the optimization of enzymes toward higher efficiency and thermostability, to enable them to be used as detergents or for the thermostabilization of therapeutic proteins (4–9). Most of the successful applications of rational protein engineering, so far, have been built on knowledge-based scoring functions (10,11), implicit solvent models (12,13), or are SVM-based (14); however, molecular-dynamics-based methods utilizing explicit solvent have come of age and are developing into a high accuracy alternative with great potential. Although the computational demand of calculating the free energy difference of a single point mutation is several orders-of-magnitude larger than with a knowledge-based scoring function, physics-based methods do have an advantage. They can be applied for those cases where a large database is not available for the derivation of statistical potentials.

Enzymes are interesting catalysts for enantioselective synthesis in chemical industry (15–20). However, their natural environment, water, is often a poor solvent for organic molecules. Many enzymes retain both structure and some functionality in organic solvents (21,22). Rational optimization under these conditions, however, is limited,

with scoring functions based on statistical potentials. Here, physics-based methods should work with the same accuracy as in aqueous solution and open the possibility to rational protein engineering in nonnatural environments.

In a recent review, Potapov et al. (23) compared six established protocols and assessed their ability to predict the thermodynamic consequences of point mutations. Their work revealed that, although the assessed methods are, on average, capable of predicting the correct trend, the accuracy of the best protocol did not exceed a correlation coefficient of 0.6. These findings indicate that the development of more accurate methods is highly desirable, particularly when considering that free energy changes resulting from point mutations are rather small in most cases.

Among the simulation protocols that aim at calculating free energy differences, perturbation approaches (24–27) have grown in popularity over the last years. Here, the Hamiltonian H is coupled to a parameter λ which is used to drive a system from a state A ($\lambda = 0$), here corresponding to the Hamiltonian of the wild-type protein, to a state B ($\lambda = 1$), corresponding to the Hamiltonian of the mutant. Free energies can either be computed by using so-called equilibrium methods such as free energy perturbation (28) or thermodynamic integration (TI) (29), or by using non-equilibrium methods such as those based on the work of Jarzynski (30,31) and Crooks (32). Although the simulation protocols and analysis methods differ substantially, their initial setup is identical. While λ is switched from 0 to 1 the system must evolve from state A to state B, corresponding, e.g., to an amino acid mutation, thereby changing intramolecular and intermolecular interactions. Depending on the particular type of mutation, atoms must be annihilated (decoupled from the system), or dummy atoms turned into real atoms. Hence, a topology for such a simulation has to

Submitted November 25, 2009, and accepted for publication January 21, 2010.

*Correspondence: bgroot@gwdg.de

Editor: Bertrand Garcia-Moreno.

© 2010 by the Biophysical Society
0006-3495/10/05/2309/8 \$2.00

doi: 10.1016/j.bpj.2010.01.051

ensure that all bonded and nonbonded interactions are correctly switched between the two states.

In this article, we present a novel library (to our knowledge) that contains hybrid residues for all amino acid mutations except proline mutations, each of them representing different amino acids in state A and state B, respectively. This protocol allows us to set up mutation free energy calculations in an automated way, thereby enabling the use of perturbation-based methods for screening of large mutation databases, and hence, computational protein engineering. To validate the approach, we calculated folding free energy differences for 109 point mutations in the microbial Ribonuclease Barnase, and compare the results with experimentally determined values.

METHODS

Construction of hybrid residues

We have compiled a database of hybrid residues that contain all possible mutations involving naturally occurring amino acids, except for proline. For residues with alternating protonation states, we treated each protonation state as a separate residue, which leads to a total number of 24 residues and, hence, a database of 552 hybrid residues (see Fig. 1). The database consists of coordinates for each hybrid residue and force-field parameters of the AMBER99sb force field (33) for both states. To enable an automated setup of free energy calculations, we developed two Python-scripts based on the PYMACS package (<http://wwwuser.gwdg.de/~dseelig/pymacs.html>) which 1), replace residues in a structure file by appropriate hybrid residues from the database; and 2), modify GROMACS topology files to ensure the correct force field parameters for each state. The scripts have been extensively tested for all possible mutations, and they allow the straightforward setup of free energy simulations.

Free energy calculations

To test our mutation library, we calculated folding free energy differences for 109 point mutations at 64 different positions in the well-studied microbial Ribonuclease Barnase (Protein DataBank No. 1bni, 110 residues). All simulations were carried out with GROMACS-4.0 (34,35), the AMBER99sb force field (33), and the TIP3P water model (36). As simula-




FIGURE 1 Hybrid residues. A database of hybrid residues has been compiled. A script replaces one or more residues in a structure file by hybrid residues that represent one amino acid in state A and another amino acid in state B.

tion protocol, we chose nonequilibrium fast-growth thermodynamic integration (FGTI) runs. Therefore, for each mutation, equilibrium ensembles at $\lambda = 0$ and at $\lambda = 1$ are required. The simulation system consists of the protein solvated in a dodecahedron box with ≈ 7400 water molecules and NaCl was added to achieve a 150 mM solution. Both the A- and B-states were sampled for 10 ns using a stochastic dynamics integrator at 298 K and constant pressure of 1 atm using the Parrinello-Rahman barostat (37). Electrostatic interactions were calculated at every step with the particle-mesh Ewald method (38), and short-range repulsive and attractive dispersion interactions simultaneously described by a Lennard-Jones potential with a cutoff of 1.1 nm and a switching function that was used between 1.0 and 1.1 nm. Dispersion correction for energy and pressure was applied. The SETTLE (39) algorithm was used to constrain bonds and angles of water molecules, and LINCS (40) was used for all other bonds, allowing a time step of 2 fs.

From these ensembles, 100 snapshots from the last 8 ns were taken and short simulations were performed in which λ was changed from zero to one, or from one to zero, respectively. For the fast-growth TI simulations, we used a double-precision version of GROMACS-4.0 with a leap-frog integrator and a velocity-rescaling thermostat (41). Energy calculations, timestep, and pressure coupling was analogous to the equilibration runs. To account for atomic overlaps occurring close to $\lambda = 0$ and $\lambda = 1$, soft-core potentials were used for both electrostatics and Lennard-Jones interactions as implemented in GROMACS-4.0 with $\alpha = 0.3$, $\sigma = 0.25$, and a soft-core power of 1. The complete switching from $\lambda = 0$ to $\lambda = 1$ was done within 50 ps (note that we did not separate the electrostatic and Lennard-Jones part) and derivatives of the Hamiltonian with respect to λ were recorded at every step. Free energies were calculated from the work distributions obtained from integration according to

$$W = \int_{\lambda=0}^{\lambda=1} \frac{\delta H_{\lambda}}{\delta \lambda} d\lambda$$

and calculating the intersection of the forward and backward work distributions according to the Crooks-Gaussian-intersection method as described in Goette and Grubmüller (42). Folding free energy differences were calculated from the difference of the free energies between the reference (unfolded) simulations and the free energies computed for the mutations in the folded protein according to the thermodynamic cycle shown in Fig. 2.

Reference state

The unfolded state of a protein chain is difficult to model, as no single unfolded conformation exists. In previous work (43), unfolded states have been approximated with short peptides that turned out to produce reasonable results. Instead of using the particular sequence around the amino acid of interest, we chose GXG peptides, where X is the amino acid of interest, with capped termini as a reference state. This has the advantage that the reference state of a particular mutation only has to be calculated once, and then, upon calculating all possible mutations, these can be stored and used




FIGURE 2 Thermodynamic cycle. The unfolded state was modeled with capped GXG peptides ($X = \text{any amino acid}$). From the thermodynamic cycle, the folding free energy difference $\Delta\Delta G = \Delta G_3 - \Delta G_2$ between the wild-type protein and the mutant can be calculated via $\Delta G_1 - \Delta G_4$.

as a reference database. Except for the size of the simulation system, which consists of the peptide, a water box with ≈ 1000 water molecules, and 150 mM NaCl, the simulation details were the same as described for the folded state simulations.

RESULTS

Experimentally determined free energy changes were taken from the ProTherm database (44), and contain urea and thermal unfolding data taken from the literature (45–60). The protein stability differences $\Delta\Delta G$ were calculated according to $\Delta G_{\text{unfolded}} - \Delta G_{\text{folded}}$, as shown in the thermodynamic cycle in Fig. 2. Therefore, destabilizing mutations have a negative $\Delta\Delta G$. In Fig. 3, a scatter plot of the experimental values versus the calculated values is shown. As can be seen, a remarkable correlation of 0.86 and an average absolute error of 3.31 kJ/mol are obtained. A quantity of 71.6% of the calculated free energy differences are within ± 1 kcal/mol of the experimental values. Experimental and calculated values are also available in Table S1 of the Supporting Material.

It should be noted that in their assessment of computational methods, Potapov et al. (23) also calculated an experimental-versus-experimental correlation for mutations of which more than one experimental value is available. They obtained a correlation of 0.86 with an average unsigned error of 1.84 kJ/mol.

Potapov et al. (23) furthermore showed that the predictive power of their tested protocols is much better for mutations into Alanine than for mutations into other amino acids. We therefore split our dataset into Alanine mutations and non-Alanine mutations, and assessed the performance difference. Fig. 4 A shows scatter plots for Alanine mutations, non-Alanine mutations, and Glycine-mutations. The results show that the accuracy for Alanine mutations is only slightly better than for non-Alanine mutations (74.2% vs. 70.5% within ± 1 kcal/mol). For mutations into Glycine, the performance drops to 60.9%. However, from the calculated dataset the most destabilizing mutations ($\Delta\Delta G < -20$ kJ/mol) are Glycine mutations. Hence, it remains to be determined whether the method in general performs worse with Glycine

mutations, or with largely destabilizing mutations (which may also result in conformational changes that are not adequately sampled within the simulation time).

In addition, the type of secondary structure does not significantly affect the prediction accuracy. In Fig. 4 B, a discrimination is presented for residues that are located in helices and sheets and for those residues with no secondary assignment according to DSSP (61). The highly destabilizing Glycine mutations, which are systematically predicted to be more destabilizing than experimentally observed, are all located in β -sheets. They therefore contribute significantly to the poor performance of only 59.1% of predictions within ± 1 kcal/mol for β -sheet residues, in contrast to 79% for helical residues and 72% for residues with no DSSP assignment.

A notable dependence on the predictive power of the presented method, however, is observed with respect to the “buriedness” of the mutated residue. In the central picture of Fig. 4 C, the structure of Barnase is shown with color-coded residues. Residues colored in blue are considered well packed according to a packing analysis used in tCONCOORD (62). Residues colored in red are regarded as loosely packed. This is only a rough classification that basically separates the protein into core residues and surface residues. However, it can be seen that the accuracy of calculated free energy changes is significantly better (78.4% within ± 1 kcal/mol) for surface residues than for core residues (65.5%). These findings are most likely caused by the enhanced difficulty of finding the correct side-chain placement within the protein core upon mutation, in comparison to surface residues.

Finally, we assess the dependence of the prediction accuracy on the size of amino acids. Therefore, amino acids are separated into small- and medium-sized types, which are composed of Glycine, Alanine, Serine, Threonine, Valine, Histidine, Asparagine, and Cystine, and big amino acids, to which residues Phenylalanine, Leucine, Isoleucine, Tryptophane, Tyrosine, Glutamine, and Methionine were assigned. Fig. 4 D shows the scatter plots for mutations of small-medium amino acids into other small-medium amino acids,




FIGURE 3 Folding free energy differences for Barnase mutations. (Left) Scatter plot of experimental values versus calculated values. The two thin lines parallel to the diagonal line represent deviations of ± 1 kcal/mol. (Right) Deviation from experimental values.

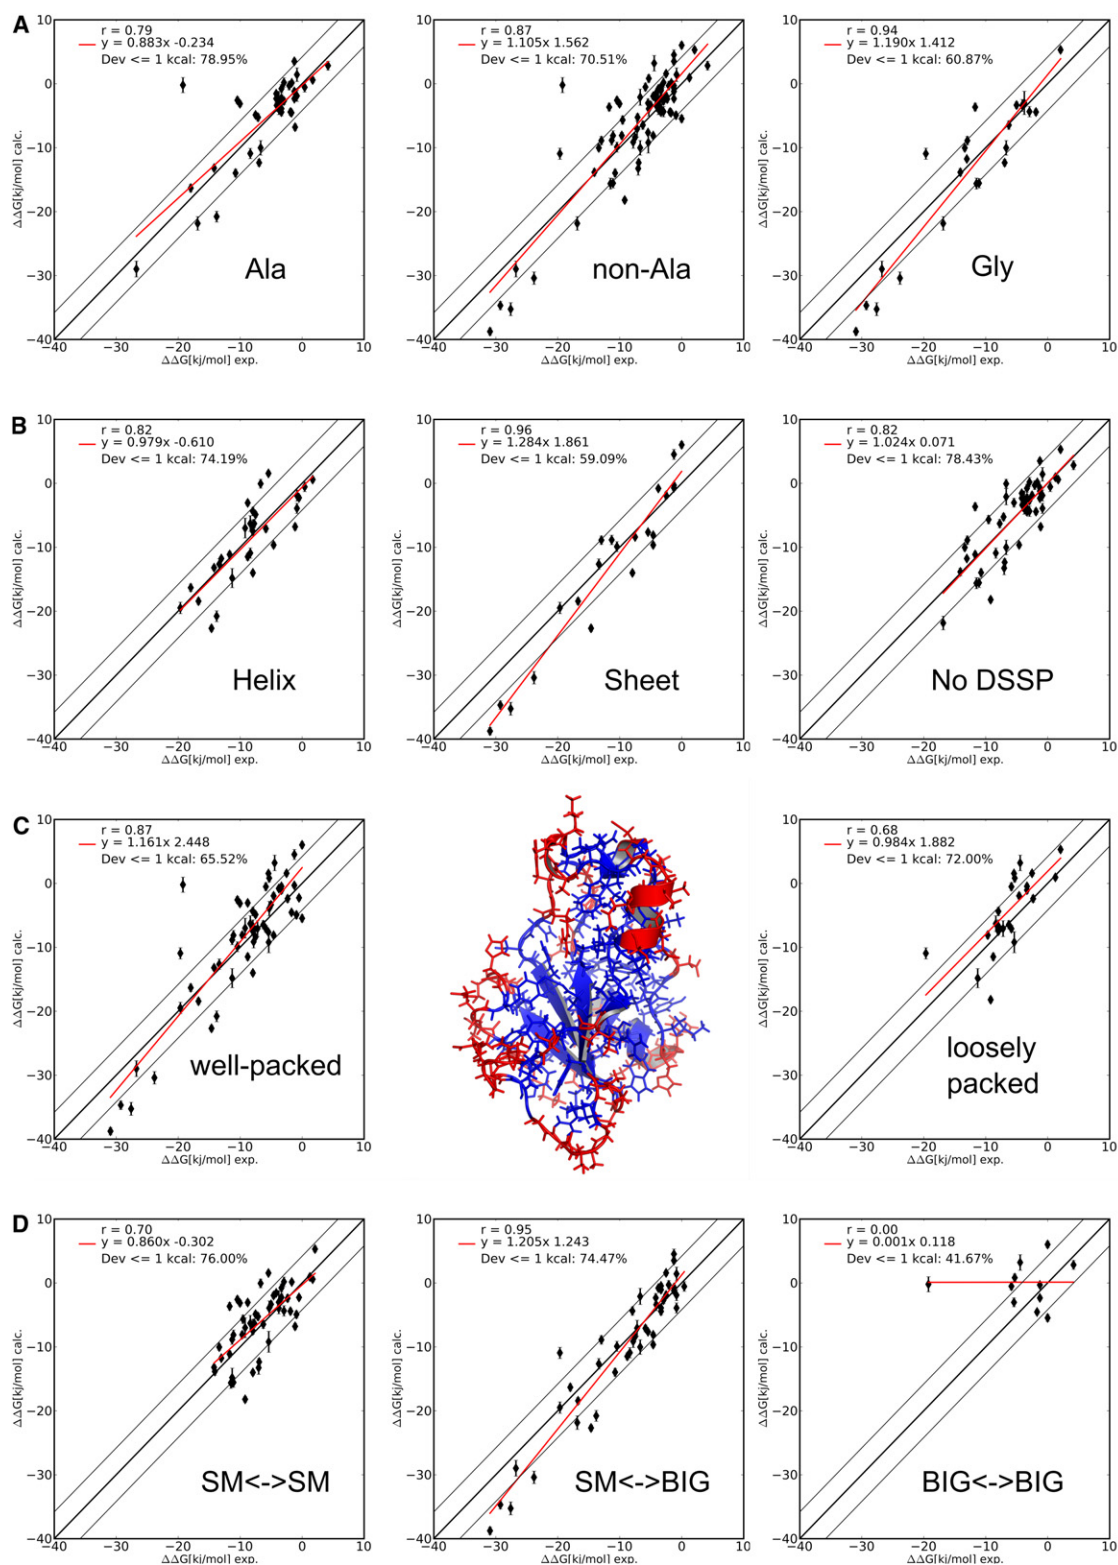


FIGURE 4 Assessment of calculation accuracy. (A) Alanine, non-Alanine, and Glycine mutations; (B) dependence on secondary structure; (C) dependence on packing properties; and (D) dependence on size.

mutations of small-medium amino acids into big amino acids (or vice versa), and the correlation for mutating big amino acids into other big amino acids.

The results show that mutating a big amino acid into another big amino acid is the most challenging case. This is reasonable if one keeps in mind that two big amino acids

are not necessarily similarly shaped. That is, due to its planarity, an aromatic residue can occupy a narrow pocket that is excluded for an aliphatic side chain. Accordingly, a pocket that tightly fits to an aliphatic side chain is not capable of accommodating an aromatic side chain without major rearrangements, and it might not be sampled within the timescales used for the presented protocol.

Because the computational demand of perturbation-based free energy calculations is quite substantial, we further assessed the dependence of the accuracy on the total simulation time. For the data shown above, a total simulation time of 30 ns per mutation was employed (26 ns, if we neglect the 2-ns equilibrations that were not used for the actual calculations). Ten nanoseconds were used for both of the two sampling runs at $\lambda = 0$ and $\lambda = 1$, and $2 \times 100 \times 50$ ps used for the FGTI runs. Fig. 5 shows the accuracy as a function of the total simulation time (the equilibration time was not taken into account). The ratio of sampling time and the number of FGTI runs is thereby kept constant. Hence, a total simulation time of 5 ns corresponds to 1.5-ns sampling for both states and $2 \times 20 \times 50$ ps = 2 ns for the FGTI runs. It is remarkable that with only 20% of the computational effort we already obtain an accuracy of 65% of the predictions within ± 1 kcal/mol. It can furthermore be seen that the accuracy increases with the computational effort, but does not seem to have reached convergence, at a total simulation time of 26 ns per mutation. We therefore expect that more computational effort would result in a further increase of the accuracy.




FIGURE 5 Accuracy versus total simulation time. With only 20% of the computational effort, an accuracy of $>65\%$ of the predictions within ± 1 kcal/mol is obtained. Convergence of the accuracy is apparently not reached at 26 ns per mutation.

To test how the protocol performs with charged mutations, we have added another set of 25 mutations involving a change in the net charge. The results are shown in the left graph of Fig. 6. As can be seen, the accuracy and the correlation with the experimental data are significantly worse than for electrostatically neutral mutations. Only 52% of the calculated values are within ± 1 kcal/mol of the experimental value and six of the 25 values (24%) deviate by >10 kJ/mol. For comparison, for the set of neutral mutations, only two of 109 mutations (1.8%) deviate by >10 kJ/mol from the experimental value. A more detailed investigation shows that the accuracy depends on the location of the residue that is mutated. The right graph of Fig. 6 shows that the deviation of the calculated free energy changes from the experimental value as a function of the relative solvent-accessible surface area of the residue (relative means compared to the solvent-accessible surface area of the residue in an extended chain). As can be seen, free energy changes of mutations at highly solvent-exposed positions are well predicted, whereas mutations of partly buried residues are not.

DISCUSSION

Free energy calculations can be a major step forward in computational protein design. The results shown here indicate that free energy calculations allow the most accurate computational protein stability predictions of mutants among methods available to date. However, the increase in accuracy does not come for free. For each mutation, we performed 30 ns of simulation time for both the folded state and the unfolded reference state. Hence, for the 109 mutations presented here, a total simulation time of $>3 \mu\text{s}$ for both states was required which, even for a small protein like Barnase, requires substantial computational resources. In practical protein design and engineering applications, however, thousands of possible mutations usually need to be scanned. This implies that free energy calculation cannot be expected to replace other methods that are based on fast optimization algorithms (63), rotamer libraries (64), and scoring functions (10,65,66). Despite these limitations, molecular-dynamics-based, free energy calculations can certainly be regarded as valuable extensions in terms of refinement and verification of designed proteins. Moreover, both computational and algorithmic advances render atomistic simulations in the order of microseconds more and more tractable. In addition, problems with limited complexity such as, e.g., Alanine scans of proteins or protein-protein interfaces, can be readily carried out using free energy calculations with more moderate resources.

For 88% of the mutations used for this study, the prediction whether the mutation results in a stabilizing or a destabilizing effect is correct, and in only two cases did the protocol predict a wrong tendency (a stabilizing effect of >3 kJ/mol where the experimental $\Delta\Delta G$ is <-3 kJ/mol). This is

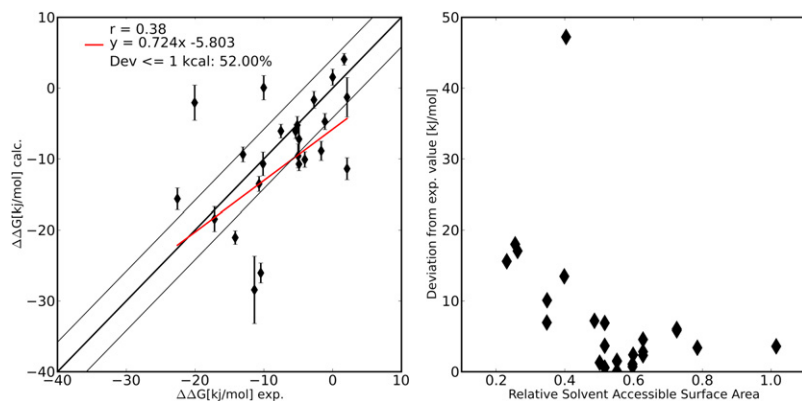


FIGURE 6 Accuracy for mutations with a net charge change. (Left) The accuracy for mutations which result in a change of the net charge of the system is significantly worse than for neutral mutations. Only 52% of the calculated values are within ± 1 kcal/mol of the experimental value, and six of the 25 mutations deviate by >10 kJ/mol (one data point is not shown). (Right) Dependence of the accuracy to the relative solvent-accessible surface area. Mutations at highly solvent-exposed positions are in favorable agreement with experimental data, whereas mutations at partly buried positions are badly predicted.

particularly remarkable if one considers that one of the two states required for the calculation of the free energy differences, the unfolded state, is completely undefined and crudely approximated by a tripeptide that does not even share the sequence with the protein fragment under consideration. These findings open the possibility for the compilation of a database of reference values for all possible mutations. For the mutations considered here, these reference values are available online and in Table S2 and Table S3. Hence, future protein-stability calculations do not require the calculation of both states, but only the calculation of the mutation in the protein of interest.

Mutations involving a charge change should be interpreted with care. As the long-range electrostatic interactions are calculated using particle-mesh Ewald with tin-foil boundaries, an effect equivalent to a uniformly distributed canceling charge is generated to compensate for the net charge due to the charge mutation. For charge mutations at highly solvent-exposed positions on the protein surface, providing a high-dielectric environment for the mutation in both the folded and unfolded (peptide) case we obtain comparable accuracy to neutral mutations. In these cases, the background charge effect performs well due to shielding effects of the high-dielectric environment. For mutations in a more heterogeneous environment as at partly buried positions, we see a distinct loss of accuracy that is most likely due to an incorrect charge distribution.

For free energy calculations aiming at protein/ligand interactions or protein/protein-interfaces, this limitation might be overcome by simultaneously mutating an oppositely charged amino acid far from the interface or binding site. However, for protein stability calculations, this is not applicable, and alternatives need to be developed. Another limitation arises from the requirement of equilibrated ensembles of the native protein and the mutated protein. It is heavily case-dependent whether a 10-ns trajectory (of which snapshots from the last 8 ns were used as starting configurations for the fast-growth TI calculations) will be sufficient to serve as an approximation of a converged ensemble. Analysis of the presented data confirmed the trivial expectation that “longer is better”. However, with only 20% of the

computational effort, we already obtain an accuracy of 65% within ± 1 kcal/mol.

The data furthermore show that a total simulation time of 30 ns per mutation, as employed here, yields a high level of accuracy. By investing more computational power, it is expected to increase accuracy even further. It is evident that special care needs to be taken for those cases involving conformational changes that occur on longer timescales.

The work performed on the system while switching λ from zero to one depends heavily on the conformation from which the FGTI simulation is started. Hence, if the distribution of snapshots taken from the ensemble only covers a subset of the conformational space in equilibrium or the relative weights of different conformations deviate from those in the true ensemble, the resulting work averages from which the free energy differences are calculated will deviate from the true averages. However, as the work values are conformation-dependent, such behavior may be detected when looking at the evolution of calculated work values over the simulation time (of the equilibrium simulation). Conformational changes usually result in a jump, indicating that longer simulation time is required, while oscillation around the mean indicates convergence.

A second aspect that needs to be taken into account is that a modeled mutation does not necessarily represent an equilibrium state. Whether this equilibrium state is reached and adequately sampled within the simulation time depends on the type of mutation and the magnitude of the resulting free energy change. Highly destabilizing mutations, as for some of the Glycine mutations in the presented dataset, may cause changes in both the conformational flexibility of the protein and the mean structure—therefore representing an additional challenge. Furthermore, mutations of well-packed residues are difficult to predict. In our setup we did not pay special attention to favorable positioning of the inserted hybrid residue or to possible rearrangements of neighboring residues. This may easily lead to an unfavorable conformation that probably will not be corrected-for by the simulation within the accessible time. For future improvements of the method, we therefore consider the utilization of rotamer libraries or established methods

such as Rosetta (10) to obtain more realistic initial configurations.

CONCLUSION

We presented the development of a mutant library based on the AMBER99sb force field and a computational framework to carry out free energy calculations in an automated fashion involving naturally occurring amino acids mutations (except for proline). The library was applied to predict the effect of 109 point mutations on the thermodynamic stability in microbial Ribonuclease Barnase. For 71.6% of the mutations, the accuracy of the calculated free energy differences was within ± 1 kcal/mol of the experimental value, and an overall correlation of 0.86 was obtained. We furthermore showed that GXG tripeptides with capped termini serve as sufficient approximations of the unfolded state, enabling the compilation of a reference database containing precomputed values for all possible amino acid mutations.

An initial version of this library is available online. The library is expected to facilitate the setup of free energy calculations for various applications, particularly for rational protein engineering and design.

SUPPORTING MATERIAL

Four tables are available at [http://www.biophysj.org/biophysj/supplemental/S0006-3495\(10\)00216-X](http://www.biophysj.org/biophysj/supplemental/S0006-3495(10)00216-X).

This work was funded by the Deutsche Forschungsgemeinschaft under grant No. GR 207914.

REFERENCES

- Kaplan, J., and W. F. DeGrado. 2004. De novo design of catalytic proteins. *Proc. Natl. Acad. Sci. USA*. 101:11566–11570.
- Jiang, L., E. A. Althoff, ..., D. Baker. 2008. De novo computational design of retro-aldol enzymes. *Science*. 319:1387–1391.
- Röthlisberger, D., O. Khersonsky, ..., D. Baker. 2008. Kemp elimination catalysts by computational enzyme design. *Nature*. 453:190–195.
- Carter, P. J. 2006. Potent antibody therapeutics by design. *Nat. Rev. Immunol.* 6:343–357.
- Chennamsetty, N., V. Voynov, ..., B. L. Trout. 2009. Design of therapeutic proteins with enhanced stability. *Proc. Natl. Acad. Sci. USA*. 106:11937–11942.
- Ditursi, M. K., S. J. Kwon, ..., J. S. Dordick. 2006. Bioinformatics-driven, rational engineering of protein thermostability. *Protein Eng. Des. Sel.* 19:517–524.
- Korkegian, A., M. E. Black, ..., B. L. Stoddard. 2005. Computational thermostabilization of an enzyme. *Science*. 308:857–860.
- Kumar, S., C. J. Tsai, and R. Nussinov. 2000. Factors enhancing protein thermostability. *Protein Eng.* 13:179–191.
- Marshall, S. A., G. A. Lazar, ..., J. R. Desjarlais. 2003. Rational design and engineering of therapeutic proteins. *Drug Discov. Today*. 8:212–221.
- Rohl, C. A., C. E. Strauss, ..., D. Baker. 2004. Protein structure prediction using Rosetta. *Methods Enzymol.* 383:66–93.
- Schymkowitz, J., J. Borg, ..., L. Serrano. 2005. The FoldX web server: an online force field. *Nucleic Acids Res.* 33(Web Server issue): W382–W388.
- Benedix, A., C. M. Becker, ..., R. A. Böckmann. 2009. Predicting free energy changes using structural ensembles. *Nat. Methods*. 6:3–4.
- Pokala, N., and T. M. Handel. 2005. Energy functions for protein design: adjustment with protein-protein complex affinities, models for the unfolded state, and negative design of solubility and specificity. *J. Mol. Biol.* 347:203–227.
- Capriotti, E., P. Fariselli, and R. Casadio. 2005. I-MUTANT2.0: predicting stability changes upon mutation from the protein sequence or structure. *Nucleic Acids Res.* 33(Web Server issue):W306–W310.
- Arnold, F. H. 2001. Combinatorial and computational challenges for biocatalyst design. *Nature*. 409:253–257.
- Chen, C. Y., I. Georgiev, ..., B. R. Donald. 2009. Computational structure-based redesign of enzyme activity. *Proc. Natl. Acad. Sci. USA*. 106:3764–3769.
- Koeller, K. M., and C. H. Wong. 2001. Enzymes for chemical synthesis. *Nature*. 409:232–240.
- Schmid, A., J. S. Dordick, ..., B. Witholt. 2001. Industrial biocatalysis today and tomorrow. *Nature*. 409:258–268.
- Toscano, M., K. Woycechowsky, and D. Hilvert. 2007. Minimalist active-site redesign: teaching old enzymes new tricks. *Angew. Chem.* 46:3212–3236.
- Walsh, C. 2001. Enabling the chemistry of life. *Nature*. 409:226–231.
- Klibanov, A. M. 2001. Improving enzymes by using them in organic solvents. *Nature*. 409:241–246.
- Mattos, C., and D. Ringe. 2001. Proteins in organic solvents. *Curr. Opin. Struct. Biol.* 11:761–764.
- Potapov, V., M. Cohen, and G. Schreiber. 2009. Assessing computational methods for predicting protein stability upon mutation: good on average but not in the details. *Protein Eng. Des. Sel.* 22:553–560.
- Jorgensen, W. 2004. The many roles of computation in drug discovery. *Sci. STKE*. 303:1813–1818.
- Knight, J. L., and C. L. Brooks, 3rd. 2009. Lambda-dynamics free energy simulation methods. *J. Comput. Chem.* 30:1692–1700.
- Jorgensen, W., and C. Ravimohan. 1985. Monte Carlo simulation of differences in free energies of hydration. *J. Chem. Phys.* 83: 3050–3054.
- Beveridge, D. L., and F. M. DiCapua. 1989. Free energy via molecular simulation: applications to chemical and biomolecular systems. *Annu. Rev. Biophys. Biophys. Chem.* 18:431–492.
- Zwanzig, R. W. 1954. High-temperature equation of state by a perturbation method. 1. Nonpolar gases. *J. Chem. Phys.* 22:1420–1426.
- Straatsma, T., and H. Berendsen. 1988. Free energy of ionic hydration: analysis of a thermodynamic integration technique to evaluate free energy differences by molecular dynamics simulations. *J. Chem. Phys.* 89:5876–5886.
- Jarzynski, C. 1997. Nonequilibrium equality for free energy difference. *Phys. Rev. Lett.* 78:2690–2693.
- Jarzynski, C. 1997. Equilibrium free-energy differences from nonequilibrium measurements: a master-equation approach. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics*. 56:5018–5035.
- Crooks, G. 1998. Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems. *J. Stat. Phys.* 90:1481–1487.
- Hornak, V., R. Abel, ..., C. Simmerling. 2006. Comparison of multiple AMBER force fields and development of improved protein backbone parameters. *Proteins: Struct. Funct. Bioinf.* 65:712–725.
- Hess, B., C. Kutzner, ..., E. Lindahl. 2008. GROMACS 4: algorithms for highly efficient, load-balanced, and scalable molecular simulation. *J. Chem. Theory Comput.* 4:435–447.
- Van der Spoel, D., H. J. C. Berendsen, ..., R. Van Drunen. 1995. GROMACS User Manual. <http://md.chem.rug.nl/~gmx>.
- Jorgensen, W. L., J. Chandrasekhar, ..., M. L. Klein. 1983. Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys.* 79:926–935.

37. Parrinello, M., and A. Rahman. 1981. Polymorphic transitions in single crystals: a new molecular dynamics method. *J. Appl. Phys.* 52: 7182–7190.
38. Essmann, U., L. Perera, ..., L. G. Pedersen. 1995. A smooth particle mesh Ewald potential. *J. Chem. Phys.* 103:8577–8592.
39. Miyamoto, S., and P. A. Kollman. 1992. SETTLE: an analytical version of the SHAKE and RATTLE Algorithms for Rigid Water Models. *J. Comput. Chem.* 13:952–962.
40. Hess, B., H. Bekker, ..., J. G. E. M. Fraaije. 1997. LINCS: a linear constraint solver for molecular simulations. *J. Comput. Chem.* 18:1463–1472.
41. Bussi, G., D. Donadio, and M. Parrinello. 2007. Canonical sampling through velocity rescaling. *J. Chem. Phys.* 126:014101.
42. Goette, M., and H. Grubmüller. 2008. Accuracy and convergence of free energy differences calculated from nonequilibrium switching processes. *J. Comput. Chem.* 30:447–456.
43. Pitera, J., P. Kollman, and M. Center. 2000. Exhaustive mutagenesis in silico: multicoordinate free energy calculations on proteins and peptides. *Proteins: Struct. Funct. Bioinf.* 41:385–397.
44. Kumar, M. D., K. A. Bava, ..., A. Sarai. 2006. ProTherm and ProNIT: thermodynamic databases for proteins and protein-nucleic acid interactions. *Nucleic Acids Res.* 34(Database issue):D204–D206.
45. Horovitz, A., J. M. Matthews, and A. R. Fersht. 1992. Alpha-helix stability in proteins. II. Factors that influence stability at an internal position. *J. Mol. Biol.* 227:560–568.
46. Kellis, Jr., J. T., K. Nyberg, ..., A. R. Fersht. 1988. Contribution of hydrophobic interactions to protein stability. *Nature.* 333:784–786.
47. Loewenthal, R., J. Sancho, and A. R. Fersht. 1992. Histidine-aromatic interactions in barnase. Elevation of histidine pK_a and contribution to protein stability. *J. Mol. Biol.* 224:759–770.
48. Matouschek, A., J. T. Kellis, Jr., ..., A. R. Fersht. 1989. Mapping the transition state and pathway of protein folding by protein engineering. *Nature.* 340:122–126.
49. Matouschek, A., J. Matthews, ..., A. Fersht. 1994. Extrapolation to water of kinetic and equilibrium data for the unfolding of barnase in urea solutions. *Protein Eng. Des. Sel.* 7:1089–1095.
50. Oliveberg, M., V. L. Arcus, and A. R. Fersht. 1995. pK_a values of carboxyl groups in the native and denatured states of barnase: the pK_a values of the denatured state are on average 0.4 units lower than those of model compounds. *Biochemistry.* 34:9424–9433.
51. Oliveberg, M., and A. R. Fersht. 1996. New approach to the study of transient protein conformations: the formation of a semiburied salt link in the folding pathway of barnase. *Biochemistry.* 35:6795–6805.
52. Sancho, J., L. Serrano, and A. R. Fersht. 1992. Histidine residues at the N- and C-termini of α -helices: perturbed pK_as and protein stability. *Biochemistry.* 31:2253–2258.
53. Serrano, L., M. Bycroft, and A. R. Fersht. 1991. Aromatic-aromatic interactions and protein stability. Investigation by double-mutant cycles. *J. Mol. Biol.* 218:465–475.
54. Serrano, L., A. G. Day, and A. R. Fersht. 1993. Step-wise mutation of barnase to binase. A procedure for engineering increased stability of proteins and an experimental analysis of the evolution of protein stability. *J. Mol. Biol.* 233:305–312.
55. Serrano, L., and A. R. Fersht. 1989. Capping and α -helix stability. *Nature.* 342:296–299.
56. Serrano, L., A. Horovitz, ..., A. R. Fersht. 1990. Estimating the contribution of engineered surface electrostatic interactions to protein stability by using double-mutant cycles. *Biochemistry.* 29:9343–9352.
57. Serrano, L., J. T. Kellis, Jr., ..., A. R. Fersht. 1992. The folding of an enzyme. II. Substructure of barnase and the contribution of different interactions to protein stability. *J. Mol. Biol.* 224:783–804.
58. Serrano, L., J. Sancho, ..., A. R. Fersht. 1992. Alpha-helix stability in proteins. I. Empirical correlations concerning substitution of side-chains at the N and C-caps and the replacement of alanine by glycine or serine at solvent-exposed surfaces. *J. Mol. Biol.* 227:544–559.
59. Strehlow, K. G., and R. L. Baldwin. 1989. Effect of the substitution Ala—Gly at each of five residue positions in the C-peptide helix. *Biochemistry.* 28:2130–2133.
60. Vu, N. D., H. Feng, and Y. Bai. 2004. The folding pathway of barnase: the rate-limiting transition state and a hidden intermediate under native conditions. *Biochemistry.* 43:3346–3356.
61. Kabsch, W., and C. Sander. 1983. Dictionary of protein secondary structure: pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers.* 22:2577–2637.
62. Seeliger, D., J. Haas, and B. L. de Groot. 2007. Geometry-based sampling of conformational transitions in proteins. *Structure.* 15: 1482–1492.
63. Allen, B. D., and S. L. Mayo. 2006. Dramatic performance enhancements for the FASTER optimization algorithm. *J. Comput. Chem.* 27: 1071–1075.
64. Dunbrack, Jr., R. L., and M. Karplus. 1993. Backbone-dependent rotamer library for proteins. Application to side-chain prediction. *J. Mol. Biol.* 230:543–574.
65. Shen, M. Y., and A. Sali. 2006. Statistical potential for assessment and prediction of protein structures. *Protein Sci.* 15:2507–2524.
66. Yang, Y., and Y. Zhou. 2008. Specific interactions for ab initio folding of protein terminal regions with secondary structures. *Proteins: Struct. Funct. Bioinf.* 72:793–803.