1

Concentration and crowding effects on protein stability from a coarse-grained model

JASON K. CHEUNG, VINCENT K. SHEN, JEFFREY
R. ERRINGTON, AND THOMAS M. TRUSKETT

Introduction

Most of what we know about protein folding comes from experiments on polypeptides in dilute solutions [1–4] or from theoretical models of isolated proteins in either explicit or implicit solvent [5–12]. However, neither biological cells nor protein solutions encountered in biopharmaceutical development generally classify as dilute. Instead, they are concentrated or "crowded" with solutes such as proteins, sugars, salts, DNA, and fatty acids [13–15]. How does this crowding affect native-state protein stability? Are all crowding agents created equal? If not, can generic structural or chemical features forecast their effects?

To investigate these and other related questions with computer simulations requires models rich enough to capture three parts of the folding problem: the intrinsic free energy of folding of a protein in solvent, the main structural features of the native and denatured states, and the connection between protein structure and effective protein–protein interactions. The model must also be simple enough to allow for the efficient simulation of hundreds to thousands of foldable protein molecules in solution, which precludes the use of atomistically detailed descriptions of either the proteins or the solvent.

We recently developed a coarse-grained modeling strategy that satisfies these criteria. It is not optimized to describe any specific protein solution. Rather, it is a general tool for understanding experimental trends regarding how concentration or crowding impact the thermodynamic stability of globular proteins.

2   J. K. Cheung *et al.*

To date, the approach has been used to study how protein concentration affects the folding transition [16], how solution demixing phase transitions (e.g., liquid–liquid phase separation) couple to protein denaturation [17], and how surface anisotropy of the native proteins relates to their unfolding and self-assembly behaviors in solution [18]. In this chapter, we review the modeling strategy and some key insights it has produced.

Coarse-grained modeling strategy

Intrinsic protein stability

A two-state protein molecule in a pure solvent has a temperature- and pressure-dependent thermodynamic preference for either its native (folded) or denatured (unfolded) form [19–24]. The free energy difference between these states ΔG_{f}^0 quantifies the driving force for folding in the absence of protein–protein or protein–solute interactions. It also determines the equilibrium probability $(1+\exp[\Delta G_{\mathrm{f}}^0/k_{\mathrm{B}}T])^{-1}$ associated with observing the native state in an infinitely dilute solution, where k_{B} is the Boltzmann constant and T is the temperature.

Interactions that influence this "intrinsic" stability of the native state include, but are not limited to, intra-protein hydrogen bonding, electrostatics, disulfide bonds, and London–van der Waals interactions, as well as effective forces due to excluded volume, chain conformational entropy, and hydrophobic hydration [25]. Here, we focus exclusively on the last three, since they are relevant not only to protein folding [26–29] but also to other self-assembly processes in aqueous solutions [30–32]. Chain conformational entropy and intra-protein excluded volume interactions favor the more expanded denatured state of a protein, while the ability to bury hydrophobic residues in a largely water-free core favors the compact native fold. Intrinsic stability characterizes how these factors for a protein in the infinitely dilute limit balance at a given temperature and pressure.

The coarse-grained modeling strategy we review here [16] calculates ΔG_{f}^0 under the assumption that a foldable protein can be represented as a collapsible heteropolymer. The effective inter-segment and segment–solvent interactions of the heteropolymer are chosen to qualitatively reflect the aqueous-phase solubilities of the amino acid residues in the protein sequence [33, 34]. One advantage of heteropolymer collapse (HPC) models is that they derive from independently testable principles of polymer physics and hydration thermodynamics. A second advantage is that their behaviors can often be predicted by approximate analytical theories or elementary numerical techniques, which allow them to be efficiently incorporated into multiscale

simulation strategies such as the one discussed here. Although HPC theories are descriptive rather than quantitative in nature, the combination of their simplicity and their ability to reproduce experimental folding trends of globular proteins [34] makes them particularly attractive for use in model calculations.

HPC theories often reflect a balance of structural detail and mathematical complexity [9, 35–37]. In our preliminary studies, we use a basic, physically insightful approach introduced by Dill and co-workers [34, 38]. This theory models each protein of N_r amino acid residues as a heteropolymer of $N_s = N_r/1.4$ hydrophobic and polar segments. As is explained in Ref. [38], the factor of 1.4 enters due to a lattice treatment of the protein in which the chain is partitioned into cubic polymer segments. The amino acids in a globular protein can be represented as occupying cubic volumes with an average edge length of 0.53 nm [38], whereas the separation of α-carbons in an actual protein is about 0.38 nm ($0.53/0.38 \approx 1.4$). The inputs to the theory include temperature T (and, more generally, pH and ionic strength [35]), the number of residues in the protein sequence N_r, the fraction of those residues that are hydrophobic Φ (e.g., based on their aqueous solubilities [33, 34]), and the free energy per unit k_BT associated with hydrating a hydrophobic polymer segment χ. A simple parameterization for χ is available that captures experimental trends for the temperature-dependent partitioning of hydrophobic molecules between a nonpolar condensed phase and liquid water at ambient pressure [16]. Although in this chapter we focus exclusively on thermal effects, we have previously introduced a statistical mechanical method for extending the parameterization for χ to also account for hydrostatic pressure [39].

To compute ΔG_f^0 using this HPC model, one first constructs an imaginary two-step thermodynamic path that reversibly connects the denatured (D) and native (N) states [34]. In step 1, the denatured state with radius of gyration R_D collapses into a randomly condensed configuration with radius of gyration, R_N. The theory assumes that the fraction of solvent-exposed residues that are hydrophobic in both the denatured and the randomly condensed states is Φ, the sequence hydrophobicity of the protein. In step 2, the native state is formed from the randomly condensed state via residue rearrangement at constant radius of gyration, so that the fractional surface hydrophobicity of the protein changes from Φ to Θ. By independently minimizing the free energies of the native and denatured states in this analysis, HPC theory predicts the values of both R_D/R_N and Θ. The intrinsic free energy of folding is the sum of the contributions from the two steps along the imaginary folding path, $\Delta G_f^0 = \Delta G_1^0 + \Delta G_2^0$. Approximate analytical solutions for this HPC theory describe cases where the hydrophobic residues have uniform [34] or patchy [18] spatial distributions on the protein surface. We

4   J. K. Cheung *et al.*

discuss below how these solutions can, in turn, be used to infer approximate nondirectional and directional protein–protein interactions, respectively.

Non-directional protein–protein interactions

While intrinsic thermodynamic stability governs whether an isolated protein favors the native or denatured state, protein–protein interactions play a role in stabilizing or destabilizing the native state at finite protein concentrations.

Protein–protein interactions reflect protein structure. Since HPC theories provide only coarse information about structure, the effects we discuss here are the most basic, generally pertaining to how protein size and surface chemistry couple to their interactions. We first examine the case where proteins display a virtually uniform spatial distribution of solvent-exposed hydrophobic residues, so that protein–protein interactions are, to first approximation, isotropic. We also limit our discussion to systems where the driving force of proteins to desolvate their hydrophobic surface residues by burying them into hydrophobic patches on neighboring proteins dominates the attractive part of the effective protein–protein interaction. The repulsive contribution to the inter-protein potential accounts for the volume that each protein statistically excludes from the centers of mass of other protein molecules in the solution. As should be expected, the structural differences between folded and unfolded protein states translate into distinct native–native NN, native–denatured ND, and denatured–denatured DD protein–protein interactions.

HPC theory correctly predicts that denatured proteins generally exclude more volume to other proteins ($R_D > R_N$) as compared to their native-state counterparts as shown in Fig. 1.1a [34]. Moreover, denatured proteins exhibit a greater fractional surface hydrophobicity than folded molecules ($\Phi > \Theta$). Mean-field approximations [16, 17] predict that the magnitudes of the average "contact" attractions between two isotropic proteins scale as

$$\varepsilon_{\mathrm{ND}} = \frac{N_s \chi(T) \Phi \Theta k_B T}{12} \left(\frac{f_e(\rho_s^*)}{[1 + \rho_s^{*-1/3}]^2} + \frac{f_e(1)}{[1 + \rho_s^{*1/3}]^2} \right) \tag{1.1}$$

$$\varepsilon_{\mathrm{NN}} = \frac{N_s \chi(T) f_e(1) \Theta^2 k_B T}{24} \tag{1.2}$$

$$\varepsilon_{\mathrm{DD}} = \frac{N_s \chi(T) f_e(\rho_s^*) \Phi^2 k_B T}{24} \tag{1.3}$$

where ρ_s^* is the effective polymer segment density, $f_e(\rho_s^*) = 1 - f_i(\rho_s^*)$ is the fraction of residues in the denatured state that are solvent exposed, and $f_i(\rho_s^*) = [1 - (4\pi\rho_s^*/\{3N_s\})^{1/3}]^3$ is the fraction of residues that are on the interior of the protein. A detailed derivation of the above equations can be found in Refs. [16, 18].
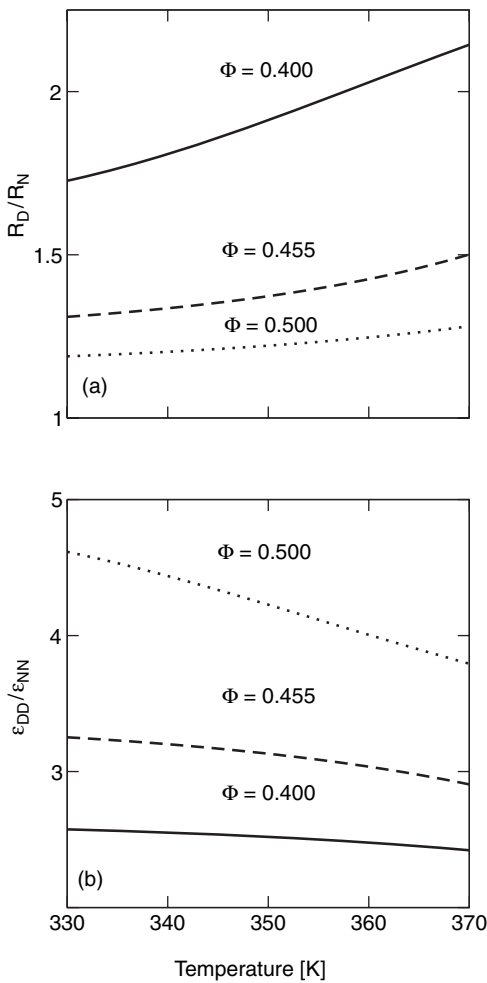
Fig. 1.1 Comparison of the (a) radius of gyration of the denatured protein relative to the native protein, R_D/R_N, and (b) the effective magnitude of the DD attraction relative to the NN attraction, $\varepsilon_{DD}/\varepsilon_{NN}$, for proteins of $N_r = 154$ residues and sequence hydrophobicity $\Phi = 0.400$ (solid), 0.455 (dash), and 0.500 (dot).

Given (1.1–1.3) and $\Phi > \Theta$, it follows that contact attractions involving dena-tured proteins will generally be stronger than those involving the native state (Fig. 1.1b). This is why denaturation often leads to aggregation and precipitation in protein solutions. Along these lines, attractions between proteins increase in strength with the hydrophobic content of the underlying protein sequence Φ.

In our coarse-grained strategy, the interprotein exclusion diameters, $\sigma_{DD}/\sigma_{NN} \approx R_D/R_N$ and $\sigma_{ND}/\sigma_{NN} \approx (R_N + R_D)/2R_N$, and the contact energies of (1.1–1.3), all of which are derived from HPC theory [34], serve as inputs into an effective protein–protein potential $V_{ij}(r)$ that qualitatively captures many

aspects of protein solution thermodynamics and phase behavior (see, e.g., Refs. [40, 41]):

$$V_{ij}(r) = \infty \quad r < \sigma_{ij}$$

$$V_{ij}(r) = \frac{\epsilon_{ij}}{625} \left\{ \frac{1}{[(\frac{r}{\sigma_{ij}})^2 - 1]^6} - \frac{50}{[(\frac{r}{\sigma_{ij}})^2 - 1]^3} \right\} \quad r \geq \sigma_{ij} \tag{1.4}$$

where *ij* corresponds to the type of interaction NN, ND, or DD.

Directional protein–protein interactions

In Dill and co-workers' original development of this HPC theory, they assume that there are no spatial correlations between solvent-exposed hydrophobic residues in either the denatured or the native state [34]. One way to relax this assumption is to allow for segregation of hydrophobic residues on the surface of the native state. For example, consider the hypothetical scenario where two symmetric "patches" form on the surfaces of native proteins during folding. The patches are distinguishable because they have a different hydrophobic residue composition than the rest of the solvent-exposed "body." As shown in Fig. 1.2, the size of each patch is defined by the polar angle α. The fractional patch hydrophobicity Θ_p and body hydrophobicity Θ_b are expressed as

$$\Theta_p = \frac{f_{ph}\Theta}{1 - \cos\alpha}$$

$$\Theta_b = \frac{(1 - f_{ph})\Theta}{\cos\alpha} \tag{1.5}$$
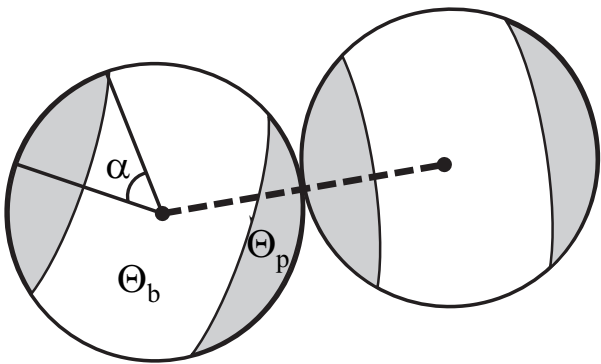
Fig. 1.2 Schematic of two anisotropic native-state proteins. The patch (shaded) and body (white) regions have different hydrophobic residue compositions. The size of the patch is defined by the angle α. The hydrophobicities of the patch Θ_p and body Θ_b are determined by (1.5). Since the dashed line connecting the protein centers passes through a patch region on each molecule, these two proteins are currently in a patch–patch alignment. Adapted from Ref. [18].

where f_{ph} quantifies the fraction of the surface hydrophobic residues that are sequestered into the patch regions on the native protein. Increasing f_{ph} increases Θ_p, which results in higher surface anisotropy and, as we see below, stronger patch–patch attractions. Knowledge of native-state structure would, in principle, allow one to formulate an approximate patch model for a given protein [42], but predicting this structure directly from sequence information using HPC theory is still not generally possible. In other words, f_{ph}, α, and patch location, along with protein sequence, are still knowledge-based inputs for the coarse-grained strategy.

The directional dependencies of the contact attractions of patchy proteins are approximated as follows [18]:

$$\varepsilon_{ND} = \frac{N_s \chi(T) \Phi \Theta_m k_B T}{12} \left(\frac{f_e(\rho_s^*)}{[1 + \rho_s^{*-1/3}]^2} + \frac{f_e(1)}{[1 + \rho_s^{*1/3}]^2} \right) \tag{1.6}$$

$$\varepsilon_{NN} = \frac{N_s \chi(T) f_e(1) \Theta_m \Theta_n k_B T}{24} \tag{1.7}$$

$$\varepsilon_{DD} = \frac{N_s \chi(T) f_e(\rho_s^*) \Phi^2 k_B T}{24}. \tag{1.8}$$

Here, Θ_m and Θ_n denote the apparent surface hydrophobicities associated with different orientational states of interacting native molecules m and n, respectively. For example, to compute the value of Θ_m for molecule m of a given pair interaction, one only needs to know the orientation of molecule m relative to that of the imaginary line connecting its center of mass to that of the other participating protein. If this line passes through a patch on molecule m's surface (see Fig. 1.2), then $\Theta_m = \Theta_p$; otherwise $\Theta_m = \Theta_b$, and so on. Equations (1.6–1.8) reduce to (1.1–1.3) for the isotropic (uniform surface hydrophobicity) case (i.e., $\Theta_p = \Theta_b = \Theta$).

As an illustration, we examine below aqueous solutions of two model proteins of molecular weight $N_s = 110$ (i.e., $N_r = 154$) and hydrophobic residue composition $\Phi = 0.4$, parameters typical for medium-sized, single-domain globular proteins [43]. The difference between the two models is that their native states display distinct surface residue distributions, which in turn lead to different protein–protein interactions: "nondirectional" (i.e., no patches) and "strongly directional" ($f_{ph} = 0.75$, $\alpha = \pi/6$). For simplicity, we refer to these models by their names shown above in quotes, rather than by the f_{ph} and α parameters that define them. As we discuss below, the behaviors of these two model systems provide insights into the mechanisms for stability in several experimental protein solutions.

Figure 1.3 shows the effect of native protein surface anisotropy on the strength of protein–protein attractions. The patch–patch attractions for the

8   J. K. Cheung *et al.*
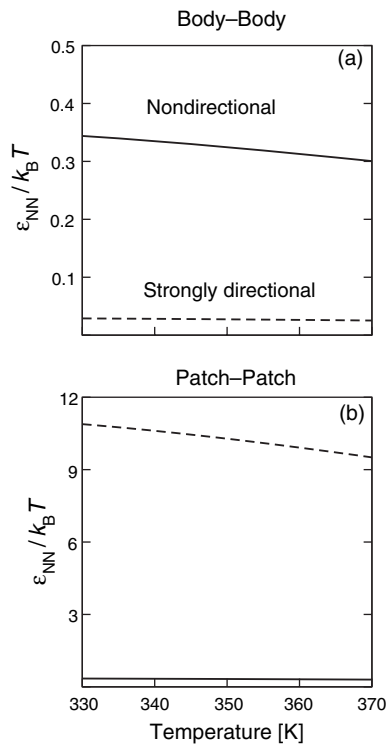
Fig. 1.3 Comparison of the contact attraction ε_{NN} relative to $k_B T$ for (a) body–body alignment and (b) patch–patch alignment of $N_r = 154$, $\Phi = 0.4$ proteins with strongly directional interactions (dash). The contact attraction for $N_r = 154$, $\Phi = 0.4$ native proteins with nondirectional interactions (solid) is also shown.

strongly directional protein are more than an order of magnitude larger than the other attractions. Pairs of directional proteins can desolvate a higher number of hydrophobic residues by self-associating (as compared to the nondirectional proteins), but only if they do so with their hydrophobic patches mutually aligned, which in turn imposes an entropic penalty. This balance between favorable hydrophobic interactions and unfavorable entropy yields the possibility of continuous equilibrium self-assembly transitions involving the native state [18]. Given the symmetric patch geometry of the native state model studied here, the morphology of the self-assembled "clusters" would resemble linear polymeric chains.

The interactions discussed above are similar in spirit to a "two-patch" description that was recently introduced to model the native–native protein interactions of the sickle cell variant of hemoglobin [44] and also to other semi-empirical anisotropic potentials developed for native proteins [42, 45–50].

However, the coarse-grained strategy described here differs significantly from these earlier models in two ways: it explicitly accounts for the possibility of protein denaturation, and it estimates the intrinsic properties of the native and denatured states using a statistical mechanical theory for heteropolymer collapse. This link to the polymeric aspect of the protein is crucial because it allows our model to be used as a tool to investigate how native-state protein anisotropy affects folding equilibria, self-assembly, and the global phase behavior of protein solutions.

Reducing protein stability to a classic chemical engineering problem

It is worth emphasizing that the coarse-grained model described above represents an effective binary mixture of folded and unfolded proteins (the aqueous solvent only entering through χ) connected via the protein folding "reaction." Links between the intrinsic native-state stability of the proteins, ΔG_{f}^0, the physical parameters defining the protein sequence (N_{r}, Φ), the native-state surface morphology ($\Theta, \alpha, f_{\mathrm{ph}}$), the interactions of hydrophobic residues with aqueous solvent χ, and the protein–protein interactions ($\epsilon_{ij}, \sigma_{ij}$) are established through the HPC model [16, 17]. As in experimental protein solutions, the fraction of proteins in the native state generally depends on both temperature and protein concentration. This fact, often neglected in other modeling approaches which ignore the polymeric nature of proteins or protein–protein interactions, arises because temperature affects the intrinsic stability of the native state ΔG_{f}^0, and both temperature and protein concentration influence the local structural and energetic environments that native and denatured proteins sample in solution.

In short, the coarse-grained approach frames the stability of concentrated protein solutions in terms of a classic chemical engineering problem: mapping out the equilibrium states of a reactive, phase-separating mixture [51]. In the next section of this chapter, we discuss how advanced Monte Carlo methods designed to efficiently solve the latter problem can also be used to address the former.

Simulation methods

The properties of the coarse-grained protein model described above can be readily studied using transition-matrix Monte Carlo (TMMC) simulation. TMMC is a relatively new simulation technique that has emerged in recent years as a highly efficient method for investigating the thermophysical properties of fluids. It is useful for a variety of applications ranging from the precise calculation of thermodynamic properties of pure and multicomponent fluids in bulk

and confinement [52–62], surface tension of pure fluids and mixtures [52, 63–65] and Henry's constants [66] to the investigation of wetting transitions [67, 68] and adsorption isotherms [69, 70]. The basic goal of TMMC is to calculate the free energy of a system along some order parameter path by determining the order parameter probability distribution. To do this in a conventional simulation, a histogram is constructed by simply counting the number of times the system takes on, or visits, a given order parameter value. In a transition-matrix-based approach, the distribution is determined by accumulating the transition probabilities of the system moving from one order parameter value to another during the course of a Monte Carlo simulation, and subsequently applying a detailed balance condition over the explored region of order parameter space. To facilitate sufficient and uniform sampling of order parameter space, a self-adaptive biasing scheme is often employed. The reader is referred to Refs. [52, 56–58, 63, 71–75] for further details. For this coarse-grained protein model, we use a particular TMMC implementation designed for multicomponent systems [57].

To obtain thermodynamic and structural quantities of interest, we perform TMMC simulations within the grand-canonical ensemble. Under these conditions (fixed chemical potentials, volume, and temperature), an appropriate choice of order parameter is the total number of molecules in the system. While the order-parameter probability distribution is unique to the chemical potentials used in the simulation, histogram reweighting can be used to determine the distribution at other chemical potentials [76]. Because the coarse-grained protein model can be regarded as a binary mixture of native and denatured proteins where the components can undergo a unimolecular chemical reaction (folding), chemical equilibrium requires that the chemical potentials of the native and denatured proteins be identical. Thus, only a single chemical potential needs to be specified in a TMMC simulation of the protein solution. The free-energy change of the reaction, that is the intrinsic free energy of folding ΔG_f^0, enters as an activity difference between the folded and unfolded proteins, a treatment which implicitly assumes that the protein's intermolecular and intramolecular degrees of freedom are separable.

Although the main output of a TMMC simulation is the total protein number (concentration) distribution, other system properties can be also calculated. This is done in a straightforward way by collecting isochoric averages during the course of a simulation. Combining these statistics with the above-mentioned histogram-reweighting technique, the fluid-phase properties of the coarse-grained protein model can be determined over a wide range of concentrations from a single TMMC simulation. In our studies, we calculate, along with other properties, the following quantities: the overall fraction of folded proteins f_N, the fraction of clustered or aggregated proteins f_{clust}, and the average fraction